Layer1Labs Silicon

LAYER1LABS SILICON INC. · DOCUMENT

ChronoHive Measured Results

What ChronoHive has measured, how each figure was obtained, and the limits published alongside it.

Version 1.0 · October 26, 2026 · Layer1Labs Silicon Inc.

This document collects the results ChronoHive has published, in one place, with the labels the measurements carry on the Layer1Labs website. ChronoHive is a deterministic execution engine for AI and data infrastructure: work is written as programs, compiled into a single specification, and executed so that the same program with the same inputs makes the same decisions, bit-for-bit, on every run.

Every figure here is printed with the basis it was measured on — the workload, the hardware, and the regime — because a speedup figure without its regime is not a result an evaluator can use. The misses are here too: the regimes where ChronoHive currently trails the conventional approach are stated in Section 7, in the same terms the website's FAQ states them.

See it live: the ChronoHive page these results are published on — `https://layer1labs.ai/chronohive/`

Documentation library: this document is published in the Layer1Labs documentation library — `https://layer1labs.ai/docs/`

Figures in this document are drawn from the results published at layer1labs.ai/chronohive/ and the limits published at layer1labs.ai/faq/, both read on October 26, 2026. Each figure is cited to its published basis where it appears.

© 2026 Layer1Labs Silicon Inc. All rights reserved. See the Notices at the end of this document.

Contents

1 How to read these results 3

2 The headline result 3

3 Results by storage regime 3

4 Determinism, integrity, and checkpointing 4

5 Admission control, and a contention demonstration 5

6 What is modeled, not measured 6

7 Limits, stated plainly 6

8 What these results do not establish 6

Colophon 7

Notices 8

1 How to read these results

The website labels every figure with how it was obtained, and this document keeps those labels attached to the figures:

- **Measured** — the figure comes from runs that were actually executed, on the hardware and in the regime stated with it.
- **Preregistered** — the test design (methodology, arms, and success criteria) was fixed in writing before any run. Most of the campaigns behind this document are preregistered; the site's results rest on 501 completed preregistered trials on real GPUs, and an internal red team recomputed every headline figure before publication.
- **Modeled** — the figure comes from a model or a sensitivity calculation, not from a run. Modeled figures are labeled **MODELED** wherever they appear (Section 6), and are never presented as speedups achieved in production.
- **Demonstrated** — the figure comes from a staged demonstration of the policy being enforced (Section 5), not from a benchmark campaign.

Three standing rules govern everything that follows. First, **the envelope**: all measurements in this document are single-node, on at most 2 GPUs. No multi-node or cluster measurement exists yet — any cluster figure in Layer1Labs materials is a labeled scenario or model, never a measurement. Second, **no stacking**: figures come from separate preregistered campaigns and are never stacked or merged — a gain measured in one regime says nothing by itself about another. Third, **the basis travels with the figure**: where a result was measured on a particular workload, that workload is stated every time the result is stated.

2 The headline result

+9.8% — end-to-end training speedup, measured on real data.

Preregistered A/B trial (ResNet-18 on CIFAR-10): six interleaved trials of 200 steps, run on RTX 5060 Ti and RTX 4070 SUPER graphics processors. 9.8% faster end-to-end on both GPUs — just under our own 10% confirmation bar, and reported exactly as measured.

This is the result the ChronoHive page leads with, and it is the right one to lead with: it was preregistered, it ran on real data, and it is reported with its shortfall included. The trial's design was fixed in writing before any run — ResNet-18 on CIFAR-10, six interleaved trials of 200 steps each, run on RTX 5060 Ti and RTX 4070 SUPER graphics processors — and the measured improvement landed at +9.8% on both GPUs, just under the 10% confirmation bar the trial set for itself. It is reported exactly as measured, not rounded up to the bar.

Checkpoint integrity held alongside the speedup in the same body of work: 400 of 400 checkpoints were read back and verified, and final training states were bit-identical in 20 of 20 trial groups (Section 4).

3 Results by storage regime

Regime decides everything. The same coordination that recovers a large share of wall-clock time when storage is slow and shared recovers a smaller share when storage is fast and local, because there is less stall to recover. The website publishes the regimes separately, and so does this document.

© 2026 Layer1Labs Silicon Inc.

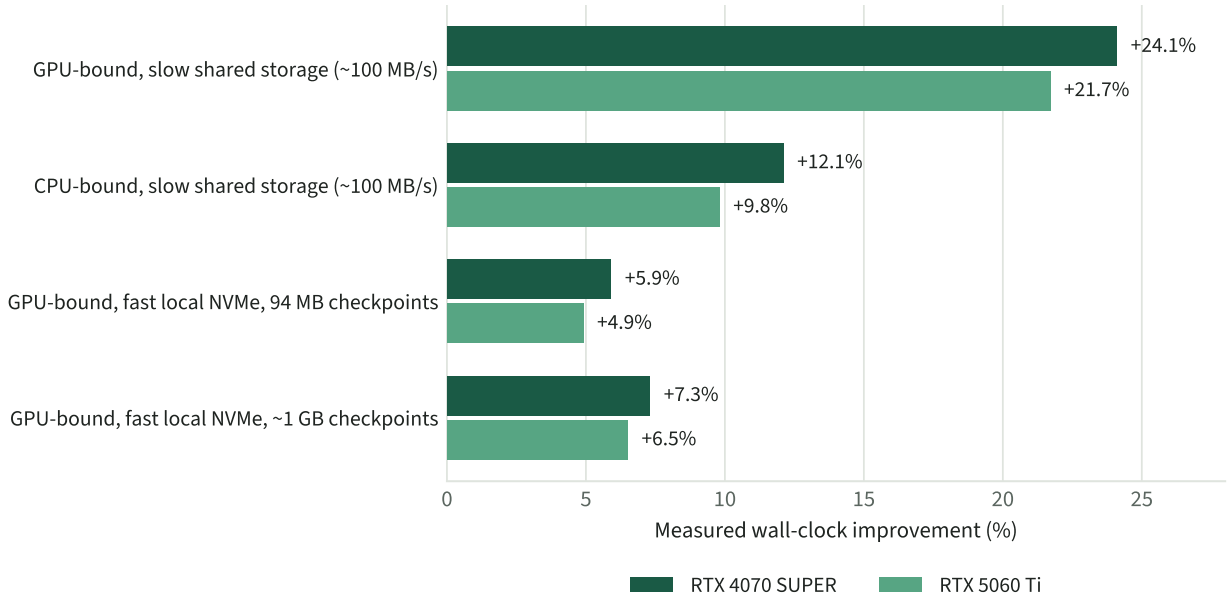


Figure 1: Measured wall-clock improvement by storage regime and workload. Separate preregistered campaigns; figures are never stacked or merged. Source: layer1labs.ai/chronohive/, "Measured Results."

Regime (all Measured)	4070 SUPER	5060 Ti
GPU-bound training, slow shared storage (about 100 MB/s)	+24.1%	+21.7%
CPU-bound workload, same slow shared-storage regime	+12.1%	+9.8%
GPU-bound training, fast local NVMe, 94 MB checkpoints	+5.9%	+4.9%
GPU-bound training, fast local NVMe, ~1 GB checkpoints	+7.3%	+6.5%

Two further measured results sit outside the regime table. Two-rank distributed training (DDP, the standard PyTorch multi-process training mode) measured **+1.7%** end-to-end job-time improvement. And the largest gains in the table — GPU-bound training over slow, shared storage — are the regime the FAQ identifies as the one where production clusters hurt most.

4 Determinism, integrity, and checkpointing

A speedup is worth little if the run cannot be trusted or repeated. The determinism and integrity results are measured with the same discipline as the speed results:

- **92 of 92 runs.** In the determinism campaign across two orchestrator harnesses, 46 of 46 runs in each reproduced exactly — zero duplicate writes, and all 460 stage outputs matched the reference fingerprint chain. Determinism here is measured, not just designed.
- **400 of 400 checkpoints** read back and verified; final training states bit-identical in **20 of 20** trial groups.

© 2026 Layer1Labs Silicon Inc.

- **Checkpoint stalls cut.** In the large-checkpoint campaign, the per-run checkpoint stall fell from 10.5–11.7 s to 0.3–1.5 s.
- **Bounded memory.** The asynchronous checkpoint writer's memory footprint is bounded at **3× the checkpoint size** — measured at every capture across 40 preregistered isolated trials. The previous transport reached ∼7× at ∼4 GB checkpoints.

Layer1Labs also holds a formal proof of the engine's determinism — that its decisions are exact by construction, with a state small enough to enforce in hardware. The proof itself is unpublished; it is available to qualified partners and investors under NDA. It is mentioned here, and only mentioned, in the same terms the website uses.

5 Admission control, and a contention demonstration

ChronoHive's admission control decides which work gets storage bandwidth, and when — against real capacity, rather than by contention. It was measured in the same preregistered A/B test as the headline result, on three platforms:

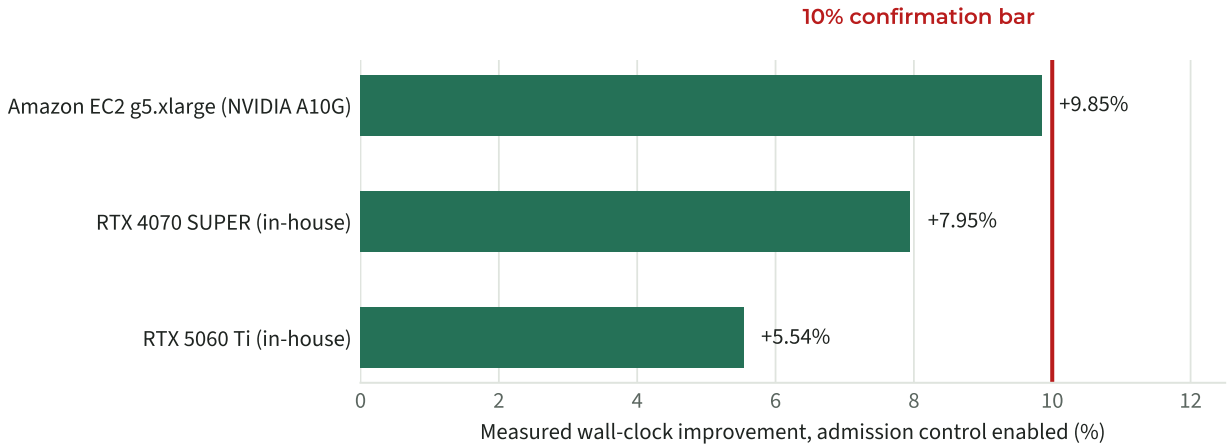


Figure 2: Measured wall-clock improvement with admission control enabled, by platform: ten interleaved trials of 200 steps on real data per platform (22 completed preregistered trials in the campaign, per the FAQ). All three platforms finished under the 10% confirmation bar. Source: layer1labs.ai/chronohive/ and layer1labs.ai/faq/.

The reading of this result is the one the website gives it: **positive on every platform, and under the 10% confirmation bar on each — a positive observation, not a proof.** The cloud instance (Amazon EC2 g5.xlarge, NVIDIA A10G) measured +9.85%; the in-house RTX 4070 SUPER measured +7.95% and the RTX 5060 Ti +5.54% in the same test.

Separately, the contention demonstration shows the policy itself being enforced (**Demonstrated**, not a benchmark): an unmanaged baseline admitted 480 operations and refused none. The coordinated policy admitted **401**, refused 119, retried 40, and dropped 79 — the programmed policy, enforced deterministically. Refusing work is the point of admission control; the demonstration exists to show the refusals happening exactly as programmed.

© 2026 Layer1Labs Silicon Inc.

6 What is modeled, not measured

Two figures in Layer1Labs materials are models. They are printed here, in this section only, with their labels on:

- **5–15% throughput uplift — MODELED.** A sensitivity range for pilot planning. It is not a production speedup, and no measured figure in this document should be read as confirming it.
- **Cross-regime median +12.5% — MODELED.** The FAQ's scalability model — a simulation, not a result — suggests a cross-regime median of +12.5%. It is an expectation to be tested by the multi-node validation described in Section 7, not a measurement.

The FAQ also carries an investor illustration of the same character: investor materials model roughly $2.7M per year for one 1,024-H100 cluster *if* a 5% end-to-end gain materializes, using published cloud GPU pricing. It is arithmetic on stated assumptions, not a result, and it is included here only so that no modeled figure circulates without its label.

7 Limits, stated plainly

The FAQ publishes the misses, and this document carries them in full:

- **CPU-bound workloads on fast NVMe trail.** CPU-bound workloads on fast NVMe still trail synchronous checkpointing by **6–14%** in Layer1Labs' measurements.
- **Very large checkpoints at short intervals are the hardest measured case.** At $\sim$4 GB checkpoints written at $\sim$3.5 s intervals, the best configuration — a coalescing writer whose memory stays bounded at 3$\times$ checkpoint size — still trails synchronous checkpointing by **about 32%**, because encoding the checkpoint takes $\sim$10 s against a $\sim$3.5 s interval. The limit is encode speed, not memory.
- **The admission-control result is below its own bar.** Positive on every platform measured, but under the 10% confirmation bar on each (Section 5) — reported as an observation, not a proof.
- **Nothing here is a cluster result.** Every measurement behind this document is single-node, on at most 2 GPUs. No multi-node or cluster measurement exists yet. The funded next step is the real test: multi-node validation from 8 to 32 GPUs on shared storage.

8 What these results do not establish

These results do not establish that ChronoHive will speed up any particular customer's workload: the measured gains are regime-dependent, and two regimes in this document are negative. They do not establish cluster-scale behavior, which has never been measured. The modeled range in Section 6 does not establish a production speedup, and the admission-control campaign does not establish its effect to the standard of its own confirmation bar.

What they do establish is narrower and, for an evaluator, more useful: on real data, in preregistered tests whose designs were fixed before the runs, coordinated admission and deterministic execution recovered measured wall-clock time in the regimes where storage contention is the bottleneck — and reproduced every run exactly while doing it. A pilot is how the remaining question — what happens on your

© 2026 Layer1Labs Silicon Inc.

cluster, at your scale — gets answered: the methodology, arms, and success criteria are fixed in writing before any run, and the deliverable is a joint benchmark report with reproducible artifacts.

A companion document, *The Traffic Demo, In Depth*, in this documentation library walks through the public demonstration of the same compile-and-execute architecture — a Lingua Franca program timing a traffic grid, live on the website — with its own measurements. Its results are stated there and are not repeated here.

Colophon

Sources: the “Measured Results” and supporting sections of `https://layer1labs.ai/chronohive/` and the technical-evaluator answers at `https://layer1labs.ai/faq/`, both fetched live on October 26, 2026, and cross-checked against the site’s production source. No other source was used. Figures 1 and 2 are generated by script from the published values. Build: Version 1.0, October 26, 2026; LuaLaTeX, set in the Layer1Labs document family; no figure in this document was hand-drawn or hand-typed into its chart.

© 2026 Layer1Labs Silicon Inc.

Notices

Copyright. © 2026 Layer1Labs Silicon Inc. All rights reserved.

Sharing. This document may be shared in its complete, unmodified form, with attribution to Layer1Labs Silicon Inc. No other reproduction, modification, adaptation, or derivative use — in whole or in part — is permitted without the prior written permission of Layer1Labs Silicon Inc.

No license. Nothing in this document grants, by implication or otherwise, any license or right under any patent, copyright, trademark, or other intellectual property of Layer1Labs Silicon Inc. Publication of this document is a disclosure of information only.

Patents. The technology described in this document is the subject of pending patent applications, including 9 provisional patent applications filed with the U.S. Patent and Trademark Office. No statement in this document is a representation as to the scope, validity, or enforceability of any application or of any patent that may issue from one.

Trademarks. Layer1Labs, ChronoHive, and ChronoFabric are trademarks of Layer1Labs Silicon Inc. Other names mentioned in this document may be trademarks of their respective owners; their use here is for identification only and claims no affiliation or endorsement. In particular, Lingua Franca is an open-source research project of its own community.

Measurements and basis. Figures in this document are drawn from the project's measurement records and are cited where they appear; where a figure is modeled rather than measured, it is labeled as modeled. This document is provided for information only, on an "as is" basis, without warranty of any kind, express or implied, including as to accuracy, completeness, or fitness for a particular purpose. It describes a demonstration system; it is not a specification, an offer, or a forecast of results in any other setting.

Current version. Documents in this family are revised in place under version numbers. The current version of this document is the one published at layer1labs.ai; earlier versions are superseded.