Columnar OLAP vs. Schema-on-Read Indexing on Security Workloads

A reproducible benchmark on a 10M-event Zeek conn.log workload

Jeremy Wiley · Security Data Works

2026-05

ABSTRACT

Abstract

This report measures query-engine performance on a real Zeek network-telemetry workload across columnar OLAP engines (ClickHouse, Dremio, Trino, StarRocks) and the dominant schema-on-read SIEM. Headline result: ClickHouse Native completes a five-query analytical suite on 10 million Zeek conn.log events in 0.19 seconds; the schema-on-read SIEM completes the identical suite on the identical events in 27.52 seconds — a 145× ratio. Compression on the ClickHouse side reduces 3.27 GB of raw JSON to 399 MB on disk (8.2× factor). The methodology, hardware spec, query suite, scaling behavior, and caveats are documented here. The reference implementation (Docker Compose, data generators, runners) is shared under NDA with engagement prospects and qualifying reviewers.

1 1. Why this benchmark exists

Most published security-tool comparisons are vendor-authored, vendor-funded, or vendor-cooperated to the point of being marketing collateral. That includes mine, until recently — early write-ups of the result below leaned on cost framing that didn’t survive contact with the lakehouse-economics view (see Section 8, “What changed”).

The benchmark exists to characterize a single load-bearing question for security data architecture decisions: **on a real Zeek network-telemetry workload, how does query latency on columnar OLAP compare to schema-on-read indexing?** The answer informs the substrate decision underneath SIEM modernization, lakehouse migration, and federated detection pipelines.

What this report is not:

- Not a feature comparison. SIEMs ship correlation rules, alerting frameworks, RBAC, parsers, and field experience that columnar OLAP engines don’t. The comparison is on query-engine performance against the same workload, not on the broader product surface.
- Not a multi-node distributed benchmark. Single-node Docker Compose. Production environments run multi-node and the relative performance shifts on multi-node JOINs in ways this benchmark doesn’t capture.

- Not a recommendation that any organization replace its SIEM with ClickHouse. The matrix of when columnar OLAP wins, when SIEM-native wins, and when both belong in the architecture is the product of an engagement, not a benchmark.

2 2. Hardware and environment

Single-node Docker Compose on WSL2:

- **CPU**: 16 cores
- **RAM**: 32 GB
- **Storage**: NVMe SSD
- **OS**: Ubuntu 22.04 LTS under WSL2 on Windows 11
- **Container runtime**: Docker Engine 27.x
- **Orchestration**: Docker Compose v2.x

Engines tested under matched configurations:

Engine	Version	Storage	Memory limit
ClickHouse Native (MergeTree)	24.x	local SSD	16 GB
Dremio + Reflections	25.x	local SSD	16 GB
Trino	433	Iceberg/S3	16 GB
StarRocks	3.3	local SSD	16 GB
Schema-on-read SIEM	latest	local SSD	16 GB

No per-tool tuning was applied beyond documentation defaults. Vendor-recommended configurations are tested as additional rows where relevant; this report covers the documentation-default baseline.

3 3. Workload

Source: Zeek `conn.log` — per-connection network telemetry record produced by the Zeek (formerly Bro) network security monitor. Among the most common high-volume security data formats in production SOCs.

Volume: 10 million events. Generated synthetically with realistic field distributions matching a small-mid enterprise Zeek deployment shape (1.4 KB average raw event size; 3.27 GB total JSON).

Field schema (selected, full spec in reference implementation):

- `ts` — timestamp (epoch float)
- `id.orig_h`, `id.resp_h` — source/destination IPs
- `id.orig_p`, `id.resp_p` — source/destination ports
- `proto` — transport protocol (tcp/udp/icmp)
- `service` — application protocol (http, dns, ssl, etc.)
- `duration` — connection duration (seconds)
- `orig_bytes`, `resp_bytes` — payload bytes per direction
- `conn_state` — Zeek connection-state code (S0, SF, REJ, etc.)

Query suite — five queries, ten iterations each for statistical stability:

1. **Time-bucketed traffic aggregation** — events grouped into 5-minute buckets, summing bytes per protocol.
2. **Top talkers** — top-25 source IPs by total bytes over the full window.
3. **Protocol distribution** — count of events per `service` field value, ordered.
4. **Distinct-host count** — distinct internal IPs observed talking to external destinations.
5. **Cross-source JOIN with simulated SIEM alerts** — connection records joined against a 50K-row alert table, returning enriched alert context.

The fifth query is the most diagnostic. It exercises the join planner, which is where columnar engines and schema-on-read indexes diverge most sharply.

4 4. Headline result

Engine	P50 (s)	Δ vs. ClickHouse
ClickHouse Native	0.19	1.0×
Dremio + Reflections	1.0	5.3×
Schema-on-read SIEM	27.52	145×

ZSTD-22 compression on the ClickHouse side reduces 3.27 GB raw JSON to 399 MB on disk — an 8.2× factor. The schema-on-read SIEM's compressed footprint on the same data is roughly 2,385 MB, a 1.4× factor.

Scaling profile (full query suite at increasing event counts):

Events	ClickHouse Native	Schema-on-read SIEM
1 M	0.04 s	3.47 s
5 M	0.11 s	13.18 s
10 M	0.19 s	27.52 s

The schema-on-read SIEM scales worse than linear (8× latency for 10× data); ClickHouse stays sub-second across the full range.

5 5. Methodology — four principles

5.1 5.1 Reproducibility before performance

A benchmark result that can't be re-run isn't a benchmark; it's an opinion. The reference implementation contains the methodology document, container definitions, data generators, query suite, and analysis JSON. A practitioner with the same hardware can re-run the experiment and verify the number independently. The reference implementation is shared under NDA — see Section 9.

5.2 5.2 Identical workload across candidates

Workload and queries are pinned before any tool runs. The same query suite runs against the same data on every candidate engine. Vendor-recommended configurations are tested as additional rows in the result table — labeled — rather than folded silently into the headline.

5.3 5.3 Documented caveats

Every result ships with what was tested, what wasn’t, and where the result generalizes vs. doesn’t. See Section 7.

5.4 5.4 Vendor cooperation invited, not required

Every vendor whose product appears in a benchmark is invited to review the methodology and propose configuration changes before publication. Vendor-proposed configurations are tested and reported as additional rows. This benchmark does not accept funded engagements, pre-publication vetoes, or vendor workload selection.

6 6. Statistical confidence

Each query runs ten iterations. Reported numbers are P50 (median). P95 and standard deviation are recorded in the analysis JSON.

For the headline ClickHouse-vs-schema-on-read comparison, the standard deviation across iterations is under 5% of the median for both engines. The 145× ratio is robust to the variance — the smallest ClickHouse iteration (~0.16s) against the largest schema-on-read iteration (~30s) still produces a >180× ratio; the largest ClickHouse against the smallest schema-on-read produces ~110×. The headline is representative, not cherry-picked.

7 7. Caveats — when this benchmark applies and when it doesn’t

This result generalizes cleanly to:

- Security workloads dominated by a small number of recurring analytical queries (the SOC dashboard / saved-search pattern).
- Time-bucketed aggregation, top-talker analysis, protocol distribution, and joins against alert tables.
- Single-node deployments and the lower end of multi-node where partitioning is straightforward.

This result generalizes less cleanly to:

- Full-text search-dominated workloads (raw-event substring matching, free-form analyst hunting). The query suite doesn’t exercise full-text indexes; the schema-on-read SIEM’s relative position improves on those workloads.
- Multi-node JOIN-heavy workloads at petabyte scale. Single-node benchmarks don’t capture cross-node coordination overhead.
- Log shapes other than Zeek conn.log. Endpoint telemetry, cloud control-plane logs, and identity events have different cardinality and field-shape characteristics that affect both engines differently.

Known limits:

- The schema-on-read SIEM is tested on documentation-default configuration. Production deployments often have specialized data models, accelerated data models, summary indexes, and report acceleration configured. Each of those changes the comparison; none of them were applied here.
- The synthetic data generator approximates production Zeek shape but doesn’t exactly replicate any specific organization’s distribution.

8 8. What changed — cost framing in

2026

For most of 2025 my recommendation language was “ClickHouse is cost-effective for security workloads” — defensible at 30–90% cost reduction against per-GB-ingested licensing models. That framing broke in May 2026 in a conversation with Lipyeow Lim (Distinguished Engineer, Databricks). His framing: “ClickHouse is expensive” — when the comparison baseline isn’t legacy SIEM licensing but Iceberg-on-S3 with a separate query engine.

Both readings are correct. The recommendation now names the baseline:

- ClickHouse is **cheap** versus per-GB SIEM licensing.
- ClickHouse is **comparable** versus Snowflake or Databricks SQL at sustained TB/day.
- ClickHouse is **structurally more expensive** than Iceberg-on-S3 with a swappable engine — the MergeTree format duplicates data already storable in open formats; compute and storage scale together rather than independently; replication overhead multiplies storage cost.

Production deployments leaning ClickHouse-first on hot tier and Iceberg-on-S3 on cold tier are the emerging pattern that captures both the latency advantage and the open-format cost economics. The Q4 2026 lab work on streaming-write maturity into Iceberg is partly aimed at characterizing this hybrid shape.

9 9. Reference implementation access

The reference implementation — Docker Compose, data generators, query runner, analysis JSON — is shared under a one-page mutual NDA with:

- **Engagement prospects** during scoping. Run it on your own hardware before signing the SOW.
- **Qualifying reviewers** — security data engineers, OCSF contributors, analyst-firm researchers — for independent verification.

- **The annual external reviewer** named on the lab page each year.

The implementation is not published openly because the comparison set includes commercial software whose licensing terms restrict third-party publication of comparative test results. The methodology, the result, and the reasoning are public; the executable artifact is gated by a one-page NDA. Counsel-reviewed NDA template available at engagement scoping.

To request access: book a 30-minute discovery call via the website (`securitydataworks.com/engagements`) or email `jeremy@securitydataworks.com` with the subject line `Benchmark NDA request`.

Document version: v1.0 — 2026-05 **Next scheduled review**: Q3 2026 alongside the catalog comparison benchmark.