Layer1Labs Silicon

LAYER1LABS SILICON INC. · GUIDE

Evaluating ChronoHive: The Pilot Process

How an evaluation runs, as published — first contact, a jointly preregistered pilot, and a joint benchmark report.

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

This guide walks through evaluating ChronoHive, the deterministic execution engine from Layer1Labs, exactly as Layer1Labs publishes it: how to make first contact, how a pilot is scoped, how its success criteria are fixed in writing before anything is measured, what you receive at the end, and how licensing follows. It also collects, in one place, the limits Layer1Labs publishes alongside its results — because a pilot designed against honest limits is the only kind worth running.

Everything in this guide comes from the published pages. Where the published material does not state a term — a timeline, a cost, a condition — this guide does not invent one. Section 6 lists those open points as the questions to settle, in writing, during scoping.

See it live: The results and terms this guide refers to — `https://layer1labs.ai/chronohive/`

See it live: This guide in the Layer1Labs Docs catalog — `https://layer1labs.ai/docs/`

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

Contents

1 What an evaluation is **3**

2 First contact **3**

3 Scoping against the published envelope **3**

4 The preregistered pilot **3**
4.1 The basis a pilot reproduces or refutes 4

5 What you receive, and what follows **5**

6 Questions to settle in writing before the pilot **5**

7 The published limits **6**

Colophon **6**

Notices **7**

1 What an evaluation is

ChronoHive is proprietary technology of Layer1Labs Silicon Inc., offered for commercial licensing by written agreement. For qualified teams, an evaluation path is available before any license: a pilot, run jointly, that tests ChronoHive against criteria both sides fix in writing before the first run.
Two properties of the pilot are published and worth underlining at the start. Pilots are **isolated** — they run apart from your production work — and **software-first**: ChronoHive runs on the clusters you already own, and no hardware purchase is involved at any point in an evaluation.

2 First contact

Evaluation starts with a conversation. The published routes are:

Route	Address
Partnerships & Licensing	licensing@layer1labs.ai
Technical & Evaluation	info@layer1labs.ai

Use the licensing address for the engagement itself and the technical address for questions about fit, integration, and the measurement record. One thing to ask for early: **the methodology and trial data behind the published results are available on request.** An evaluator who reads the methodology before scoping will write better success criteria in Section 4's step — the criteria are only as good as the measurement design they name.

3 Scoping against the published envelope

Scoping fits your workload to what has actually been measured. The published envelope is explicit, and a pilot scoped outside it would be testing something Layer1Labs has not claimed:

> **The published measurement envelope.** Every published ChronoHive result is 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. Results come from separate preregistered campaigns and are never stacked or merged.

Inside that envelope, the published record spans storage regimes an evaluator can match to their own: slow, shared storage (about 100 MB/s), and fast local NVMe with checkpoints of about 94 MB and about 1 GB. The full results, each with its label and basis, are collected in the companion document *ChronoHive Measured Results*; the headline basis is summarized in Section 4 below, because it is the basis a pilot sets out to reproduce or refute.
Integration scoping is short, by design. You keep your infrastructure: work is authored as Lingua Franca programs, compiled by chronoc into a single constraint specification, and their connection points bind to your storage and compute through software callables. There is no hardware purchase and no platform migration in an evaluation.

4 The preregistered pilot

The pilot follows a fixed pattern, shown in Figure 1. Its defining feature is the order of operations: **validation is preregistered jointly — fixed in writing before any run.** Before anything is measured, both sides agree and write down:

© 2026 Layer1Labs Silicon Inc.

- **Methodology** — how the measurement will be run;
- **Arms** — what runs against what (for example, your current setup against ChronoHive, interleaved);
- **Success criteria** — the numbers that will count as success, named before the data exists.

Only then does the measurement run. The verdict is read against the criteria that were fixed in advance — not against a bar chosen afterward.

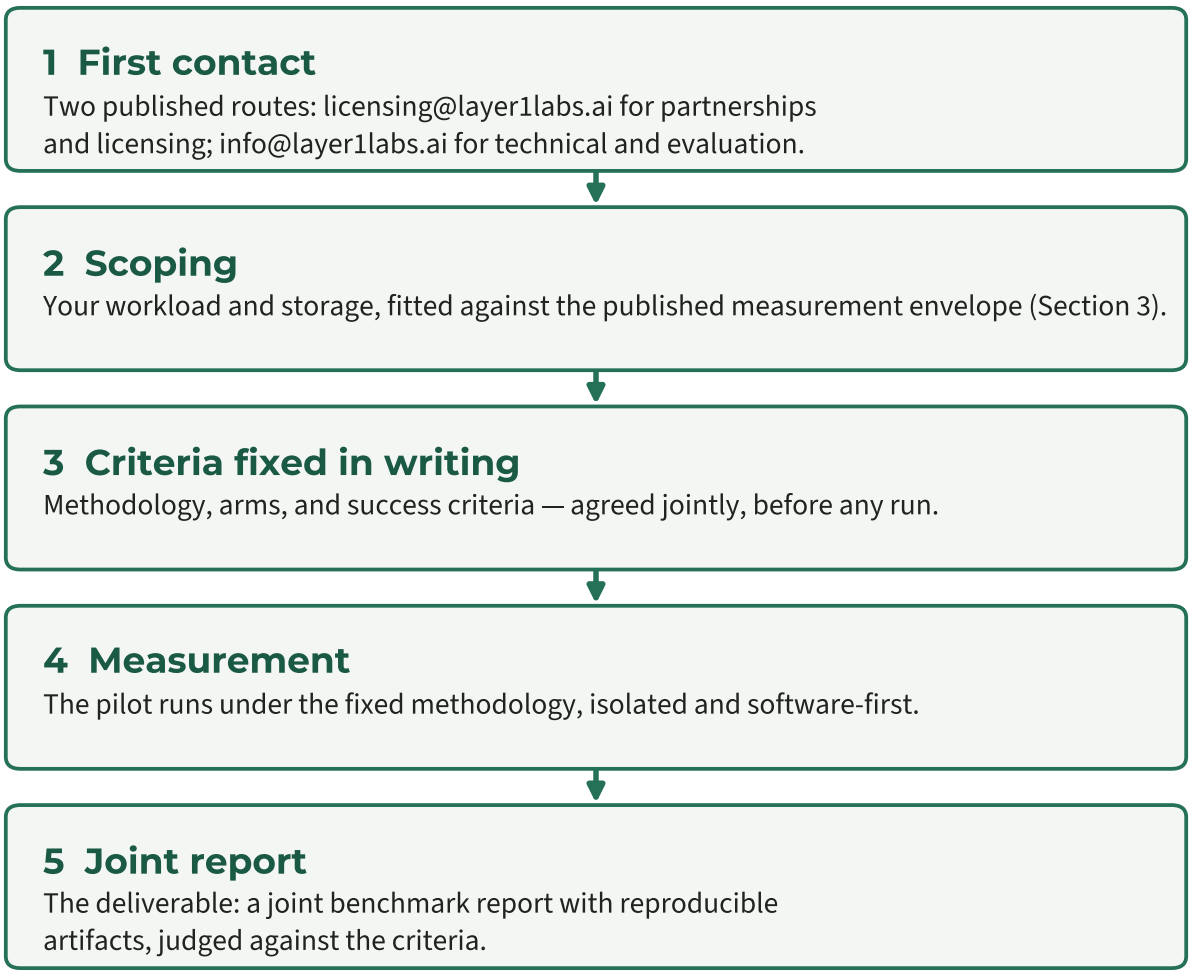


Figure 1: The pilot process as published: contact, scoping, criteria fixed in writing, measurement, and the joint report. Source: layer1labs.ai/chronohive/, “Engage.”

4.1 The basis a pilot reproduces or refutes

A pilot is meaningful because the published results were produced the same way — preregistered, on real data, reported exactly as measured. The headline an evaluator is really testing:

+9.8% end-to-end training speedup — Measured. Preregistered A/B trial on real data (ResNet-18 on CIFAR-10): six interleaved trials of 200 steps, run on RTX 5060 Ti and RTX 4070 SUPER graphics processors. The result came in just under the trial's own 10% confirmation bar, and is reported exactly as measured.

That last sentence is the house standard a pilot inherits: behind the published figures stand 501 completed preregistered trials on real GPUs, and an internal red team recomputed every headline figure before publication. A figure that lands just under its bar is published as it landed. Your pilot's report follows the same rule — which is what makes a pilot worth running at all: the verdict, favorable or not, is one both sides agreed in advance to be bound by.

One range appears in the published material with a different label, and an evaluator should keep it in its place: a **5–15% throughput uplift is labeled Modeled** — a sensitivity range for pilot planning, expressly not a production speedup. It is a reasonable planning assumption for scoping a pilot. It is not a result a pilot is expected to "confirm"; the pilot measures your workload, in your regime.

5 What you receive, and what follows

The deliverable of the pilot is **a joint benchmark report with reproducible artifacts**: the results, the methodology they were produced under, and the artifacts another team would need to reproduce them. Because the criteria were fixed before the runs, the report's verdict needs no interpretation — it is read directly against the success criteria both sides signed.

If the evaluation leads to adoption, ChronoHive is licensed commercially **by written agreement**. That is the published route in full: evaluation by joint pilot, licensing by written agreement. The published material does not state license pricing, term lengths, or support terms — those are settled in the agreement itself, and Section 6 lists the points to raise.

6 Questions to settle in writing before the pilot

A careful evaluator closes the gaps the published material leaves open, during scoping and in writing. The published pages do not state:

- **Timeline** — how long scoping, the pilot runs, and the joint report take, start to finish.
- **Cost** — whether any fee attaches to the pilot itself, and what each side supplies.
- **Where the pilot runs** — whose cluster, whose accounts, and what access each side needs.
- **Confidentiality and data** — whether an agreement is required before the pilot, how your workloads and data are handled during it, and what happens to pilot data and artifacts afterward, beyond the joint report itself. (One published exception: the formal determinism proof Layer1Labs holds is unpublished and available to qualified partners and investors under NDA — if your evaluation needs it, raise it here.)
- **Support during the pilot** — who responds when a run fails, and how quickly.
- **The license to follow** — the pricing, term, and support shape of the written agreement a successful pilot leads to, so the pilot's success criteria can be judged against the real cost of adoption.

None of these gaps is hidden; they are simply matters for the agreement rather than the website. Settling them in writing before the pilot is the same discipline as the preregistration itself: nothing important gets decided after the data is in.

© 2026 Layer1Labs Silicon Inc.

7 The published limits

Layer1Labs publishes the limits alongside the results. A pilot scoped with these in view will not be surprised:

- **CPU-bound workloads on fast NVMe** still trail synchronous checkpointing by 6–14% in Layer1Labs' measurements.
- **Very large checkpoints are the hardest measured case.** Checkpoints of about 4 GB written at short intervals (about 3.5 s) still trail synchronous checkpointing by about 32% in the best configuration — a coalescing writer whose memory stays bounded at 3× checkpoint size. The limit is encode speed (about 10 s to encode against a 3.5 s interval), not memory.
- **Admission gating is positive everywhere, and reported as an observation, not a proof.** Measured +9.85% on an Amazon EC2 g5.xlarge instance (NVIDIA A10G), +7.95% on RTX 4070 SUPER, and +5.54% on RTX 5060 Ti, across 22 completed preregistered trials — positive on every platform, and under the 10% confirmation bar on each.
- **The validated scope is single-node, at most 2 GPUs.** No multi-node or cluster measurement exists yet. A scalability model — a simulation, not a result — suggests a cross-regime median of +12.5%, labeled Modeled; the funded next step is the real test: multi-node validation from 8 to 32 GPUs on shared storage.

Colophon

Sources: the engagement terms, pilot pattern, and contacts as published at layer1labs.ai/chronohive/, and the limits as published at layer1labs.ai/faq/, read on October 26, 2026. Every figure in this guide is printed with the basis and label it carries on those pages. This is the first published edition of this guide (Version 1.0).
Build: LuaLaTeX (TeX Live 2026) on the Layer1Labs document template family. The figure is generated by script from the published steps; no figure value is hand-typed.

© 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.