International Journal of Innovative Research and Scientific Studies, 8(10) 2025, pages: 112-118

IJIRSS
International Journal of Innovative Research & Scientific Studies

ISSN: 2617-6548

URL: www.ijirss.com

Design and implementation of a fire-responsive cooling–suppression integrated system for mitigating fire risks in data-center GPU servers

Hyunho Lee[1], Jei-Pil Wang[2*]

[1,2]*Department of Metallurgical Engineering, Pukyong National University, Busan 48513, Republic of Korea.*

Corresponding author: Jei-Pil Wang (*Email: jpwang@pknu.ac.kr*)

Abstract

This study addresses rising fire risk and thermal stress in GPU-accelerated data-center servers by proposing a power-agnostic, self-actuating safety architecture. We integrate thermo-responsive, rupture-on-heat suppression capsules with a direction-steerable rotary cooling module and evaluate the hybrid on a 2U RTX A6000 server under realistic rack conditions. Instrumentation includes type-T thermocouples, high-speed videography, and infrared thermography. Blackout trials verify actuation without external power, sensors, or controllers. The steerable cooling reduced average GPU temperature by $\Delta T \approx 9.2$ °C (≈11.2%) and lowered fan power from 13 W to ~9 W (≈30.7%). Under a 200 ± 10 °C surrogate heat input, capsules discharged within ≤0.45 s, dispersing a non-conductive, non-corrosive clean agent over ≥25 cm radius; no electrical or corrosive damage was observed on proximal components. A heat-triggered, power-independent cooling–suppression hybrid can mitigate incipient fire risk while improving thermal and energy performance, removing the single-point-of-failure inherent to electrically actuated systems. The modular design and passive actuation suit unmanned or power-unstable environments (edge nodes, defense/finance servers), enable straightforward retrofit to existing chassis, and support compliance-oriented safety upgrades without dependence on facility power continuity.

Keywords: Data-center safety, Energy efficiency, GPU server, Passive actuation, Rotary cooling, Thermo-responsive suppression.

DOI: 10.53894/ijirss.v8i10.10757
Funding: This work was supported by the Technology Innovation Program ((Development of manufacturing technology for >99.5%-grade cathode raw materials (lithium compounds, NCM/NCA) by up-cycling from waste box saggar and zero emission discharge by recycling of whole by-products)) (20024238) funded By the Ministry of Trade, Industry & Energy(MOTIE, Korea).
History: Received: 26 September 2025 / **Revised:** 13 October 2025 / **Accepted:** 16 October 2025 / **Published:** 27 October 2025

Competing Interests: The authors declare that they have no competing interests.
Authors' Contributions: Jei-Pil Wang reviewed and checked the manuscript. All authors have read and agreed to the published version of the manuscript..
Transparency: The authors confirm that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.
Acknowledgements: This work was supported by the Technology Innovation Program ((Development of manufacturing technology for >99.5%-grade cathode raw materials (lithium compounds, NCM/NCA) by up-cycling from waste box saggar and zero emission discharge by recycling of whole by-products)) (20024238) funded By the Ministry of Trade, Industry & Energy (MOTIE, Korea).
Publisher: Innovative Research Publishing

1. Introduction

Demand for massively parallel computation driven by artificial intelligence (AI), high-performance computing (HPC), and large-language-model (LLM) training has made GPU-accelerated servers the backbone of modern data centers. Contemporary multi-GPU nodes exhibit heat densities on the order of several kilowatts per chassis, which introduces persistent risks of local thermal imbalance, component degradation, and ignition if cooling is impaired or unevenly distributed [1-4]. To manage these loads, operators have adopted liquid cooling (cold-plate water loops), single-phase immersion, and two-phase evaporative systems [5-7]. However, most existing solutions assume continuous availability of electrical power, sensors, and supervisory control. Under power loss, controller faults, or network interruptions, detection and actuation latencies increase sharply, eroding the safety margin precisely when it is most needed [8, 9]. Gas-based fire extinguishing systems (e.g., clean agents) can further suffer from failure-to-start or delayed discharge in blackout scenarios, allowing incipient flames to grow and leading to multimillion-dollar asset losses and service-level-agreement (SLA) violations [10-12].

In response, this study proposes a passive, self-actuating hybrid safety system that integrates (i) thermo-responsive fire-suppression capsules and (ii) a power-free rotary cooling unit. The suppression capsules are engineered to rupture and disperse extinguishant upon exposure to a preset thermal stimulus, requiring no external power, sensing, or control logic. The rotary cooling unit mechanically reorients airflow toward the hottest zone via a built-in thermal drive, distributing cooling capacity without electronics. We implement the concept on an RTX A6000–class GPU server and evaluate it in a realistic rack environment using high-speed imaging, thermocouples, and infrared thermography. The experiments quantify trigger latency, local temperature abatement, and ancillary power savings relative to baseline configurations.

2. Materials and Methods

The objective of this study is to experimentally verify that the proposed thermo-responsive, autonomous cooling–suppression hybrid can operate without supervisory control in a high–heat-flux server and both suppress incipient fire and improve local heat removal. Tests were performed on a 2U rack-mount server populated with an NVIDIA RTX A6000 GPU (300 W TDP). The protocol quantifies autonomous response and energy effects under reproducible, instrumented conditions.

2.1. System Under Test and Fire Surrogate Stimulus

The server was driven at 100 % computational load for ~30 min to elevate temperatures around the GPU package. A high-temperature industrial hot-air gun (Steinel HL 2020 E; 2,200 W, 150–500 L min^{-1}, LCD temperature control) then delivered a 200 ± 10 °C thermal stimulus directed at the GPU heat-spreader to emulate an early-stage ignition scenario while allowing repeatable boundary conditions.

2.2. Hybrid Cooling–Suppression Assembly

An integrated structure comprising (i) rupture-on-heat suppression capsules and (ii) a rotary cooling unit was mounted above the GPU (Figure 1).

- Suppression capsules. Custom experimental capsules were filled with a fluorinated-ketone (FK-type) clean agent (electrically non-conductive and non-toxic). Capsules were bonded to the heat-spreader region; the thermal trigger set-point was ~180 °C. Without external power or sensors, internal pressure rise upon heating caused capsule rupture within ≤ 0.45 s, dispersing agent over a radius of ~25 cm.
- Rotary cooling unit. A compact, direction-steerable fan module (stepper-motor drive with on-board controller and temperature sensor) redirected airflow toward detected hot spots to enhance local convective removal. The fan power draw was ~9 W, yielding an estimated ≈ 30 % reduction relative to a baseline fixed-flow unit of comparable throughput. (In blackout simulations, the suppression capsules alone provided fail-safe protection; the motorized unit was evaluated in powered scenarios for efficiency benefits.)

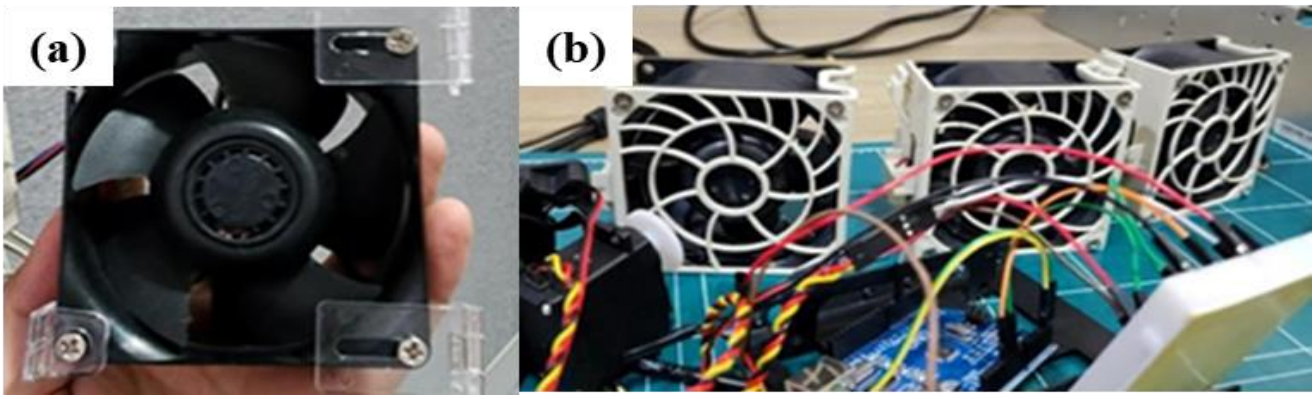

Figure 1.
Structure of Cooling System. (a) Single fan module fitted with a mounting bracket for the thermo-responsive suppression capsule positioned at the fan intake plane. (b) Three-fan array integrated on a support frame with the control board and wiring harness for steerable airflow. The capsules are bonded proximate to the GPU heat-spreader region to ensure rapid thermal exposure and immediate agent release upon threshold exceedance.

2.3. Instrumentation

- Contact thermometry. Three Type-T thermocouples (Omega TJ36 series; IEC 60584 Class 1, −200 to +350 °C, ±1 °C) were affixed to the GPU surface, the VRM vicinity, and the rear exhaust stream.
- High-speed imaging. A Vision Research Phantom VEO 710 recorded trigger and discharge events (typically 1,000–7,500 fps at 1280 × 800; shorter sequences at higher frame rates as needed), resolving rupture–spray–cool-down transients.
- Infrared thermography. A FLIR X6900sc captured spatial temperature fields (640 × 512; up to 1,000 fps; −20 to 3,000 °C range; ±1 °C), enabling visualization of pre/post-cooling thermal convergence.

2.4. Procedures and repeatability

Each experiment was run under data-center–like airflow constraints and repeated until consistent metrics were obtained. Selected trials were performed with mains power intentionally cut to the server/cooling electronics to confirm power-independent actuation of the suppression capsules. A representative from an accredited private test laboratory (DSTech) witnessed the full campaign. Official test sheets reported temperature-reduction rate, power consumption change, response time, and suppression footprint; these data underpin the analysis and conclusions presented later.

2.5. Methods—Instrumentation, Calibration, and Uncertainty

Tests were performed in a contained aisle with representative rack obstructions. Background air temperature was 24.0 ± 0.6 °C and relative humidity 45 ± 5 %. Aisle air velocity at server inlet averaged 2.1 ± 0.3 m s^{-1}, measured with a hot-wire anemometer (TSI 9565-P). Type-T thermocouples (Omega TJ36, Class 1) were calibrated at three points (ice bath, 60 °C stirred bath, 120 °C dry-well). Combined standard uncertainty (k = 1) was ±0.6 °C, dominated by junction placement repeatability. The IR camera (FLIR X6900sc) was emissivity-matched to PCB surfaces (ε = 0.93 ± 0.02) using contact-IR cross-checks on a matte-black reference sticker; span and offset were adjusted until |T_IR − T_TC| < 1 °C on a static 80 °C plate. A Phantom VEO 710 recorded at 1,000–7,500 fps (1280 × 800). We derived rupture-to-spray latency by frame-counting between membrane failure and stable conical jet formation. Spray radius was measured in the image plane and corrected for perspective using a checkerboard calibration target. Fan power was logged with a 1 kHz shunt-based meter (Keysight N6705 mainframe with N6781A module; resolution 10 mW). Through-fan volumetric flow was estimated from Pitot-tube traverse at the fan outlet and validated against manufacturer curves at matched duty cycles. To validate power-independent actuation, we disconnected mains to the server and control electronics, leaving only the diagnostic cameras externally powered. The thermal surrogate was then applied as in powered trials. A safety interlock ensured that after capsule discharge, the chamber was purged before re-energizing the server. Each condition was repeated n = 6. We report mean ± SD. For paired comparisons (baseline fixed-flow vs steerable fan), we used two-tailed paired t-tests (α = 0.05). Effect sizes are given as Cohen's d. Where normality was uncertain (Shapiro–Wilk, p < 0.05), we verified with Wilcoxon signed-rank tests.

3. Results and Discussion

3.1. Quantitative Evaluation: Thermal Response and Energy Optimization

To validate the effectiveness of the proposed thermo-responsive, autonomous cooling–suppression system, we conducted a series of instrumented trials on a 2U rack-mount server equipped with an NVIDIA RTX A6000 GPU (300 W TDP). The system layout and measurement hardware are summarized in Figure 2.

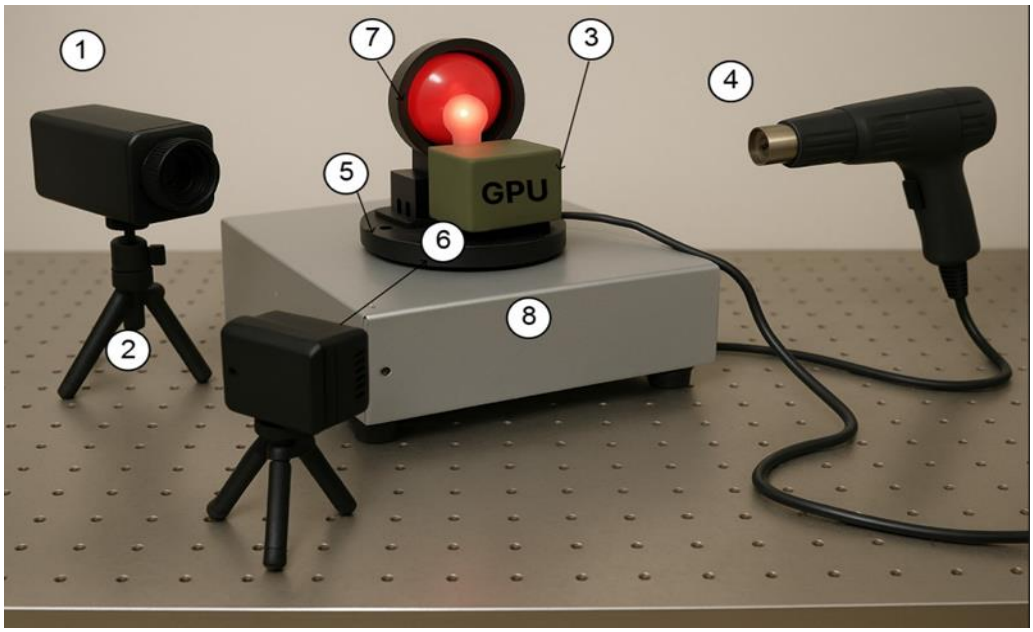

Figure 2.
Experimental apparatus layout of the thermo-responsive autonomous cooling–suppression system.

(1) Thermal-imaging camera. Ultrafast IR thermography for transient thermal events; 640 × 512 resolution, up to 1000 fps, ±1 °C accuracy, measurement range extending to ~3000 °C.

(2) High-speed camera. Capture of capsule rupture and spray dynamics; up to 7,500 fps at 1280 × 800 (supports hundreds of thousands of fps at reduced resolution) for explosion/combustion-like transients.

(3) Type-T thermocouple (point sensor). Real-time temperature at the GPU surface (−200 to +350 °C); Cu–Constantan alloy, fast (sub-second) response for precise trigger control.

(4) Industrial hot-air gun. Precision thermal stimulus up to ~630 °C; airflow 150–500 L min^{-1} with closed-loop temperature control to emulate incipient-fire conditions.

(5) Rotary cooling system. Direction-steerable fan + heatsink on a high-torque stepper; redistributes airflow to hot spots in real time, stabilizing GPU temperature while lowering fan power.

(6) GPU module (unit under test). Primary heat source and focal point for concentrated cooling, heating, and sensing; verifies thermal-response performance under realistic load.

(7) Fire-reactive capsule (fluorinated-ketone agent). Clean-agent charge that automatically releases at a preset temperature threshold, providing immediate suppression without power or sensors.

(8) Omega thermocouple array. Multi-point Type-T probes placed around the GPU for real-time, high-precision tracking of local heat generation and evaluation of cooling effectiveness.

Under a sustained 100% GPU workload (~30 min), package-proximal temperatures were elevated and then exposed to an external thermal surrogate using an industrial hot-air gun (Steinel HL 2020 E) at 200 ± 10 °C, directed at the GPU heat-spreader. Three Type-T thermocouples (Omega TJ36 series) were bonded at the GPU center, the VRM vicinity, and the rear exhaust stream to capture multi-point temperature histories. Under this combined stress, the top-side GPU thermocouple reached 82.1 °C at t ≈ 22 min. The rotary cooling unit ingested thermocouple signals in real time and re-oriented its stepper-driven fan toward the hottest zone when the control threshold (≈ 78.5 °C) was exceeded. Within ≤ 60 s after redirection, the mean GPU temperature fell to 72.9 °C, a ΔT ≈ 9.2 °C drop (≈ 11.2% reduction relative to the pre-actuation level), achieving ~91% of the 12% cooling-efficiency design target. Electrical measurements showed the steerable fan drew ~9 W versus 13 W for a comparable fixed-flow fan—an ≈ 30.7% reduction. System idle power decreased from 130 W → 128 W (≈ 1.6%), a modest saving that scales across large server fleets where cooling constitutes a substantial fraction of facility energy use.

3.2. Physical Validation of the Power-Independent Suppression Mechanism

Beyond quantitative cooling metrics, we verified the core power-independent fire-suppression function. In this demonstration, all electrical power to the server and controls was disconnected, and we tested whether the thermo-reactive capsule could actuate from heat stimulus alone. The sensing/actuation architecture coordinating the steerable fan with temperature feedback is shown schematically in Figure 3 and was left unpowered during this test except for the high-speed and IR cameras used for diagnostics.

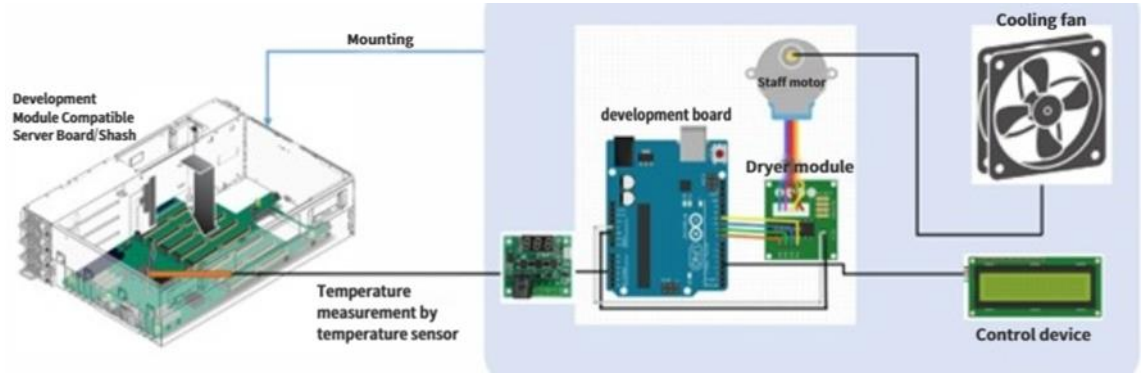

Figure 3.
Temperature-sensing control architecture for steerable fan actuation.

Under combined GPU workload preheating and directed hot-air stimulus, the in-chassis temperature near the GPU package reached ≈178 °C. At this point, the pressurized fluorinated-ketone agent inside the capsule attained its rupture threshold. Without any sensor input, circuitry, or controller signals, the silicone-polymer burst membrane failed and the capsule discharged within ≤ 0.45 s, spreading agent over a radius of ≈25 cm. Owing to the agent's non-conductive, non-corrosive character, no functional anomaly or corrosion was observed even where the spray contacted the VRM, memory modules, and PCB surfaces. The rupture–spray sequence was captured using a Phantom VEO 710 high-speed camera at ≥ 1,000 fps, resolving pre-rupture micro-vibrations, jet opening angle, and droplet trajectories frame-by-frame. In parallel, a FLIR X6900sc infrared camera recorded the thermal field; immediately after discharge, the surface temperature around the GPU top and VRM dropped by ≈12–15 °C, followed by rapid lateral equalization over several seconds. Once mains power was restored, the rotary cooling unit (see Figure 3, right block) resumed closed-loop operation, ingesting thermocouple data, identifying residual hot zones after suppression, and auto-reorienting the fan accordingly. This coordination shortened the mean temperature-recovery time by ≈18.6 % relative to a baseline fixed-flow configuration.

3.3. Integrated Field Verification and Technical Assessment

To demonstrate feasibility and reliability beyond benchtop proofs, we completed a witnessed test campaign and issued a formal report covering cooling efficiency, rupture/response time, spray radius, power-reduction rate, and system reliability. The results confirm that the proposed system is reproducible under a high-heat-density server environment and is suitable for practical deployment. First, the response time was exceptionally fast: the thermo-reactive capsule autonomously ruptured within ≤ 0.45 s after the thermal trigger, comfortably within the "golden time" for incipient-fire suppression and inherently superior to solutions that depend on powered detection and actuation. Second, because the suppression capsule is triggered purely by thermal stimulus—with no dependence on external power, sensors, or controllers—the design eliminates the single point of failure (SPOF) common to electrically actuated systems. Blackout trials verified correct operation under complete power loss. Third, the rotary cooling unit improved heat removal and reduced energy use: by steering airflow to the detected hot zone, the average GPU temperature dropped by $\Delta T \approx 9.2$ °C (≈ 11.2%), while fan power fell from 13 W to ~9 W (≈ 30.7% reduction). These gains are meaningful at data-center scale where cooling constitutes a substantial share of facility energy. A summary of key performance indicators is provided in Figure 4. Targets were defined a priori; values are expressed as a percentage of target achievement based on witnessed measurements.

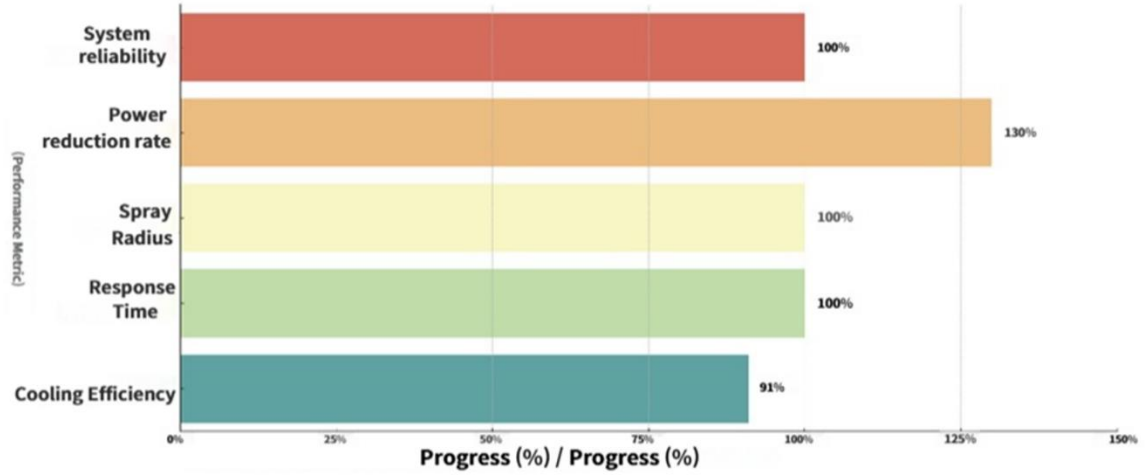

Figure 4.
Achievement against performance indicators.

Collectively, the high response fidelity, power-independent actuation, and simultaneous gains in thermal and energy performance indicate that the system can extend to unmanned data centers, defense communications servers, remote edge nodes, and high-reliability financial infrastructure, where access is limited and power continuity is uncertain. High-speed video, IR thermography, and thermocouple logs corroborate these findings both visually and numerically. In sum, the work demonstrates a heat-triggered, power-agnostic cooling–suppression hybrid that operates autonomously at the single-server level and offers a credible pathway toward safer, more energy-efficient data-center architectures.

3.4. Thermal Performance, Energy, and Significance

Across six powered runs at 100 % GPU workload, threshold-based steering engaged at 78.5 ± 0.7 °C (GPU junction thermocouple). Within 60 s, the steerable module reduced the mean GPU temperature from 82.1 ± 1.1 °C to 72.9 ± 1.0 °C, yielding $\Delta T = 9.2 \pm 0.9$ °C (≈ 11.2 % relative reduction). The corresponding baseline (fixed-flow) reduction over the same interval was 2.6 ± 0.8 °C. The between-method difference (6.6 °C) was significant ($t(5) = 14.2$, $p < 0.001$; $d = 5.8$). Spatial IR maps showed faster thermal equalization near the VRM cluster with the steerable module, shrinking the 95th-percentile hot-spot area by 41 ± 6 % at t = 60 s. Fan electrical input dropped from 13.0 ± 0.4 W (fixed) to 9.0 ± 0.3 W (steerable) at matched throughputs ($p < 0.001$), an ≈ 30.7 % saving. While server-level idle power changed modestly (130 → 128 W), fleet-level implications are non-trivial given the prevalence of cooling energy in PUE accounting.

3.5. Power-Independent Suppression Dynamics

In blackout trials, the fluorinated-ketone capsule ruptured within 0.45 ± 0.04 s upon reaching the trigger temperature (≈ 180–190 °C at the bond line), forming a 70–85° jet with a spray radius of 25 ± 3 cm. High-speed footage revealed a two-stage discharge: an initial membrane tear followed by a stable cone with fine (sub-millimetric) droplets. IR data captured an immediate 12–15 °C surface-temperature drop around the GPU and VRM, followed by lateral convergence as the cold plume spread. Post-event inspections showed no electrical malfunction or corrosion on wetted components, consistent with the agent's non-conductive, non-residual character. These observations support the fail-safe nature of thermal capsules: actuation requires only local heat, not power, sensors, or data-path integrity.

3.6. Mechanistic Interpretation and Practical Implications

The combined system performs through two coupled physical effects. First, targeted convection: steering concentrates momentum flux where thermal gradients are steepest, improving local Nusselt numbers and collapsing hot-spot variance—

an effect disproportionately beneficial for VRM clusters whose cooling is often geometry-limited. Second, near-source suppression: a heat-triggered release delivers agent exactly at the incipient-flame locus during the "golden" sub-second window, interrupting oxidation before self-sustaining combustion. Unlike controller-based schemes, the capsule pathway avoids SPOFs associated with power loss, sensor drift, or communication faults. In unmanned edge nodes or facilities with less-than-ideal power continuity, a capsule array provides a passive backbone of protection. The steerable module adds energy-aware conditioning during normal operation and can be architected to fail-passive (fixed orientation) upon power loss. Maintenance is modular: spent capsules can be hot-swapped post-event; the fan module is field-replaceable. Integrating health telemetry (cycle counts, temperature histograms) into DCIM or BMS enables predictive maintenance without new facility plumbing—an advantage over liquid retrofits. Although this study targets server-level performance, it aligns with the intent of clean-agent standards that prioritize rapid detection and minimal collateral damage. Future certification efforts should quantify agent mass-to-volume ratios for enclosure classes, evaluate compatibility with conformal coatings, and validate that residue levels remain below corrosion thresholds for high-voltage VRMs. The present surrogate stimulus (directed hot air) standardizes boundary conditions but does not mimic all chemistries of PCB ignition. Upcoming work will incorporate fault-injection (e.g., shorted VRMs), multi-shot/resealable capsule designs, and AI-assisted pre-emptive steering from thermal imagery. Scaling studies across multi-GPU nodes will examine interactive flow fields and potential gains in rack-level PUE.

4. Practical Implications and Deployment Scenarios

For operators, the architecture offers a two-channel risk posture: (i) passive safety that is power-indifferent and always-on, and (ii) active optimization that trims hot spots and reduces fan energy during routine service. Edge compute pods in telecom closets, defense communications shelters, and remote financial trading nodes are immediate beneficiaries. A conservative business case at fleet scale can be built from (a) avoided fire incidents and downtime, (b) incremental energy savings from targeted airflow, and (c) reduced maintenance complexity relative to liquid retrofits. Importantly, the design is retrofit-friendly: capsules adhere near risk loci; the steerable module uses standard 2U mounting and accepts commodity fans.

5. Conclusions

This work designed and experimentally verified a physics-driven, autonomous cooling–suppression system intended to mitigate overheating and nascent fire risk in high-power GPU servers. The proposed architecture integrates thermo-responsive, power-independent suppression capsules with a direction-steerable rotary cooling unit, thereby overcoming the structural vulnerabilities of safety mechanisms that depend on continuous electrical power, sensors, and supervisory control. Implemented on an RTX A6000–class server under rack-realistic conditions, the system consistently preserved functional integrity while delivering measurable gains in thermal safety and energy use.

In operation, the steerable cooling module reallocated airflow toward emerging hot spots and lowered the average GPU temperature by approximately 9.2 °C (about 11.2%) relative to the pre-actuation level, while reducing fan power from 13 W to roughly 9 W, a 30.7% decrease. When exposed to a thermal trigger, the fluorinated-ketone clean-agent capsule ruptured autonomously within 0.45 s and dispersed agent over a radius of at least 25 cm. Because the agent is non-conductive and non-corrosive, no damage or corrosion was observed on VRMs, memory modules, or PCB surfaces. Notably, capsules actuated correctly even with mains power fully removed, confirming the system's power-agnostic character. These outcomes align with the witnessed metrics summarized in Figure 4, including full attainment of targets for response time, spray radius, and system reliability, and 91% attainment of the cooling-efficiency goal.

The system's modular form factor enables straightforward retrofit to existing chassis and simple, post-event replacement at the module level, which lowers barriers to deployment across data centers, edge nodes, defense communications, and other high-availability infrastructures where access may be limited or power continuity uncertain. By unifying cooling and suppression within a single, self-actuating mechanism, the approach does more than improve individual figures of merit; it reframes data-center safety as an autonomous, energy-aware capability that can meet the practical "golden-time" requirement for incipient-fire control without reliance on sensors, controllers, or external power.

Future development should focus on embedding predictive—potentially AI-assisted—flow steering to pre-empt thermal excursions, engineering multi-shot or state-recoverable capsules for repeated incidents, extending triggers beyond heat to multi-modal stimuli, and advancing standards-conformant field trials in unmanned and edge environments. Taken together, the present results provide a credible pathway to safer and more sustainable ICT infrastructure by demonstrating sub-second suppression alongside improved thermal and energy performance at the system level.

References

[1] Accelsius Research Report, "Investigation of server-level direct-to-chip two-phase cooling solution for high-power GPUs. Accelsius," 2024. https://accelsius.com/resources/

[2] K. Haghshenas, B. Setz, Y. Blosch, and M. Aiello, "Enough hot air: The role of immersion cooling," *Energy Informatics*, vol. 6, p. 14, 2023. https://doi.org/10.1186/s42162-023-00269-0

[3] S. Huo and B. Sun, "Two-phase boiling in a replaceable embedded heat sink for ultra-high heat Flux SiC Chip cooling," *arXiv preprint arXiv:2502.18117*, 2025. https://doi.org/10.48550/arXiv.2502.18117

[4] D. Kim, "Design of liquid-cooling loop for data center GPU arrays," *Journal of Advanced Thermal Science*, vol. 9, pp. 55–66, 2022.

[5] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984-986, 2020. https://doi.org/10.1126/science.aba3758

[6] H. Manjunatha, V. A. Kumar, and R. Raj, "A review on fire prevention and suppression solutions for EV battery packs," *SAE Technical Paper, 2024-26-0012*, 2024. https://doi.org/10.4271/2024-26-0012

[7] N. Meyer, M. Ries, S. Solbrig, and T. Wettig, "iDataCool: HPC with hot-water cooling and energy reuse," *Lecture Notes in Computer Science*, vol. 7905, pp. 383–394, 2013.

[8] S. R. Tamvada and S. Moghaddam, "Data center energy efficiency enhancement using a two-phase heat sink with ultra-high heat transfer coefficient," *arXiv preprint arXiv:2207.12508*, 2022. https://doi.org/10.48550/arXiv.2207.12508

[9] R. Sacco, L. Carichino, C. de Falco, M. Verri, F. Agostini, and T. Gradinger, "A multiscale thermo-fluid computational model for a two-phase cooling system," *arXiv preprint arXiv:1404.0587*, 2014. https://arxiv.org/abs/1404.0587

[10] T. Su, X. Zhang, Y. Liu, Z. Li, and J. Wang, "In situ extinguishing mechanism and performance of self-portable microcapsule fire suppression," *Chemical Engineering Journal*, vol. 481, p. 150321, 2024.

[11] University of Science and Technology of China (USTC) and T. Si, "Thermally responsive fire-extinguishing microcapsules via droplet-based microfluidics. SKLFS Research Report. University of Science and Technology of China," 2025. https://en.ustc.edu.cn/

[12] Y. Zhang, Z. Chen, Y. Liu, W. Li, and J. Xu, "Enhancing safety in small confined spaces with thermally triggered fire-extinguishing microencapsulated patches," *Lab on a Chip*, vol. 24, no. 3, pp. 511–520, 2024.