ANNOUNCEMENT

TETRA AND HORNESCI

Bringing high-performance compute to business software. Today.

Storing more data makes software more expensive, and the problem gets worse the more you store. TETRA CA PBC and HorneSci have entered a cooperative agreement to change that. TETRA builds matrix data systems that hold the same dataset in a fraction of the space. HorneSci builds the sparse tensor mathematics that make matrix operations run faster on the same silicon. Together we deliver high-performance compute to ordinary business workloads.

Cooperative agreement · 2026-09-05
TETRA CA PBC · HorneSci

SECTION 01

THE COST OF SCALE

Why growth costs more than it should, and the three ways companies usually pay for it.

Software has a well-known problem: storing more data makes the software more expensive, and the more data you store the worse the problem gets. A company that wants to serve more customers has to scale, which means getting the software to do more work for more people.

The usual options

Scaling a business workload comes down to three engineering choices:

- **Optimize the code** that handles the workload.
- **Add hardware**, so the workload has more to run on.
- **Add software**, the caches and layers that make an existing system go further.

Each option buys a fixed amount of headroom and then needs doing again. Optimized code slows down once the data grows past what it was tuned for. Hardware runs out of the room it just bought. A cache layer becomes one more system to size, to operate and to keep correct. The cost of scaling therefore recurs, and it recurs faster the bigger you get.

TETRA CA PBC and HorneSci entered a cooperative agreement to break that pattern, and to reach levels of scalability and flexibility that those three options cannot.

> **HIGH-PERFORMANCE COMPUTE, NOW**
>
> HPC has belonged to national laboratories and research institutions for decades. We bring it to your business problems today.

SECTION 02

TETRA CA PBC

Matrix data systems, and what representing data as mathematics buys you.

TETRA CA PBC builds ultra high-performance data systems around our flagship product, TETRAbase. TETRAbase is a matrix database: it encodes data to and from a matrix.

That representation lets TETRA use mathematical expressions to store and retrieve information. The CA in TETRA CA stands for Collapsive Architecture, because we collapse the cost, the complexity and the time it takes to build, operate and scale your systems.

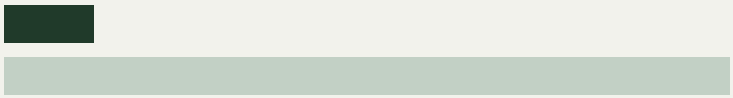
What the representation buys

The same dataset occupies roughly an eighth of the disk it takes on a traditional data system, and TETRAbase holds far less RAM while serving it. No datapoint needs an index. Everything in the system is instantly locatable, and the data stays queryable while compressed.

How you connect to it

TETRAbase offers two interfaces onto the same matrix: traditional SQL, and a labeled property graph that uses openCypher over the Bolt protocol. Both are industry standards, so the drivers, dashboards and client libraries your teams already run work against TETRAbase unchanged.

THE SAME DATASET ON DISK — 8.1x smaller



211.5 MiB TETRAbase

1.683 GiB the industry leader

The LDBC Social Network Benchmark at scale factor 1: 3.2 million nodes and 17.3 million edges, held by TETRAbase and by the leading open-source graph database. Lower is better. The full comparison is in TETRAbase Performance, LDBC SNB, 2026-09-03.

> **TWO FILES, NO INDEX**
>
> TETRAbase keeps a dataset in two files. One holds the data as loaded, and the other records every change since. It writes each value once and encrypts the whole store. There is no index to build after a load, to update on every write, or to rebuild after a restore.

SECTION 03

HORNESCI

Sparse tensor mathematics, and the people who build it.

HorneSci researches and implements deep high-performance compute technology. Their flagship product, Hierarchical Sparse Tensors, optimizes the throughput of matrix mathematical operations. Their targeted-workload benchmarks show orders-of-magnitude improvement over industry-standard HPC libraries.

Business data is sparse. Most of the entries in a matrix that describes real records hold nothing, and an operation that treats them as present pays for work that produces no answer. Sparse tensor methods organize the entries that do hold something, so the arithmetic touches only those.

Hunter Horne, the founder of HorneSci, has years of experience supporting contracted work inside the U.S. National Laboratories ecosystem.

What they contribute

Hunter and his team reduce the work a microchip performs on every calculation cycle, and they have breakthrough results. Their pedigree in HPC systems is what makes the business systems we deliver mathematically as efficient as they can be.

Reducing the work per cycle shows up twice on a bill. The query returns sooner, and the machine draws less power to return it. Across a fleet the second effect compounds with the first.

> WHY THE TWO FIT
>
> TETRA uses the same matrix mathematics that HorneSci optimizes. The two products meet at the level of the arithmetic, so a gain in the mathematics carries straight through to a running system.

SECTION 04

WHAT THE AGREEMENT DELIVERS

What customers on each side get out of the collaboration.

- **TETRA customers** get lower latencies and lower energy consumption, and both of those lower cost.
- **HorneSci customers** get TETRAbase as a general-purpose interface to Hierarchical Sparse Tensors.

Together the two organizations serve business customers with boutique, highly efficient HPC solutions.

Where it fits in your stack

The combined expertise lets you interface with or replace your current data storage and retrieval, your business logic APIs, and the auxiliary systems around them. TETRA runs alongside what you already have, and you plan a full cutover whenever you are ready. In the meantime your teams, your software and your agentic processes get HPC data access.

> ARCHITECTURAL COLLAPSE
>
> For the full version, TETRA hosts your data systems and you operate nothing. If you would rather host your own, business systems built on HPC need little operational support: you get the reliability without the extra headcount.

Who you are working with

TETRA CA PBC brings decades of enterprise software experience. HorneSci brings academic and research-grade discipline. If secure, consistent and scalable data systems matter to your organization, the two of us can help.

Signatories

	ROLE	CONTACT
K. "Mickey" Mickey	Founder and CTO, TETRA CA PBC	mickey@tetra-ca.com
Reid Johnson	Co-Founder and CEO, TETRA CA PBC	reid@tetra-ca.com
Hunter Horne	Founder, HorneSci	hunter@hornesci.com

TETRA CA PBC is a public benefit corporation. HorneSci researches and implements high-performance compute technology.

Start with your workload

The system that matters is the one running your data.

Tell us what you store, what you query and what it costs you today. We will tell you what the same workload looks like on a matrix system, and we will show you the measurements behind the answer.

TALK TO US

tetra-ca.com

inquiries@tetra-ca.com

TETRAbase Performance, LDBC SNB, 2026-09-03 carries the full head to head benchmark and the scaling results.

TETRA CA PBC