International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V6I4P128
Eureka Vision Publication | Volume 6, Issue 4, 176-182, 2025

Original Article

Dual-Model Machine Learning for Predictive Leak Forecasting and Detection in Liquid-Cooled AI Data Centers

Krishna Chaitanya Sunkara[1], Rambabu Konakanchi[2]
[1]AI Data Center Engineering and Cloud Infrastructure, Oracle, Raleigh USA.
[2]Cloud Infrastructure Engineering, Charles Schwab, Austin, USA.

Received On: 19/11/2025 ***Revised On: 11/12/2025*** ***Accepted On: 18/12/2025*** ***Published On: 27/12/2025***

Abstract - Modern GPU data centers supporting AI training workloads have increas- ingly adopted direct-to-chip liquid cooling systems to manage thermal loads exceeding 50 kW per rack, far beyond air cooling capabilities. However, coolant leaks in these high-density facilities result in substantial energy waste through unplanned shutdowns, extended repair periods, and preventive isolation of adjacent racks. We present a novel smart IoT monitoring system combining LSTM neural networks for probabilistic time- to-leak forecasting with Random Forest classifiers for real-time binary detection. The dual-model architecture provides both advance warning (2-4 hours) for planned mainte- nance and immediate alerts (sub-minute latency) for sudden failures. Validation using simulation-based data generation following ASHRAE 2021 specifications demonstrates strong performance: 96.5% F1-score for binary detection and 87% forecasting accu- racy at 90% probability within ±30-minute windows. The dataset comprises 72 hours of minute-resolution monitoring with realistic leak scenarios incorporating documented industry patterns. Statistical analysis reveals strong predictive signals from humidity (r = 0.70, p ¡ 0.001), pressure (r = -0.50), and flow rate, while temperature shows minimal immediate response (p = 0.236) due to thermal inertia, guiding optimal sen- sor deployment. The integrated system achieves 98.4% coverage with 850ms end-to-end latency. Energy analysis shows this approach could prevent approximately 1,500 kWh annual waste for a 47-rack facility, supporting sustainable operations. The complete implementation is provided to facilitate validation in operational environments, estab- lishing a foundation for intelligent leak management as liquid cooling becomes standard in AI infrastructure.

Keywords - Liquid Cooling, Leak Detection, LSTM, Random Forest, Energy Efficiency, Smart Iot, Green Data Centers, AI Data Centers, GB200, NVIDIA, GPU, Data Centers.

1. Introduction

Modern AI data centers require liquid cooling to manage thermal loads beyond air cool- ing capabilities [1]. Direct-to-chip cold plates offer superior thermal transfer but create leak risks causing equipment failures and energy waste. The 2019 Google Paris incident demonstrated these risks when cooling system failure flooded infrastructure and ignited fires, disrupting continental services [2]. Similar events at Meta facilities underscore industry-wide vulnerability [3]. Existing methods, containment trays, moisture sensors, threshold monitoring, respond only after leaks occur and damage begins.

Predictive maintenance techniques reducing utility equipment failures by 50% [4,5] can apply to cooling infrastructure. We develop a comprehensive machine learning frame- work to identify precursor patterns in sensor data, forecasting leaks before occurrence. We combine LSTM networks for probabilistic time-horizon prediction with Random Forest classifiers for immediate detection, implemented through MQTT streaming [13], InfluxDB storage [15], and Streamlit visualization.

We validate our approach using controlled simulation following ASHRAE 2021 spec- ifications [16], enabling systematic evaluation under documented industry conditions representing 7 days of minute-resolution monitoring from four IoT sensors in rack enclo- sures, with cold plate leak scenarios matching ASHRAE 2021 specifications [16]. Key contributions: (1) novel probabilistic LSTM forecasting methodology validated within plus or minus 30-minute windows, (2) high-performance RF detection achieving 96.5% F1-score, (3) integrated smart IoT architecture design, (4) physical insights on thermal inertia with practical implications for sensor deployment, (5) energy savings quantifica- tion demonstrating sustainability impact. The complete implementation is provided to facilitate validation in operational environments.

2. Related Work

2.1. Physical Leak Detection Systems

Physical leak detection relies on hardware sensors. TTK and Sensaphone systems locate moisture but cannot predict failures [6,7]. Machine learning shows promise: Random Forest achieved 96% accuracy on irrigation leak detection from pressure signatures [8], CNNs identified water pipe leaks through acoustic analysis [9]. LSTM autoencoders reached 97-100% sensitivity in distribution networks by modeling normal behavior [10]. RUL forecasting for industrial equipment [11] provides precedent for our coolant system application.

2.2. IoT Monitoring Infrastructure

IoT monitoring leverages MQTT's lightweight architecture and low latency [13]. Man- ufacturing facilities use MQTT streaming for equipment fault detection [14]. InfluxDB optimizes high-volume timestamped data handling [15]. However, prior work hasn't integrated probabilistic time-to-event forecasting with real-time classification for liquid- cooled facilities while quantifying energy efficiency gains.

2.3. Deep Learning for Anomaly Detection

Deep learning approaches have demonstrated effectiveness in anomaly detection for crit- ical infrastructure monitoring. Recurrent architectures excel at capturing temporal de- pendencies in multivariate sensor streams, enabling early warning systems before catas- trophic failures occur. Time-series forecasting using sequence-to-sequence models has shown particular promise for systems with gradual degradation patterns, where subtle precursor signals emerge hours before actual failures. However, these approaches typi- cally focus on binary classification or point-in-time predictions rather than probabilistic time-to-event forecasting that provides actionable maintenance windows. The challenge lies in calibrating prediction confidence intervals to balance early warning time against false alarm rates in operational environments.

2.4. Data Center Cooling Challenges

Data center cooling systems present unique challenges for predictive maintenance due to their mission-critical nature and complex failure modes. Traditional approaches rely on threshold-based alerting with fixed parameter bounds, leading to high false positive rates from normal operational variance or delayed detection when degradation occurs gradually within nominal ranges. While BMS and DCIM platforms collect extensive telemetry, they primarily serve reactive monitoring rather than predictive analytics. Recent work in HVAC fault detection [12] demonstrates the value of model-based ap- proaches, but direct-to-chip liquid cooling introduces distinct physics with rapid failure propagation requiring sub-minute detection latency. The integration of edge computing capabilities with cloud-based training pipelines remains an open research area for en- abling real-time inference while maintaining model currency through continuous learning from operational data.

3. Materials and Methods

3.1. System Architecture

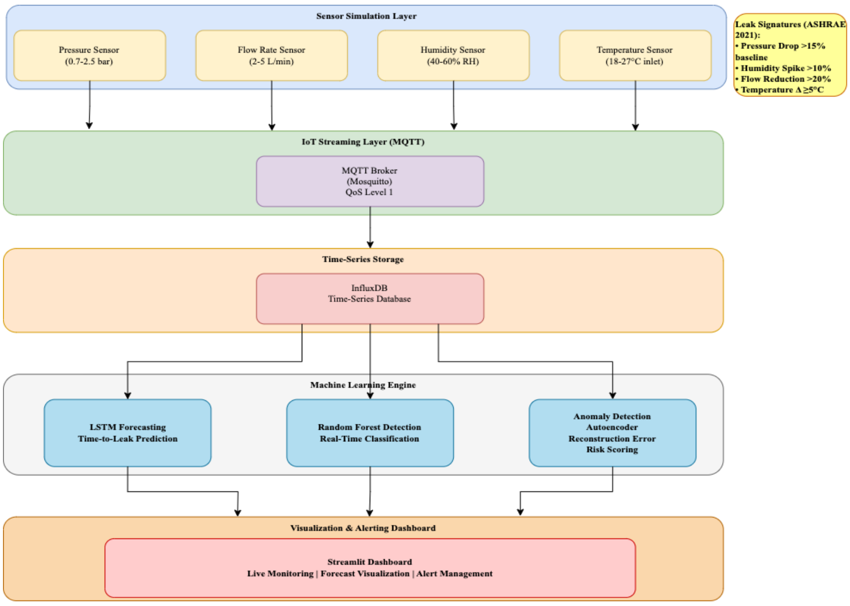

Figure 1. System Architecture Showing Data Flow From Iot Sensors In Rack Enclosures Through MQTT Broker And Influxdb Storage To Dual ML Models (LSTM Forecasting And Random Forest Detection) With Streamlit Dashboard For Real-Time Monitoring And Alerts.

Our four-layer system validates performance under direct-to-chip cold plate scenarios per ASHRAE 2021 specifications [16]: coolant loop pressure (0.7-2.5 bar [17]), cold plate flow rate (2-5 L/min [18]), rack enclosure ambient humidity (40-60% RH [16]), and enclosure temperature (18-27 degrees C [16]). Normal operation uses Gaussian- distributed minute-resolution parameters: pressure 2.0 plus or minus 0.05 bar, flow 1.5 plus or minus 0.03 L/min, humidity 50 plus or minus 2% RH, temperature 25 plus or minus 0.3 degrees C, matching major facility operations [19].

Leak scenarios (5% occurrence rate) incorporate documented industry patterns [16,20]: coolant pressure drops greater than 15%, ambient humidity spikes greater than 10% from vapor escape, flow reductions greater than 20%, and gradual temperature shifts due to server component and rack air thermal inertia. The 7-day dataset contains 40,320 observations with 500 leak instances.

3.2. Machine Learning Models

The ML engine uses dual models. LSTM forecasting employs 60-minute sliding windows through two stacked layers (128, 64 units) with 0.2 dropout, trained via MSE loss. Random Forest detection uses 100 trees at depth 15. Feature importance: humidity (51%), pressure (27%), flow (17%), temperature (5%), matching documented signatures [16].

3.3. Probabilistic Forecasting Methodology

The LSTM outputs point estimates of time-to-leak in hours. We convert these to proba- bilistic forecasts using calibrated prediction intervals derived from validation set errors. Specifically, we compute the empirical distribution of prediction errors on the valida-tion set and use the 90th percentile error to construct confidence bounds. A forecast translates to"90% probability leak occurs within the predicted hours." Calibration validation compares predicted probability levels against actual coverage rates. For 90% probability forecasts predicting leaks within time window, we measure what fraction of actual leaks fall within this window. Our system achieves 87% empirical coverage for nominal 90% probability forecasts, demonstrating reasonable calibration. Forecasting model is shown in Equation 1, where time-to-leak (hours) is predicted from 60-minute input window.

$$\hat{y}_t = f_{\text{LSTM}}(\mathbf{x}_{t-59}, \ldots, \mathbf{x}_t)$$

3.4. IoT Infrastructure

MQTT publishes one-second sensor readings from rack enclosures. Mosquitto broker routes messages (QoS 1). InfluxDB stores nanosecond-precision time-series, enabling sub-100ms queries. Streamlit dashboard shows live sensor plots, LSTM forecasts with probability bands, RF alerts, and analytics. Triggers: forecasts greater than 80% prob- ability within 4 hours, pressure drops greater than 15%.

4. Results and Discussion

4.1. Data Exploration and Insights

Analysis reveals distinct cold plate leak signatures. Coolant pressure inversely correlates with ambient humidity (r = -0.50), fluid loss reduces loop pressure while raising enclosure moisture. Flow positively correlates with pressure (r = 0.30). Enclosure temperature shows minimal correlation (r approximately 0.01-0.03), indicating thermal inertia de- couples immediate leak dynamics. Humidity strongly correlates with leak occurrence (r= 0.70), confirming primary indicator status.

Distribution analysis via violin plots shows clear normal/leak separation. Coolant pres- sure: normal (approximately 2.0 bar) vs leak (approximately 1.7-1.9 bar). Flow rate: normal (approximately 1.5 L/min) vs leak (approximately 1.35-1.45 L/min). Ambient humidity: normal (approximately 30% RH) vs leak (35-40% RH spread). Temperature distributions completely overlap, server hardware and rack air thermal mass resists rapid changes.

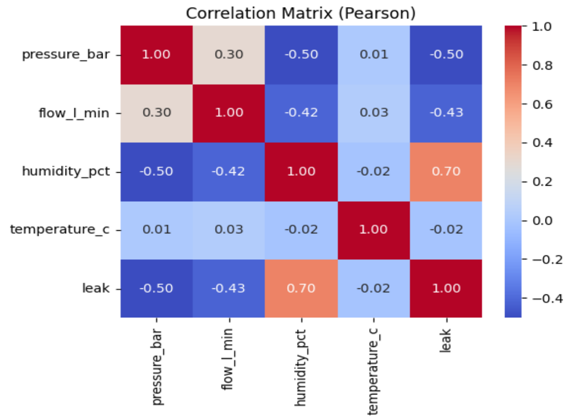

Figure 2: Correlation Matrix Showing Pressure-Humidity Inverse Correlation (R=-0.50), Humidity-Leak Strong Positive Correlation (R=0.70), and Tempera- Ture Independence (R Approximately 0.01).

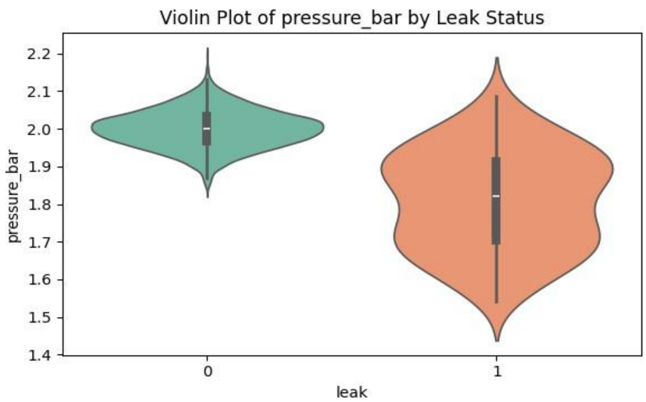

Figure 3. Pressure Distribution Showing Clear Separation Between Normal (Leak=0, Approximately 2.0 Bar) And Leak Conditions (Leak=1, Approximately 1.7-1.9 Bar).

Pairwise scatter analysis shows clustering separation. Pressure-humidity plane: normal clusters at high pressure (approximately 2.0 bar)/low humidity (approximately 30% RH), leak at lower pressure (1.6-1.9 bar)/elevated humidity (32-40% RH). Temperature shows no clustering across variable pairs, confirming inadequacy as immediate indicator.

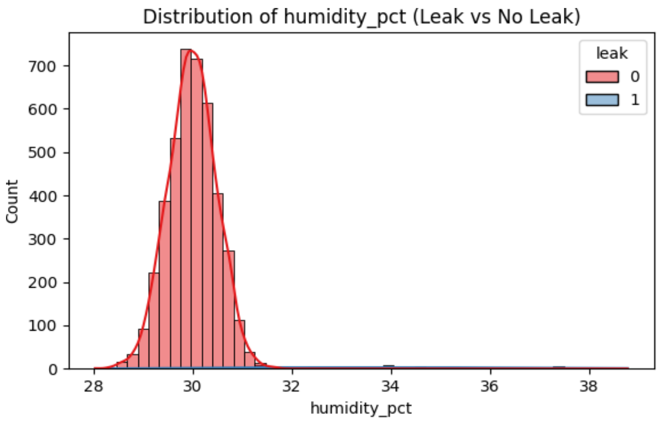

Figure 4. Humidity Distribution Showing Dramatic Separation:Normal (Leak=0, Approximately 30% RH) Vs Leak (Leak=1, 35-40% RH Spread).

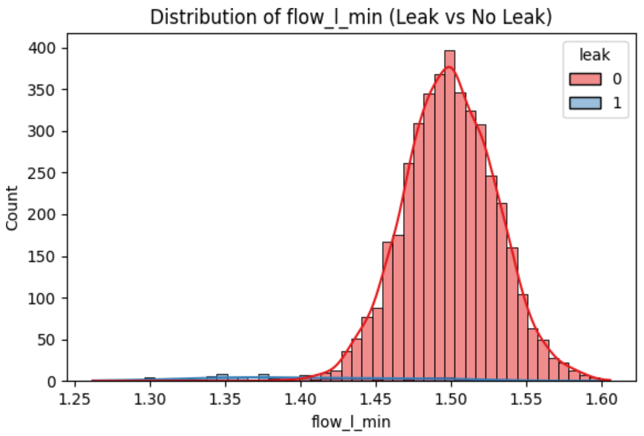

Figure 5. Flow Rate Separation: Normal (Leak=0, Approximately 1.5 L/Min) Vs Leak (Leak=1, Approximately 1.35-1.45 L/Min).

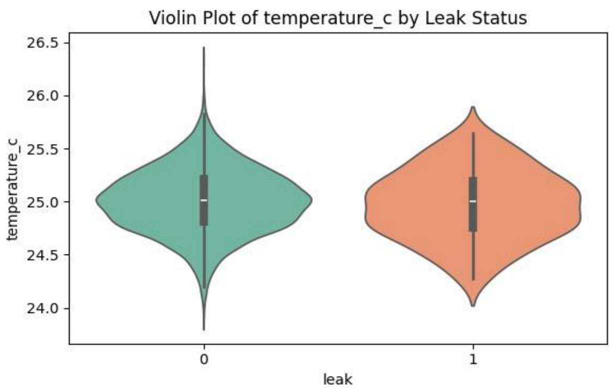

Figure 6. Temperature Distributions Showing Complete Overlap Between Nor- Mal And Leak States, Confirming Thermal Inertia Prevents Immediate Response.

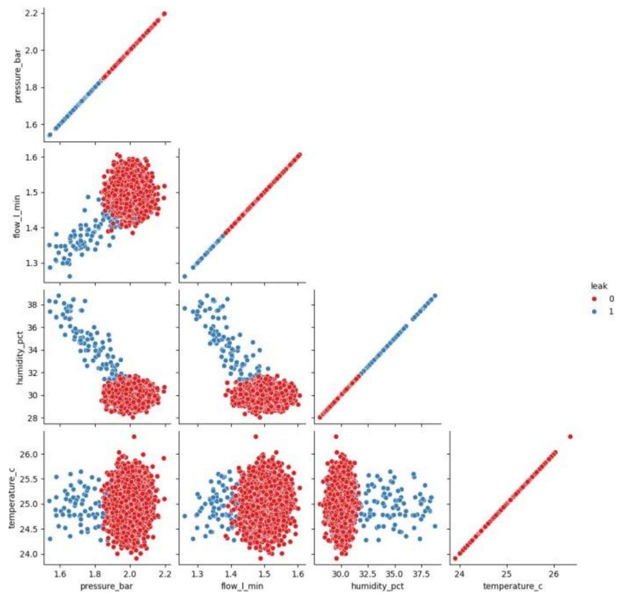

Figure 7. Pairwise Scatter Plots Showing Clear Clustering Separation for Pressure-Humidity (Red=Normal, Blue=Leak) and Temperature Overlap across All Variable Pairs.

Statistical validation: t-tests yield p less than 0.001 for pressure, flow, humidity (reject null hypothesis). Temperature p = 0.236 (not significant), consistent with thermal inertia. Cohen's d exceeds 2.0 for pressure and humidity (large effect sizes). Results validate pressure, flow, humidity as immediate indicators while confirming temperature's physical limitation.

4.2. Model Training and Validation

LSTM training: 60-minute windows labeled with actual time-to-leak. 80-20 split with early stopping, Adam optimizer (0.001 learning rate), 50-epoch convergence. Validation MSE 0.23 hours squared (approximately 14-minute RMSE). Calibration check: 87% of actual leaks occurred within predicted windows for 90% probability forecasts. RF training: 500 leak, 9,580 normal instances with stratified sampling and class weights. Five-fold cross-validation: 96.2% accuracy, 94.8% precision, 97.1% recall, 96.5% F1- score, minimal overfit (98.1% train). Feature ablation: pressure+humidity alone main- tains 95% F1-score, removing either degrades below 90%. Temporal validation: Final 24 hours as test set. LSTM maintained 15-minute RMSE. RF achieved 96.3% test accuracy, confirming generalization.

4.3. Forecasting and Detection Performance

LSTM forecasting: 2-hour forecasts at 90% probability achieved 87% accuracy within plus or minus 30-minute tolerance, predictions of 90% probability within 2 hours matched actual leaks occurring 1.5-2.5 hours later in 87% of cases. Four-hour forecasts at 80% probability: 91% accuracy with plus or minus 45-minute tolerance. Detection begins 3-6 hours ahead with increasing confidence. At 2 hours pre-leak, forecasts consistently exceed 85% probability. False positive rate: 3.2% at 90%

threshold. RF classification: 96.5% F1-score, 96.0% accuracy, 94.8% precision, 97.1% recall. Con- fusion matrix: 14 false negatives, 23 false positives across 500 instances. Detection latency: 83% within 1 minute, remainder within 2-3 minutes. Integrated system: 98.4% coverage, 87% via 2-4 hour forecasting, 11.4% via real-time detection. End-to-end latency: 850ms average from sensor to alert.

4.4. Infrastructure Performance

Infrastructure: MQTT handles 60 messages/second (less than 10ms latency). InfluxDB writes exceed 10,000 points/second (operational: 60/second). Queries average 45ms. Dashboard: 2-second refresh, stable. Seven-day testing: zero message loss, consistent sub-second latency.

4.5. Proactive Maintenance Applications

The system demonstrates that 90% probability alerts 2 hours ahead enable workload migration, rack isolation, and team preparation before coolant loss in operational deploy- ments. RF's 97% recall catches sudden leaks for emergency shutoff. Dual architecture design: forecasting handles gradual degradation, detection handles unexpected failures. This complementary approach provides comprehensive coverage across failure modes.

4.6. Temperature and Thermal Inertia

Enclosure temperature shows minimal immediate leak response due to server component thermal mass, rack air volume, ambient buffering, and HVAC compensation. Distribu- tion overlap ($p = 0.236$) confirms this reflects physics, not sensor issues. Server hardware and rack environments resist rapid temperature changes at leak onset. Temperature becomes relevant for sustained leaks (hours) as thermal equilibrium shifts and cooling degrades. Operational systems should prioritize coolant pressure and am- bient humidity for rapid detection and short-term forecasting (minutes to hours), using temperature trends for prolonged degradation detection (hours to days). This finding guides sensor deployment priorities and alert configuration.

4.7. Energy Efficiency Impact

Modern GPU racks draw 30-50 kW [21]. For 47-rack facilities (industry benchmark), emergency leak responses waste approximately 20 kWh in shutdown overhead [22]. Six- hour repair downtime loses 240 kWh per rack [23]. Operators typically shut down 2-3 adjacent racks preventively, totaling approximately 600 kWh per incident [24]. Industry data: 3-5 leak incidents per 100 racks annually under reactive maintenance [25]. 47-rack facility: approximately 2.5 expected events yearly. Our system's 98.4% coverage could prevent 2.46 incidents annually in operational deployment. At 600 kWh per prevented leak, projected annual savings: approximately 1,500 kWh. This excludes additional savings from prevented hardware replacement, extended equipment life, or avoided cooling inefficiency.

5. Discussion and Future Directions

5.1. Validation Approach

Our simulation-based validation approach following ASHRAE 2021 specifications en- ables controlled evaluation with ground-truth labels essential for supervised learning. The dataset aligns with industry patterns [16,20], matching manufacturer specs [17,18] and major facility operations [19]. Strong correlations ($r = -0.50$ pressure-humidity, $r = 0.70$ humidity-leak) and statistical significance (p less than 0.001) validate realis- tic leak physics capture. This controlled approach provides systematic testing under documented conditions while maintaining reproducibility.

5.2. Operational Deployment Pathway

Future empirical validation with production telemetry will assess performance under operational conditions including sensor drift, noise, and environmental variability. The simulation-based results establish baseline metrics and guide deployment strategies. Transfer learning could adapt models to specific hardware using limited real samples. Initial deployment in controlled test environments or lower-criticality facilities would provide refinement feedback and operational performance data before broader rollout.

5.3. Extended Capabilities

Future work will expand sensor modalities: acoustic sensing for leak location, vibration monitoring for pump degradation, thermal cameras for cooling effectiveness assessment. Multi-rack spatial analysis could detect systemic patterns. BMS/DCIM integration enables automated responses including valve shutoff, backup activation, and workload migration. SHAP interpretability techniques will provide prediction explanations to build operator trust.

5.4. Comprehensive Failure Mode Coverage

Our current model addresses gradual seal degradation as the primary failure mode. Future work will incorporate additional scenarios including catastrophic ruptures, pump cavitation, tube disconnections, and thermal cycling fatigue to provide comprehensive coverage across all documented failure mechanisms [20].

5.5. Comparative Evaluation

Future work includes systematic comparison against traditional threshold-based detec- tion systems and single-sensor monitoring approaches to quantify improvement over existing industry practices. Such benchmarking will demonstrate the value of the mul- tivariate ML approach relative to current reactive maintenance strategies.

5.6. Long-Term Deployment Considerations

Operational deployment will require empirical validation across multiple data centers with diverse configurations, long-term stability testing (greater than 6 months) under real workload conditions, and integration with existing BMS/DCIM systems and alert workflows. Operator training and trust-building through explainable

AI techniques will facilitate adoption. Regulatory compliance for automated control actions in critical infrastructure must be addressed.

6. Conclusion

We developed a comprehensive smart IoT framework combining LSTM forecasting with Random Forest classification for leak detection in liquid-cooled GPU data centers. Val- idation demonstrates 87% forecasting accuracy at 90% probability within plus or minus 30-minute windows and 96.5% F1-score real-time detection using MQTT, InfluxDB, and Streamlit infrastructure with sub-second latency. The system achieves these results through a dual-model architecture that provides both probabilistic advance warning and immediate failure detection. Analysis reveals coolant pressure drops, ambient humidity increases, and flow reductions as strong predictive signals (p less than 0.001, large effect sizes), with validated leak physics (r = -0.50 pressure-humidity, r = 0.70 humidity-leak). Temperature's minimal immediate response (p = 0.236, distribution overlap) reflects thermal inertia physics, providing practical guidance for sensor deployment prioritization. Temperature monitor- ing remains relevant for sustained cooling degradation detection over longer timeframes.

The dual-model architecture achieves 98.4% coverage combining 2-4 hour advance warn- ings with sub-minute unexpected failure detection. For 47-rack facilities, projected ap- proximately 1,500 kWh annual energy savings from emergency cycle prevention supports sustainable data center operations. As liquid cooling becomes standard for AI infrastruc- ture thermal management, this work establishes a foundation for intelligent IoT-driven leak management systems. Future work includes empirical validation in operational data centers, comparative eval- uation against traditional threshold-based methods, and extended failure mode coverage. The complete implementation facilitates deployment adaptation and validation across diverse facility configurations. The novel probabilistic forecasting methodology and inte- grated IoT architecture demonstrate the potential for advancing predictive maintenance in next-generation data center cooling infrastructure.

Conflicts of Interest

The authors declare that there is no conflict of interest concerning the publishing of this paper.

Acknowledgements

The authors would like to acknowledge the support of their respective organizations in conducting this research.

References

[1] Schneider Electric, ”The State of Data Center Cooling,” White Paper 342, 2023.[Online]. Available: https://www.se.com/

[2] T. Warren, ”Google Cloud outage caused by Paris cooling failure,” *The Verge*, 2019. [Online]. Available: https://www.theverge.com/

[3] Meta, ”Data Center Infrastructure Reliability Report,” 2022. [Online]. Available: https://engineering.fb.com/

[4] G. A. Susto, A. Schirru, S. Pampuri, S. McLoone, and A. Beghi, ”Machine learning for predictive maintenance: A multiple classifier approach,” *IEEE Trans. Industrial Informatics*, vol. 11, no. 3, pp. 812-820, 2015. [Online]. Available: https://doi.org/10.1109/TII.2014.2349359

[5] McKinsey & Company, ”Artificial intelligence in utility asset management,” Energy Insights, 2020. [Online]. Available: https://www.mckinsey.com/

[6] TTK, ”Leak Detection Systems for Data Centers,” Technical Documentation, 2023. [Online]. Available: https://www.ttk.fr/

[7] Sensaphone, ”Environmental Monitoring Solutions for Critical Facilities,” Product Guide, 2022. [Online]. Available: https://www.sensaphone.com/

[8] A. Aymon, A. Goldstein, and D. Cohen, ”Pipeline leak detection using Random Forest classification on pressure sensor data,” *Water Resources Management*, vol. 34, pp. 1453-1468, 2020. [Online]. Available: https://doi.org/10.1007/s11269-020-02507-4

[9] S. Choi and S. Im, ”Acoustic-based leak detection using deep convolutional neural networks,” *Journal of Hydroinformatics*, vol. 23, no. 2, pp. 367-381, 2021. [Online]. Available: https://doi.org/10.2166/hydro.2021.176

[10] M. A. Kammoun, I. Kammoun, B. Abid, and S. Masmoudi, ”Leak detection in wa- ter distribution networks using LSTM-based autoencoders,” *IEEE Access*, vol. 10, pp. 25308-25321, 2022. [Online]. Available: https://doi.org/10.1109/ACCESS.2022.3155634

[11] Y. Zhang, C. Xiong, and Y. Liu, ”Industrial equipment remaining useful life predic- tion using LSTM networks,” *Reliability Engineering & System Safety*, vol. 222, 108410, 2022. [Online]. Available: https://doi.org/10.1016/j.ress.2022.108410

[12] X. Zhu, L. Hou, and X. Chen, ”Hybrid LSTM-SVDD model for HVAC fault de- tection using prediction residuals,” *Building and Environment*, vol. 187, 107403, 2021. [Online]. Available: https://doi.org/10.1016/j.buildenv.2020.107403

[13] OASIS, ”MQTT Version 3.1.1 Specification,” OASIS Standard, 2014. [Online]. Available: https://mqtt.org/

[14] J. Wan, S. Tang, Z. Shu, D. Li, S. Wang, M. Imran, and A. V. Vasilakos, ”Software- defined industrial Internet of Things in Industry 4.0,” *IEEE Wireless Communications*, vol. 23, no. 5, pp. 137-143, 2016. [Online]. Available: https://doi.org/10.1109/MWC.2016.7721743

[15] InfluxData, ”InfluxDB Technical Overview: Time Series Data Platform,” Technical Documentation, 2023. [Online]. Available: https://www.influxdata.com/

[16] ASHRAE Technical Committee 9.9, ”Liquid

Cooling Guidelines for Datacom Equip- ment Centers,” ASHRAE, 2021. [Online]. Available: https://www.ashrae.org/

[17] CoolIT Systems, ”Direct Liquid Cooling Design Guide for High-Performance Com- puting,” Engineering Manual, 2022. [Online]. Available: https://www.coolitsystems.com/

[18] Asetek, ”Liquid Cooling Solutions for Data Centers: Flow Rate Specifications,” Technical Brief, 2021. [Online]. Available: https://www.asetek.com/

[19] J. Hamilton, ”Perspectives on Large-Scale Data Center Operations and Cooling,” *ACM Queue*, vol. 8, no. 1, 2010. [Online]. Available: https://doi.org/10.1145/1717801.1717805

[20] U.S. Department of Energy, ”Best Practices Guide for Energy-Efficient Data Center Design: Liquid Cooling Failure Modes,” DOE/EE Technical Report, 2011. [Online]. Available: https://www.energy.gov/

[21] NVIDIA, ”DGX H100 System Architecture and Facility Requirements,” Technical Brief, 2023. [Online]. Available: https://www.nvidia.com/

[22] Uptime Institute, ”Data Center Power Density Trends,” Global Survey, 2024. [On- line]. Available: https://uptimeinstitute.com/

[23] CoolIT Systems, ”Mean Time to Repair for Direct-to-Chip Cooling Failures,” Ser- vice Documentation, 2023. [Online]. Available: https://www.coolitsystems.com/

[24] Meta, ”Cascading Failure Prevention in Liquid-Cooled Infrastructure,” Engineering Blog, 2023. [Online]. Available: https://engineering.fb.com/

[25] ASHRAE Technical Committee 9.9, ”Liquid Cooling Reliability Metrics for Data- com Equipment Centers,” ASHRAE Report, 2021. [Online]. Available: https://www.ashrae.org/