PURETENSOR
INTELLIGENCE DISTILLED

Thermally Blind by Design

A Super I/O register limit hides memory temperature from fan control

Paper: PT-TN-2026-004 v1.0

Author: PureTensor Ltd

Date: 11 February 2026

Status: Published

Abstract

The ASUS Pro WS TRX50-SAGE WIFI, paired with a Threadripper 7970X and 256 GB of DDR5, ships with a fan controller that cannot see memory temperature. The Nuvoton NCT6798 Super I/O exposes the AMD Temperature Sensor Interface as sensor 13; the kernel driver clamps the fan-source selector at 12. The BIOS SmartFan implementation drives all PWM channels from PECI Agent 0 (the CPU package), so the DIMMs sit outside every hardware thermal loop. Under a 240 GB memory-stress workload the CPU held 35–44 °C while the aggregate DIMM channel reached 85 °C or higher, and the fans, obeying a correct curve on the wrong sensor, never exceeded roughly 40 %.

Five stress runs on 9 February 2026 bracketed the problem. A user-space daemon, dimm-fanctl, reading TSI0 every five seconds and driving PWM directly, brought peak memory temperature to 54.5 °C under sustained 64-bit writes, 26 °C below the 80 °C acceptance threshold. The daemon has since become permanent infrastructure on this node.

A secondary study on a different workstation (WRX90E-SAGE SE, dual RTX PRO 6000 Blackwell open-air cards) demonstrated a related blind spot: GPU 1 reached 93 °C and throttled while GPU 0, in the cooler slot, peaked at 82 °C. Removing the side panel and applying an earlier custom fan curve for GPU 1 dropped its peak to 87 °C and eliminated throttling, but the two interventions were made together; the improvement cannot be attributed to either alone.

The finding is architectural. A firmware register one index short of the sensor that matters leaves 256 GB of DDR5 with no hardware thermal protection on a board sold for memory-intensive workloads. The failure is invisible to any check that reads temperatures rather than asking what the controller is allowed to read.

1. The machine and its sensors

Workstation A is an ASUS Pro WS TRX50-SAGE WIFI carrying an AMD Ryzen Threadripper 7970X (32 cores, 64 threads) and four 64 GB Samsung M321R8GA0PB0-CWMXJ DDR5-5600 DIMMs running at 5200 MT/s in quad-channel configuration, one DIMM per channel. The kernel is 6.17.4-2-pve under Proxmox VE 8.x; the chassis is a Silverstone rackmount, freshly re-housed and racked before these measurements.

Before any stress run the sensor inventory was established. The Nuvoton NCT6798 Super I/O exposes a single aggregate memory-temperature channel, TSI0_TEMP, as temp13. TRX50 does not break out per-DIMM TSI channels to the Super I/O. The k10temp driver provides CPU die temperatures: Tctl (aggregate) plus CCD2, CCD3 and CCD4. CCD1 reads 0 °C on this SKU, either unpopulated or not reporting; this did not affect the test. The spd5118 driver loads, but the AMD SMBus does not expose the individual DIMM thermal hubs, so TSI0 is the only memory-temperature channel available on this platform.

A prior-day post-rack validation had already noted "no dedicated DIMM temp sensors on this board (NCT6798 chipset)". The aggregate TSI0 channel was found only on a second pass through the sensor tree.

2. The register-range limit

The BIOS SmartFan implementation sets enable=5 and temp_sel=8 for all PWM channels, binding them to PECI Agent 0 (the CPU package). The nct6798 kernel driver clamps temp_sel at 12. TSI0 is sensor 13. Memory temperature is therefore structurally unreachable by hardware fan control: not a misconfiguration, not a bad curve, but a register-range limit one index short of the sensor that matters.

Under memory-bound load the consequence is predictable. The CPU package sits in the mid-30s to mid-40s; the fans, reading a comfortable CPU, hold roughly 31 % PWM; the DIMMs climb past 80 °C with no corrective action.

3. Five stress runs

All runs used stress-ng --vm 8 --vm-bytes 30G --vm-keep, a 240 GB working set, with five minutes of stress followed by a two-minute cool-down. A safety abort triggered if the DIMM channel exceeded the danger threshold. Sampling was via the nct6798 TSI0 channel at 30-second intervals.

Run	Fan control	Stress method	Outcome	Peak / abort
1	BIOS SmartFan (temp_sel=8, PECI)	--vm-method all	Completed 5 min, MARGINAL	82.7 °C at 90 s
2	Manual lock, pwm1–5 at 55 % (140/255)	--vm-method all	ABORT	87.5 °C at 150 s
3	BIOS SmartFan (clean retest)	--vm-method all	ABORT	85.6 °C at 150 s
4	dimm-fanctl (TSI0 source)	--vm-method all	ABORT	88.3 °C at 150 s
5	dimm-fanctl (TSI0 source)	--vm-method write64	Completed 5 min, PASS	54.5 °C at 300 s

Run 1 completed, but only because --vm-method all cycles non-deterministically through lighter methods mid-test. TSI0 oscillated across roughly 51–83 °C; the fans held 31 % throughout. Run 3, an identical configuration, aborted at 85.6 °C. The difference was pure chance in the method sequence. A single passing thermal run proves nothing when the stress method is non-deterministic.

Run 2 was the intuitive intervention: lock the fans at a moderate duty cycle and see whether stable airflow helps. It produced the worst result of the SmartFan runs. The manual 55 % cap sat below what SmartFan's auto curve would have delivered during CPU-heavy phases (the curve reaches 100 % PWM at 70 °C on PECI). A fix in the wrong layer made things worse.

Run 4 introduced dimm-fanctl, a user-space daemon reading TSI0 every five seconds and mapping it to PWM. The fans correctly maxed at 100 %; fan2 (a Noctua Industrial 3000) reached

2528 RPM. The run still aborted at 88.3 °C after a 19 °C spike between two 30-second samples. The conclusion was about the instrument, not the machine: --vm-method all includes synthetic patterns that produce unreproducible thermal transients. It is unsuitable for validation.

Run 5 changed the method to write64, consistent sustained 64-bit memory writes. The thermal series:

Elapsed	TSI0	Tctl	CCD2	CCD3	CCD4	CPUTIN
30 s	49.2	49.1	47.0	40.1	41.0	38.0
60 s	51.5	51.3	48.3	42.1	42.0	39.5
90 s	52.3	52.2	49.6	42.6	42.3	41.0
120 s	53.2	53.2	50.5	43.1	42.8	42.0
150 s	53.2	53.1	50.5	43.1	43.2	43.0
180 s	53.7	53.3	47.7	43.5	43.7	43.5
210 s	53.8	53.8	50.6	43.5	43.7	44.0
240 s	54.0	54.0	48.7	44.0	43.3	44.5
270 s	53.7	53.7	49.6	44.0	43.8	44.5
300 s	54.5	56.6	61.8	44.0	43.8	45.0

The DIMMs plateaued at 53–54 °C from 120 s onward. Fans held roughly 55 % PWM at plateau. Cool-down returned to within 1.2 °C of the 40.5 °C baseline by 120 s. Peak was 54.5 °C, 26 °C below the 80 °C acceptance threshold and 30.5 °C below the 85 °C figure used elsewhere in the source material. All acceptance criteria passed.

4. The daemon

dimm-fanctl is installed as a systemd service, dimm-fanctl.service. It polls TSI0_TEMP every five seconds and maps the reading to PWM on a linear curve:

TSI0	PWM (0–255)	Fan %
≤ 35 °C	51	20 %
50 °C	127	50 %
65 °C	204	80 %
≥ 70 °C	255	100 %

A warning logs at 82 °C, a critical alert at 90 °C. On service stop the daemon restores BIOS SmartFan control; it is enabled for boot persistence. An unrelated node probe on 20 February 2026 showed the service still running in the standing set alongside Docker, NFS, promtail and node-exporter. The workaround became permanent infrastructure, not a test artefact.

Open follow-ups remain: version-control the daemon rather than leaving it hand-installed; check whether other nodes on the same board family need it; retire it if a BIOS release ever exposes TSI0 to SmartFan. One caveat belongs here rather than at the end: if the daemon dies, the DIMMs are unprotected again, and nothing in these measurements establishes what happens on daemon failure as distinct from a clean service stop.

5. Why the CPU-side validation missed this

Earlier the same day, before the DIMM work, a post-rack validation ran stress-ng matrixprod on all 64 threads for 30 seconds. Idle Tctl was 40.0 °C; under load it climbed to 83.1 °C peak at 25 seconds, roughly 12 °C below the 95 °C Tjmax. A 30-second memory stress (8 workers at 3.12 GiB each) peaked at Tctl 85.0 °C. Cool-down returned Tctl to 40.6 °C in 30 seconds. The chassis was judged excellent.

That judgement was reached with no memory-temperature data in hand at all. The CPU was fine; the fans were responsive; the sensors the operator checked were in range. The DIMMs were simply not in the picture, and nothing prompted the question.

6. Secondary study: open-air GPU recirculation

Workstation B is an ASUS Pro WS WRX90E-SAGE SE with an AMD Threadripper PRO 9975WX (32 cores, 64 threads, Zen 5), 503 GB DDR5 across eight DIMMs, two NVIDIA RTX PRO 6000 Blackwell 98 GB cards, and two 3.6 TB NVMe drives. The GPUs are open-air coolers, not blowers: they dump heat into the case rather than out of the rear bracket.

On 23 January 2026 a three-minute combined stress (64 CPU workers, 8 VM workers at 32 GB each, gpu-burn on both cards at 600 W, 4 disk workers) produced the following peaks: CPU Tctl 82.4 °C; GPU 0 (OEM part 900-5G144-2200-000, PCIe slot 1) 83 °C; GPU 1 (retail part 900-5G144-2500-000, PCIe slot 4) 92 °C. The 9 °C gap between nominally identical cards is confounded by their different part numbers and cooling solutions, so slot position and recirculation cannot be cleanly separated from hardware difference.

DIMMs A1–D1, closer to the GPU stack, rose from 43–49 °C to 62–66 °C, above the 55 °C "high" threshold. DIMMs E1–H1 stayed 34–36 °C at idle and reached only 42–45 °C under load. All subsystems passed.

On 13 February 2026, after swapping the 360 mm AIO's stock fans for three Noctua NF-A12x25 G2 PWM 120 mm units, a 120-second CPU stress peaked at 71.5 °C with 14,588,486 total bogo ops and 64/64 workers passing. The operator confirmed the fans inaudible under full load.

On 16 February 2026 the recirculation study proper ran two 120-second tests. The thermal throttle limit for both GPUs was confirmed at 93 °C via T.Limit headroom. All four GPU fans were at a 30 % minimum (roughly 1200 RPM) over a 30–100 % range, with GPUFanControlState = 0 (NVIDIA auto).

Test 1 (closed case, NVIDIA auto fans): CPU Tctl peaked at 87.8 °C against 73.9 °C on a CPU-only 60-second run, the difference being roughly 1200 W of GPU heat in the chassis. GPU 0 reached 82 °C at 600 W. GPU 1 reached 93 °C at 600 W and thermally throttled at the limit. DIMM D2:51 rose from 50.5 °C to 55.5 °C. The Mellanox CX-6 dropped from 61 °C to 57 °C under load, presumably from improved airflow. NVMe drives rose by 2 °C. The run completed 7200 GPU iterations.

Custom fan curves were then added via nvidia-settings (Coolbits 12): GPU 0 from 33 % idle to 100 % at 82 °C; GPU 1 from 35 % idle to 100 % at 78 °C, deliberately 5–8 °C earlier than GPU 0. Per-GPU hysteresis was set with instant ramp-up and ramp-down after 15 seconds sustained. Existing chassis-fan logic was preserved.

Test 2 (side panel removed, custom curve):

Sensor	Test 1 (closed / auto)	Test 2 (open / custom)	Delta
GPU 0 peak	82 °C	75 °C	−7 °C
GPU 1 peak	93 °C (throttling)	87 °C (no throttle)	−6 °C
CPU Tctl peak	87.8 °C	86.7 °C	−1.1 °C
DIMM D2:51	55.5 °C	56.5 °C	+1 °C

Sensor	Test 1 (closed / auto)	Test 2 (open / custom)	Delta
GPU iterations	7200	7250	+0.7 %

GPU 1 throttling was eliminated, with 6 °C of headroom below the limit under worst-case synthetic load. GPU 0 headroom rose from 11 °C to 18 °C.

The confound must be stated exactly: the side panel was removed and the fan curve was changed between Test 1 and Test 2. The 7 °C and 6 °C improvements cannot be attributed to either change alone. Open-case operation was adopted as the recommended configuration for this dual open-air-GPU setup, with periodic dust cleaning required.

Two further instrumentation gaps appeared on Workstation B. Only five of the eight DIMM SPD5118 sensors report through the kernel driver; bus 1 addresses 0x51 and 0x53 return empty. Fan6 is a 3-pin DC fan on a PWM-only header, running at a fixed roughly 2500 RPM regardless of duty cycle.

7. Limits of the evidence

The primary finding (the register-range limit) rests on a single board, a single kernel version, and five stress runs with n = 1 per configuration. The register clamp was observed in the nct6798 driver; whether other boards using the same Super I/O exhibit the same limit is not established.

TSI0 is an aggregate channel. The "85 °C+" figure is an aggregate, not a per-module reading, and the hottest module may be hotter still. No per-DIMM data exists for Workstation A.

All thermal numbers come from synthetic stress. Real workloads on Workstation B were expected in the 50–65 °C GPU range against the 75–87 °C synthetic peaks reported here. The synthetic results bound the problem; they do not characterise production behaviour.

The Workstation B comparison changes two variables at once (side panel and fan curve). The improvement is real; the attribution is not clean.

The daemon fix is user-space. If dimm-fanctl dies, the DIMMs lose their only thermal protection. The measurements do not establish what happens on daemon failure as distinct from a clean service stop. The daemon's five-second poll is faster than the 30-second sampling used in the stress runs, but transients inside that window remain unobserved.

Sampling at 30-second intervals missed the 19 °C spike in Run 4 until after the abort. The measurement cadence is coarse relative to the thermal dynamics of --vm-method all.

8. The methodological result

The method change from --vm-method all to --vm-method write64 is not a footnote. --vm-method all cycles non-deterministically through dozens of patterns, some of which produce unreproducible thermal transients. Run 1 completed; Run 3, identical, aborted. The difference was luck in the method sequence. A thermal validation that passes once and fails on retest is not a validation.

write64 produces consistent sustained 64-bit memory writes. It is reproducible, representative of heavy memory-write workloads, and does not mask a failure behind a lucky method cycle. The thermal series in Run 5 shows a clean plateau, a predictable cool-down, and a result that would replicate.

9. What to audit

Any thermal control loop should be checked for what its selector can address, and not only for whether its curve is sensible. The question is not "does the fan ramp up when things get hot?" but "can the controller see the thing that gets hot?" On this board the answer was no, and every check that read temperatures rather than examining the selector would have missed it.

Every hardware-managed cooling domain deserves the question: which sensor drives this, and is the one I care about in range?

PureTensor operates a sovereign AI infrastructure fleet (on-premises GPU compute, storage, and serving) run day-to-day by autonomous agents under human direction. PT-TN-2026-004. Measurements: 9 February 2026 (primary); 23 January, 13 February and 16 February 2026 (secondary thermal series on Workstation B). Raw artefacts are retained.