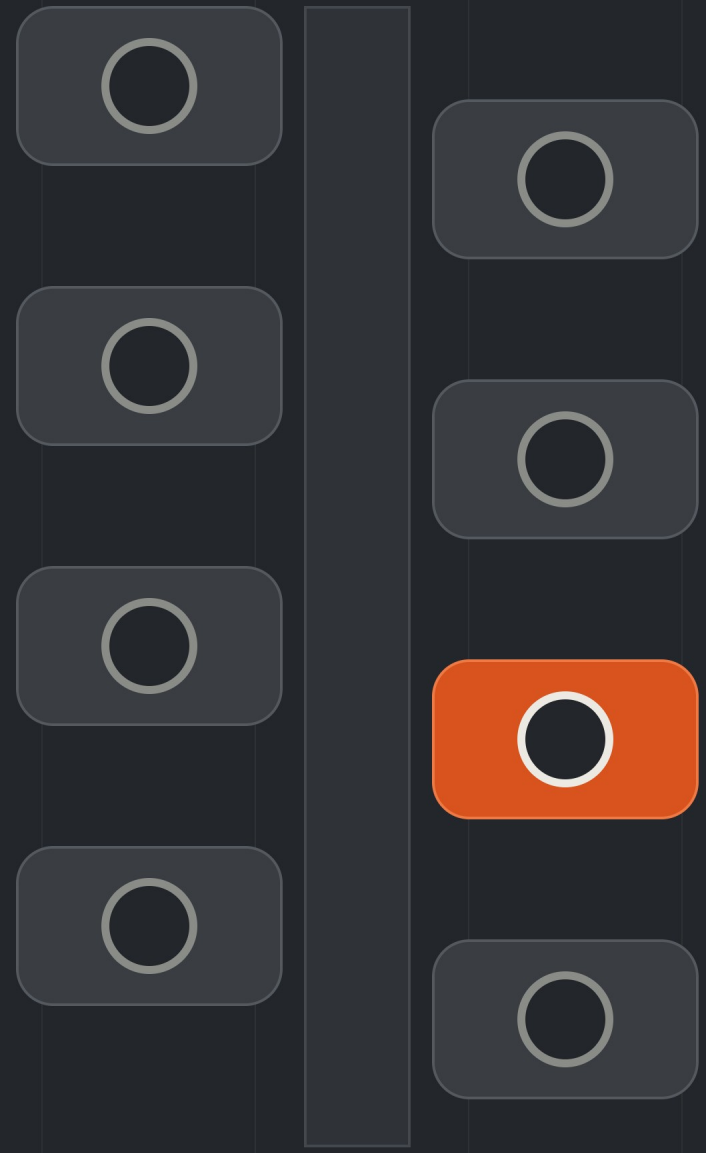
OSA CON 2026 · ROOM 2

The Biggest Dataset Nobody Modeled

Building an open semantic layer for telemetry with DuckDB and DataFusion

Clay Smith · Tero



How much did this {code change} cost or save the business?

The observability/monitoring pitch to business people has been about this for 15+ years. It never really worked out (except for eCommerce) for technical and organizational reasons.

Org issue: The answer, tools, and data are split across two teams that don’t eat lunch together.

Platform & SRE

- Deploys and feature flags
- Latency and errors
- Traces, logs, metrics

NO JOIN

Data & analytics

- Revenue and billing
- Usage and funnels
- Customers and plans

Technical issue: telemetry is the biggest dataset nobody modeled (for good reasons).

service.name	k8s.pod.name	scope.name	severity	body	other attrs
checkout	checkout-7d9f8-x2k4q	io.opentelemetry.jdbc	WARN	slow query 2.4s	+41
payments	payments-5c6b-q9z1m	otel-go/net/http	ERROR	card_declined	+37
cart	cart-8f2a-l0p3d	@opentelemetry/redis	INFO	cache miss	+29
ingest-gw	ingest-gw-66d1-r7t8n	otel-rust/tonic	ERROR	upstream timeout	+52
checkout	checkout-7d9f8-x2k4q	com.acme.billing	INFO	order.created	+18

Petabytes of metrics, logs and traces. Potentially thousands of log shapes or attributes.

Blocker 1, the tech, is fading fast in late 2026

01

OTel mostly won

Imperfect, but its conventions give you a free starting schema.

02

Open tables (and Clickhouse)

Iceberg or DuckLake: increasing production success stories for logs/traces.

03

Easier Maintenance

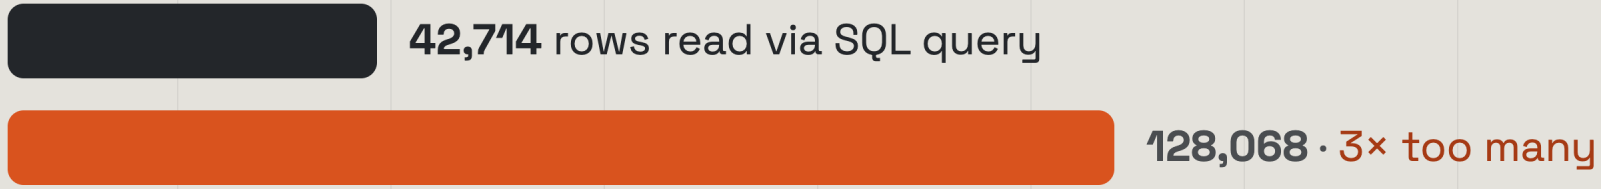
Compaction and snapshots getting easy buttons (or vendor-managed options).

OTEL IN DUCKDB IN 200 CHARACTERS.

```
INSTALL otlp FROM community;
LOAD otlp;
SELECT * FROM read_otlp_traces('traces.pb');
```

DuckLake: some trace query numbers

SPANS SINCE OCT 25 · SAME TABLE · THREE WAYS TO READ IT



Knows the current file

Rewrites leave old copies in the bucket. The catalog skips them.

Fresh rows, no tiny files

New spans wait as rows in Postgres until the next flush.

Maintenance built in

Compacts the table to one 17 MB file every ~10 minutes.

Reasonable 2026 options for your telemetry.

Option	Format	Maintenance	Worth knowing
Amazon S3 Tables	Iceberg	Auto compaction, snapshot expiry	CloudWatch Logs mirror in natively
Cloudflare R2 Data Catalog	Iceberg	Auto compaction, opt-in per table	Logpush → Pipelines → Iceberg
Google BigQuery Observability Analytics	BigQuery	Managed tables optimize storage	Logs and traces in BigQuery with a button click
MotherDuck	DuckLake	Managed DuckLake	
ClickHouse	ClickHouse	Automatic-ish	Catalogs in beta
Run it yourself	Any open table	All on you	Most control, most work
Build your own on DataFusion	Vortex or Parquet	All on you, plus the engine	LangChain's SmithDB: DataFusion + Vortex, 92 ms P50 trace loads

Sources: vendor docs and changelogs, 2026; LangChain SmithDB launch · 07 / 17

What's left is an organizational problem.

TEAM Platform & SRE	≠ LUNCH	TEAM Data & analytics
TOOL Legacy APM vendor	≠ BUDGET	TOOL data lake, duckdb, Datafusion, ClickHouse
DATA Vendor's closed store	≠ JOIN	DATA OTel, your bucket

Basically: Conway's Law applied to observability + analytics.

Observability + analytics: where we're headed

Load

OTel traces and billing data, side by side in DuckDB/Datafusion/Clickhouse

Agent asks question

An agent writes the join across telemetry and business tables

Answer or dashboard created

"This change cost $[__]"

Issue: Agents still struggle with real data work.

38%

Best frontier model's pass@1 on DataAgentBench: realistic, multi-database data tasks.

Where agents fail

- Integrating multiple databases
- Ill-formatted join keys
- Unstructured text
- Domain knowledge

Telemetry is messier than the data in DABench.

Agents need a map to understand and query observability data

Semantic layer	Vendor-neutral, built for agents, on open tables
Inferred from data	Cheap AI/ML models read and reason about telemetry data shapes and relationships as it is streamed in.
OTel conventions	OTel semantic conventions (tags) give consistency across infrastructure and applications.

`Tero benchmark: in progress` Like DAB's authors: we're seeing a significant benefit with an index of telemetry shapes.

Reasonable question: why not just use ClickHouse?

The model guesses wrong when...

- Inconsistent or existent tags
- Joins across telemetry silos
- Needle-in-the-haystack questions

Pick ClickHouse when...

- You primarily care about metrics.
- Very latency sensitive
- You already have a ClickHouse contract or team that maintains it.

Pick open tables when...

- Greenfield project
- Mostly logs or traces
- Have ~somewhat structured logs
- You have a provider that handles some of the maintenace
- **Agents are querying the data.**

[Your one-line answer]

Elephant in the room if you want to replace {legacy tool} with DuckLake/Iceberg: alerting

AN ALERT NEEDS DATA WITHIN

[~few s]

OUR OPEN-TABLE COMMITS LAND EVERY

[minutes]

DuckDB 2.0 + other approaches technically possible, just not much prior art. Likely a build-your-own-story (for now).

Your security team will likely be on board if you want to move data into object storage you own.

48%

of breaches involved a third party in the 2026 Verizon DBIR, up 60% from the year before.

“Black October,” Oct 17–22

- LiteLLM supply-chain attack stole AI API keys
- Two critical PHP remote-code bugs
- GitLab self-managed RCE
- Eight weak CA keys revoked

My estimate: getting a new vendor through security review gets 5–10× harder.

Sources: Verizon 2026 DBIR; public reports on “Black October” 2026 ·

With an open foundation and good agents, what is the observability vendor for?

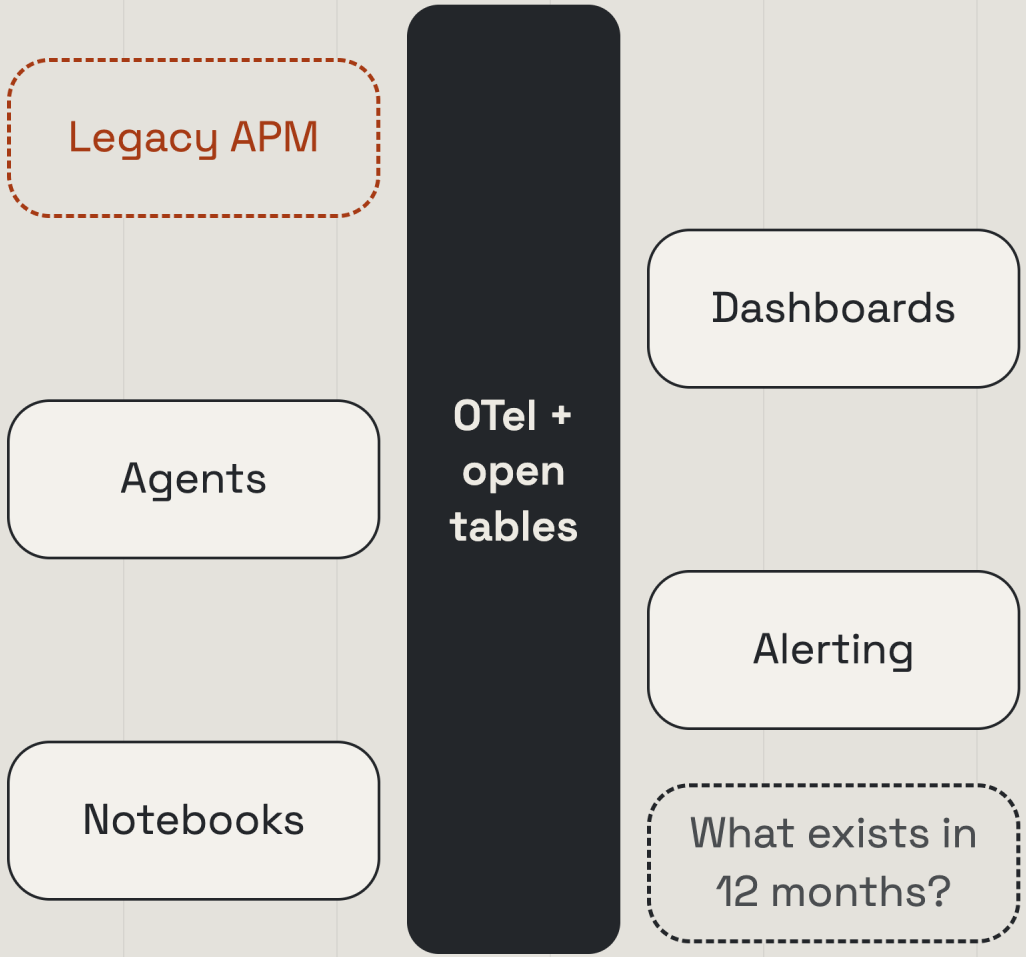
Job	Legacy APM	Open foundation + agents
Store	Vendor's cluster, priced per GB	Open tables in your own bucket
Query	Proprietary query language	SQL, joined with business data
Explore	Prebuilt dashboards	Agents build the chart you need
Alert	Mature and real-time	Still catching up (slide 13)

So what's left?

[Expensive monitoring tool] - something you can unplug in 2027?

Like a Metabolist tower: the core is permanent, and the capsules plug in and get replaced.

Modeled telemetry on open tables is the core. Every tool on top, including the APM, is optional.



05 · ONE ASK

Eat lunch with your SRE/DevOps/platform team and talk about data.

@smithclay - thanks!
linkedin.com/in/smithclay