

FULL-STACK CLICKHOUSE INFRASTRUCTURE OPERATIONS · 24×7×365

ClickHouse engineering, consulting, support and remote DBA, built for real-time analytics at petabyte scale

More than 200 organisations entrust their ClickHouse estates to ChistaDATA. One senior engineering bench delivers architecture, performance engineering, 24×7 enterprise support, managed operations and remote DBA, so the engineer who designs the sort key is the engineer who answers the pager. Every recommendation is tied to a named metric in `system.*`, delivered in writing, and proven against production telemetry rather than benchmarks.

200+

organisations running ClickHouse with ChistaDATA

15 min

Severity 1 response, 24 hours a day, 365 days a year

11

global offices, follow-the-sun from the San Francisco Bay Area

100%

open-source ClickHouse (Apache 2.0), zero vendor lock-in

THE ESTATE WE ENGINEER AND OPERATE



Engineering

Performance, scalability, HA/DR, Data SRE, security and compliance

Consulting

Strategy, architecture, audits and migration from 12 source platforms

Support

24×7×365 enterprise support and fully managed ClickHouse operations

Remote DBA

Named senior ClickHouse DBAs from 8 hours a month to 24×7

SIX DISCIPLINES, ONE STANDARD OF EVIDENCE

Six engineering disciplines, each measured in production

Outcomes are stated in p95 and p99 latency, sustained ingestion throughput, storage efficiency and headroom under peak load; each discipline has a telemetry surface, a target and a written deliverable.

01 · PERFORMANCE

Query and storage performance

Sort-key and primary-index granule selection, skip-index design, projections, materialised-view denormalisation, `EXPLAIN indexes = 1` and `EXPLAIN PIPELINE` analysis, SIMD and vectorised-execution profiling from `system.trace_log`.

02 · SCALABILITY

Sharding and replica topology

Shard-key choice, Distributed table design, replica routing through `cluster()` and `remote()`, parallel replicas, Keeper sizing, object-storage tiering sized to the workload the business actually runs.

03 · HIGH AVAILABILITY

HA, DR and continuity

ReplicatedMergeTree configuration, 3- or 5-node Keeper quorum design on isolated disks, multi-region replication, cross-DC failover, and DR topologies tested against defined RPO and RTO in timed drills.

04 · DATA SRE

Reliability engineering

SLOs and error budgets, alerting on replication lag, parts count, merge backlog, Keeper latency and disk growth, runbook automation, chaos tests and quarterly restore drills from object storage.

05 · SECURITY

Security and compliance

Encryption at rest and in transit, RBAC with row and column policies, per-role quotas and settings profiles, GRANT-level audit logging, change control aligned to GDPR, HIPAA, SOX, PCI DSS and SOC 2.

06 · DATA STRATEGY

Data strategy and operations

Pipeline design with Kafka, Debezium, Kinesis, RabbitMQ and Flink, federation with the wider lakehouse, governance through Unity Catalog and Lake Formation, and the operating model around the cluster.

ENGINEERING SCOPE · EVERY SUBSYSTEM WE DESIGN, TUNE AND OPERATE

Storage and table engines

MergeTree ReplicatedMergeTree
ReplacingMergeTree SummingMergeTree
AggregatingMergeTree
CollapsingMergeTree
VersionedCollapsingMergeTree
GraphiteMergeTree

Distributed and replication

Distributed tables ClickHouse Keeper
ZooKeeper multi-region replication
parallel replicas `cluster()` / `remote()`
cross-shard JOIN

Ingestion and pipelines

Kafka engine Debezium CDC RabbitMQ
Kinesis Flink ClickHouse Pipes
Vector Fluent Bit HTTP / native

Storage tiering and cloud

S3 / object storage S3 disk
tiered storage TTL move policies
ClickHouse Cloud Aiven Kubernetes

Query performance and caches

vectorised execution SIMD
skip indexes Bloom filters
materialised views projections
query cache mark cache
uncompressed cache

Observability and operations

`system.query_log` `system.trace_log`
`system.parts` `system.merges`
`system.mutations` Prometheus exporter
OpenTelemetry Grafana
chistadata-agent

MEASURED IN THE CHISTADATA LAB

2,875 → 28 ms

Same 106 rows, with and without the leading sort-key column in the predicate

10,155 → 319 ms

Daily rollup served from a materialised view instead of the raw table

8.6 → 0.64 GB

Peak memory for a wide ORDER BY ... LIMIT with lazy materialisation enabled

1.34 B rows/s

Single-column aggregate scan rate, server-side, on a 2 vCPU node


THE FIRST QUERY ON EVERY ENGAGEMENT

```
SELECT -- cost per query fingerprint, last 24 h
      normalizedQueryHash(query) AS fingerprint,
      count() AS runs,
      quantile(0.99)(query_duration_ms) AS p99_ms,
      formatReadableSize(max(memory_usage)) AS peak_mem
FROM system.query_log
WHERE type = 'QueryFinish'
      AND event_time >= now() - INTERVAL 1 DAY
GROUP BY fingerprint
ORDER BY sum(query_duration_ms) DESC
LIMIT 20;
```

ChistaDATA lab, 1-billion-row table, ClickHouse 26.7 and 26.8 LTS, 2 vCPU / 7 GB, server-side timings from `system.query_log`, as published on chistadata.com. Test on your own workload.

FROM ASSESSMENT TO STEADY-STATE OPERATIONS

Four steps, each ending in a written, decision-grade deliverable



CONSULTING PORTFOLIO

Service	What we deliver
ClickHouse strategy	Architecture, schema and sort-key design, capacity models, SLO definition, self-managed vs vendor-managed decision
Performance audit	Top-N fingerprint review, EXPLAIN analysis, MergeTree settings, OS, filesystem and storage tuning, ranked fix list with expected gain
Real-time analytics	Materialised-view topology and projections that hold p99 below 100 ms at programmatic-auction scale
Data foundation and fabric	Kafka, CDC, lakehouse and federation around the ClickHouse core; cloud-native ClickHouse on Kubernetes
Data archiving	RDBMS archival to ClickHouse with TTL tiering to object storage and retention policies
DBaaS optimisation	Cost and performance review of managed ClickHouse and cloud data services, FinOps right-sizing
Generative AI	Vector search, feature tables and private, in-VPC RAG on ClickHouse with governance controls
Break-fix engineering	Emergency diagnosis and recovery for replication, Keeper, merge, mutation and disk-full incidents

MIGRATION AND MODERNISATION

End-to-end migration programmes into ClickHouse

Schema translation, sort-key and partition design for the target, dual-write or CDC cut-over, row-count and aggregate reconciliation, and a rollback path at every phase.

Amazon Redshift

Snowflake

Google BigQuery

Apache Druid

Apache Pinot

Vertica

Teradata

Hadoop / Hive

Elasticsearch

MySQL

MariaDB

PostgreSQL

Validation built into every cut-over. Source and target are compared per partition on row counts and checksummed aggregates before traffic moves.

```
-- Cut-over gate: rows and checksum per partition
SELECT toYYYYMM(event_time) AS month,
       count()               AS rows,
       sum(cityHash64(order_id)) AS checksum
FROM analytics.orders
GROUP BY month ORDER BY month;
```

STARTUP AND SCALE-UP CONSULTING PLANS

Avg. hours / month	Quarterly	Six-monthly	Annually
4	US\$7,500	US\$10,500	US\$25,500
8	US\$10,800	US\$15,500	US\$30,500
16	US\$15,500	US\$22,500	US\$40,000
40	US\$44,500	US\$58,500	US\$85,000

Selected tiers; plans are available in 4-hour steps from 4 to 40 hours a month.

THREE ENGAGEMENT MODELS, SIZED TO THE BUSINESS OUTCOME

Project engineering

Fixed-scope audits, migrations, HA/DR builds and performance remediation with acceptance criteria.

Consulting retainer

Principal-engineer time for architecture, schema review and upgrades: remote US\$450/h, on-site US\$600/h.

Managed services

Fully managed operations under contract: 24×7×365 monitoring, 15-minute Severity 1. See page 4.

24×7×365 CLICKHOUSE ENTERPRISE SUPPORT

A named principal engineer, a Technical Account Manager and a follow-the-sun incident bench

Support covers self-managed ClickHouse on bare metal, virtual machines and Kubernetes, and vendor-managed services including ClickHouse Cloud and Aiven. Every customer gets an incident bridge, a written root-cause analysis, and an engineer who already knows the cluster.

Sev	Definition	Response
S1	Production down or data at risk: inserts rejected, replicas read-only, Keeper quorum lost	15 minutes
S2	Severe degradation: p99 breach, replication lag growing, merges stalled	12 hours
S3	Partial impact: a query class, a shard or a pipeline affected, workaround available	24 hours
S4	Questions, configuration and upgrade guidance, feature requests	48 hours

Response clocks run 24 hours a day, 7 days a week, 365 days a year.

ENTERPRISE SUPPORT PLAN


US\$75,000

 per year

Unlimited ClickHouse instances, 24×7×365 reactive coverage with the full Severity 1-4 SLA.

- Named principal ClickHouse engineer and Technical Account Manager
- Incident bridges, written RCA with the evidence queries attached
- Upgrade and patch guidance across LTS lines with rollback plans
- Configuration, schema and query reviews on request
- Pairs with a Quarter or Half DBA plan for proactive engineering

HOW A SEVERITY 1 INCIDENT RUNS



CLICKHOUSE MANAGED SERVICES

Operations as a service, under contract

Monitoring and on-call

24×7×365, paged-engineer response, 15-minute Severity 1

Backup and DR

Backup validation and timed restore drills to RPO and RTO

Capacity reviews

Ingestion velocity, disk growth, Keeper and merge headroom

Versions and patching

Staged LTS upgrades with tested rollback

Schema-change review

Sort keys, partitions, TTL, codecs, views before deploy

Quarterly architecture audit

Security hardening and compliance reporting

HEALTH CHECKS THAT RUN ON EVERY CLUSTER

```
-- Replication: stuck or failing queue entries
SELECT database, table, type, num_tries,
       substring(last_exception, 1, 80) AS err
FROM system.replication_queue
WHERE num_tries > 10 ORDER BY num_tries DESC;

-- Parts pressure per partition, before inserts fail
SELECT database, table, partition, count() AS parts
FROM system.parts WHERE active
GROUP BY database, table, partition
HAVING parts > 300 ORDER BY parts DESC;
```

Order of magnitude

p99 latency reductions delivered through skip-index, projection and view engineering

Multi-million rows/s

sustained ingestion across Kafka, Debezium, Kinesis, RabbitMQ and HTTP paths

60–90% lower

operating cost reported by managed-services customers versus building an in-house ClickHouse DBA function

Audit-ready by default. Every cluster ships with documentation for GDPR, HIPAA, SOX, PCI DSS and SOC 2 controls: encryption posture, RBAC matrix, row and column policies, quotas, audit-log retention and change history.

SENIOR REMOTE DBAS AND DATA SRE FOR PRODUCTION CLUSTERS

A named ClickHouse DBA inside your workflow within days, not a six-to-twelve-month hire

Every DBA on the bench is a career ClickHouse specialist in MergeTree internals, replication and Keeper, backed by the 24×7×365 bench across eleven offices.

THE OPERATING MODEL: TELEMETRY FIRST, EVERY STEP WRITTEN DOWN

