



Thermal Excursion Event Case Study

Physics-Informed Detection and Prescriptive Cooling
Response in GPU Compute Infrastructure

DGX H100 · StuxGrid Mesh · Entropy Platform · Vertiv PDX

CLASSIFICATION

Confidential — Investor Distribution

Contents

- 1 Abstract**
- 2 Facility and Instrumentation**
- 3 Event Timeline**
 - 3.1 Baseline and Test Initiation
 - 3.2 Thermal Excursion
 - 3.3 Shutdown and HVAC Response
- 4 Entropy Platform Response**
 - 4.1 Anomaly Detection
 - 4.2 Failure Mode Discrimination
 - 4.3 Prescribed Cooling Modifications
- 5 Discussion**
- 6 Appendix**

1 Abstract

During a scheduled stress test of an 8xH100 DGX Station cluster on April 9, 2026, a GPU power fault in DGX02 produced an acute thermal excursion: +4.8°C at the Pod-1 exhaust sensor in a single 60-second interval, a subsequent peak of 35.5°C during HVAC oscillation, and an 8.6°C collapse over seven minutes following full cluster shutdown. The event triggered sustained compressor short-cycling in the Vertiv Liebert PDX cooling unit.

The HyperBase Entropy platform, operating its physics-informed neural network (PINN) observation loop across the StuxGrid sensor mesh, detected the thermal anomaly within the first measurement interval. Using a convection-diffusion model trained on 30M+ historical data points, Entropy performed real-time root cause classification, identified the event as a GPU power fault rather than a cooling system failure, and prescribed three modifications to the HVAC control strategy to mitigate compressor damage and prevent thermal overshoot during future load transients.

< 60s

DETECTION WINDOW

+4.8°C

PEAK THERMAL SPIKE

8.6°C

POST-SHUTDOWN COLLAPSE


30M+

TELEMETRY DATA POINTS

2 Facility and Instrumentation

The facility houses two DGX Stations in adjacent pods sharing a common hot aisle, cooled by a single Vertiv Liebert PDX precision unit. DGX01 and DGX02 each contain 8×H100 GPUs (640GB HBM3 aggregate per node). DGX02 had a known GPU fault (slot 6, POST code 02; 7/8 GPUs healthy) prior to the test.

Environmental telemetry is provided by the StuxGrid sensor mesh: ESP32 nodes with Bosch BME280 sensors publishing via MQTT to a Mosquitto broker, ingested through Telegraf into InfluxDB v2.7.12 at approximately 25 messages/second aggregate.



Entropy Platform Data Flow — Sensor to Prescription
Figure 1. Entropy platform data pipeline, from physical sensors through PINN inference to prescriptive output.

Sensor ID	Desig.	Position	Event Status
SN-EB0i-004	T4	DGXpod-1 Front Intake	Active
SN-EB0i-002	T5	DGXpod-1 Back Exhaust	Active (primary signal)
SN-EBD-001	T6	DGXpod-2 Front Intake	Active
SN-EB0i-001	T7	DGXpod-2 Back Exhaust	Excluded — calibration drift*
SN-EB0i-003	T8	PDX Intake	Excluded — calibration drift*

Table 1. StuxGrid sensor deployment. *T7 and T8 exhibited static readings throughout the event window, consistent with calibration drift identified by Entropy's health monitor in prior telemetry windows. Both were excluded from active analysis and queued for recalibration.

Entropy's sensor health monitor had flagged T7 and T8 as operating outside nominal calibration bounds in routine sweeps prior to this event. The PINN model was operating on a reduced but known-good sensor set at the time of the incident, rather than incorporating unreliable data. The event validated the health monitor's earlier assessments.

3 Event Timeline

All timestamps CDT (UTC-5). Duration: ~90 minutes.

3.1 Baseline and Test Initiation

Prior to the test (17:00–18:31), T4 ≈ 30.3°C, T5 ≈ 30.1°C, T6 ≈ 28.8°C — steady-state under background load. Entropy registered all active sensors within two standard deviations of their 7-day rolling baselines. At 18:31, SSH to DGX02 initiated the stress test. By 19:00, T6 began rising as Pod-2 GPUs engaged. At 19:28, DGX01 completed; T4 = 27.97°C, T5 = 27.92°C.

3.2 Thermal Excursion (19:31–19:34 CDT)




Figure 2. T4 and T5 thermal profiles during the excursion window (00:26–00:36 UTC). Shaded region: GPU fault interval. T5 spike of +4.8°C in 60 seconds at 00:31, followed by intake collapse at 00:34.

Time (UTC)	T4 (Pod-1 In)	T5 (Pod-1 Out)	T6 (Pod-2 In)	Note
00:28	27.97°C	27.92°C	28.15°C	DGX01 completes
00:29	28.09°C	27.90°C	28.20°C	
00:30	28.19°C	27.91°C	28.22°C	
00:31	28.21°C	30.16°C ↑↑	rising	Entropy alerts
00:32	28.34°C	32.75°C ↑↑↑	rising	Peak transient
00:33	28.06°C	29.70°C	rising	
00:34	26.72°C ↓	25.49°C	–	Intake collapse

Table 2. Minute-resolution data during excursion.

The T5 step function — +4.8°C in 60 seconds — is inconsistent with gradual load increase. Simultaneous T4 collapse (28.3 → 26.7°C) indicates fan-driven airflow cessation, consistent with compute node power loss.

3.3 Shutdown and HVAC Response (19:34–20:02+)

T5 oscillated 25–35.5°C for ~27 minutes (peak 35.5°C at 19:52). At 20:01, exhaust crossed below intake, confirming zero heat output. Subsequent ~10-minute oscillation is compressor short-cycling: the PDX chasing a cold, heatless room.

Time (UTC)	T5 (Pod-1 Out)	T4 (Pod-1 In)	Interpretation
01:00	33.97°C	25.75°C	HVAC oscillation
01:01	28.65°C ↓↓	25.71°C	Cluster power loss
01:02	25.20°C	25.44°C	T5 ≈ T4: zero heat generation
01:06	22.16°C	24.76°C	Below supply temperature

01:08	20.60°C	23.94°C	Full offline confirmed
-------	---------	---------	------------------------

Table 3. Cluster shutdown. T5 collapse: 8.6°C in 7 minutes.

4 Entropy Platform Response

4.1 Anomaly Detection

At each telemetry interval, the PINN predicts the expected thermal field given current boundary conditions. The physics residual — the difference between predicted and observed temperatures — is the primary anomaly signal. During baseline, residuals averaged 0.12°C. At 19:31, Entropy registered a T5 residual of +4.6°C — exceeding 30 standard deviations — and classified the event as a Level 1 Thermal Excursion.

Note: Conventional threshold monitoring would flag T5 only after crossing 35°C. Entropy detected the anomaly at 30.16°C because the model predicted 27.9°C. The signal is not that temperature was high; it is that temperature diverged from physics.

4.2 Failure Mode Discrimination

Within the same detection cycle, Entropy performed differential analysis of the spatial thermal response:

Hypothesis	Expected Signature	Observed	Match
HVAC failure	All sensors rise in phase	T5 spikes; T4 stable	No
Gradual overload	Monotonic ramp over minutes	4.8°C step in 60s	No
GPU power fault	Exhaust spike + intake drop	T5: +4.8°C; T4: -1.6°C	Yes

Table 4. Failure mode differential. Anti-correlated T4/T5 = power fault; in-phase = cooling.

The PINN's R_{ij} coupling matrix quantifies T4–T5 cross-pod coupling at $R = 0.28$. This learned coefficient enabled Entropy to confirm the anti-correlation signature as a GPU power fault within the same measurement cycle.

4.3 Prescribed Cooling Modifications

Entropy generated three prescriptions targeting the immediate compressor risk, the HVAC control vulnerability, and the physical infrastructure.

P1: Compressor Short-Cycle Prevention

Transition PDX to load-adaptive dead-band mode when exhaust drops below intake for ≥ 2 consecutive intervals. The observed ~10-minute short-cycling accelerates compressor wear and risks thermal protection lockout. Dead-band mode allows idle when the supply-return differential falls below efficient operating range.

P2: Thermal-Inertia-Aware Setpoint Management

PINN-driven dynamic setpoint: when sudden load reduction is detected via dT/dt , preemptively raise setpoint 2–3°C to approximate new steady-state. The model's learned thermal time constants (~5.3-hour Liebert limit cycle) enable prediction within seconds. Projected: substantial attenuation of the 35.5°C peak overshoot.

P3: Exhaust-Path Isolation

Deploy blanking panels or active dampers. $R_{ij} = 0.28$ indicates 28% cross-pod coupling. Physical isolation reduces this toward zero, improving per-pod fault localization and PINN spatial resolution.

5 Discussion

Sensor mesh resilience. Entropy operated on 3 of 5 active room sensors, having previously excluded T7 and T8 due to calibration drift. The PINN's ability to detect, classify, and prescribe using only the known-good subset demonstrates graceful degradation under partial observability.

HVAC control limitations. The PDX's PID controller assumes quasi-static thermal loads. Abrupt transitions produce step-changes PID handles poorly by design. The 27-minute oscillation is a predictable consequence. P2 addresses this with physics-informed targets.

Detection versus actuation. Entropy detected and classified within the first measurement interval. Practical impact depends on the actuation path: P1/P2 require PDX control interface integration (in active development). P3 requires mechanical modification. The current deployment demonstrates detection and prescription; closed-loop actuation is the next milestone.

Metric	Observed (Pre-Entropy)	With Entropy
Detection latency	> 5 min (threshold)	≤ 60s (physics-residual)
Root cause method	Manual forensic	Automated differential (≤ 60s)
HVAC oscillation	~27 min sustained	Substantially reduced*
Peak overshoot	35.5°C	Attenuated*
Compressor stress	Short-cycling observed	Dead-band prescribed*
Degraded sensor handling	Discovered post-incident	Pre-flagged; excluded

Table 5. Comparative analysis. *Pending closed-loop validation.

6 Appendix

A. StuxGrid Node Specifications

Component	Specification
Microcontroller	Espressif ESP32-WR00M-32
Sensor	Bosch BME280 (T, RH, P)
Temperature Accuracy	$\pm 1.0^{\circ}\text{C}$ (typ. $\pm 0.5^{\circ}\text{C}$)
Interval	~ 2.5 s (configurable)
Protocol	[REDACTED]
Pipeline	[REDACTED]

Table 6. StuxGrid hardware.

B. Entropy Platform Stack

Layer	Technology	Function
Telemetry	[REDACTED]	Environmental sensing
Transport	MQTT (Mosquitto)	Real-time data bus
Storage	[REDACTED]	Time-series (30M+ pts)
Physics	PINN [REDACTED]	Thermal model
Reasoning	[REDACTED]	Synthesis / planning
Embedding	[REDACTED]	Vector search
Memory	[REDACTED]	Agent state
Orchestration	JAWS [REDACTED]	Agent coordination
Compute	NVIDIA DGX Spark (GB10)	Training / inference

Table 7. Entropy stack. Entire pipeline on-premises.

Confidential. Quantitative projections are conservative; values noted as pending validation require closed-loop deployment for empirical confirmation.

© 2026 HyperBase Ventures. All rights reserved.