

FULL-STACK DATABASE INFRASTRUCTURE · VENDOR-NEUTRAL · 24×7

Architecture, Engineering, Operations and Analytics for every database you run

MinervaDB is the vendor-neutral database infrastructure company that more than 900 enterprises rely on to architect, engineer, operate and optimize the data platforms their businesses run on: PostgreSQL, MySQL, MariaDB, Microsoft SQL Server, Oracle, IBM Db2, MongoDB, SAP HANA, ClickHouse, Trino, Apache Cassandra, Redis, Valkey, Milvus and every major cloud DBaaS. One accountable team, from the architecture decision to the 3 a.m. incident.

900+

enterprises served globally

46

cities with delivery presence

15+

years of database engineering depth

200+

years of combined leadership experience

15 min

Severity-1 response, 24×7×365

FOUR DISCIPLINES, ONE PARTNER

The four pillars of full-stack database infrastructure

01 / ARCHITECTURE

Correct at scale

Workload characterisation, engine selection, HA and multi-region topology, sharding and replication strategy, RPO/RTO engineering, capacity models built for 100× growth.

02 / ENGINEERING

Automated and hardened

Infrastructure as code, zero-downtime migrations and major upgrades, backup and PITR automation, observability and SLO instrumentation, security hardening.

03 / OPERATIONS

Reliability you can measure

24×7 monitoring and incident response, consultative support, Remote DBA, patch and version lifecycle, capacity management and executive SLO reporting.

04 / ANALYTICS

Data into decisions

Real-time and batch pipelines, columnar analytics on ClickHouse, federated query with Trino, lakehouse and warehouse engineering, vector search and RAG data layers.

SERVICE LINES

SERVICE	WHAT YOU GET
Engineering	Architecture, performance engineering, HA/DR, migrations, modernization off Oracle and Db2, IaC and automation
Consulting	Architecture reviews, performance audits, health checks, data strategy and governance, Fractional CDO
24×7 Support	Consultative support on every engine with S1 15 min · S2 12 h · S3 24 h · S4 48 h response targets
Remote DBA	Fractional or full-time senior DBA capacity: monitoring, backups, patching, tuning, capacity and security
Managed Services	End-to-end operation of the estate against explicit SLOs, with monthly reviews and quarterly drills

WHO WE SERVE

CIOs, CTOs, CDOs, GCC leaders, data architects and data scientists

Banking and fintech · insurance · healthcare and life sciences · telecommunications · retail, e-commerce and CPG · manufacturing · SaaS and ISVs · media · airlines and airport operations · supply chain and logistics · public sector and defence.

We sell no licences and earn no referral fees, so every recommendation follows the workload, including when the right answer is to keep what you have.

ENGINE-AGNOSTIC BY DESIGN

Production depth across every major database engine

Each engine is covered for consulting, engineering, 24×7 support and Remote DBA. We go deep on internals and say plainly when an engine is the wrong choice for a workload.

ENGINE	ROLE	ENGINEERING DEPTH
PostgreSQL	OLTP and HTAP	Planner and EXPLAIN analysis, autovacuum and bloat, XID wraparound prevention, logical and streaming replication, Patroni HA, PgBouncer, Citus, PostGIS, pgvector, TimescaleDB, FDWs; versions 13 to 18
MySQL	OLTP at scale	InnoDB internals, performance_schema, GTID replication, Group Replication and InnoDB Cluster, ProxySQL and MySQL Router, online schema change, 8.0 to 8.4 LTS upgrades, HeatWave
MariaDB	Open-source RDBMS	Galera clustering, MaxScale routing and failover, ColumnStore, system-versioned tables, MySQL-to-MariaDB migrations
SQL Server	Enterprise OLTP	Always On availability groups, Query Store, wait-stats analysis, columnstore, In-Memory OLTP, TDE, Azure SQL migrations; 2016 to 2025
Oracle / Exadata	Enterprise RDBMS	AWR/ASH (where licensed) and Statspack, SQL Plan Management, RAC, Data Guard, RMAN, Exadata tuning, licence-aware modernization to PostgreSQL
IBM Db2	LUW and z/OS	LUW: MON_GET monitoring, BLU, pureScale, HADR. z/OS: data sharing, utilities, plans and packages, CDC offload from the mainframe
MongoDB	Document store	Shard-key design, replica sets and write concern, WiredTiger cache, aggregation tuning, oplog sizing, Atlas operations
SAP HANA	In-memory ERP	Column-store memory, delta merge, NSE, system replication and takeover, backint backups, revision lifecycle
ClickHouse	Columnar OLAP	MergeTree sort keys and partitioning, projections, materialized views, Keeper, replication and sharding, S3 tiering, Kafka ingestion (with ChistaDATA)
Trino	Federated SQL	Connector strategy for Hive, Iceberg and Delta Lake, memory and resource groups, fault-tolerant execution
Cassandra	Wide-column	Query-first data modelling, consistency levels, multi-DC replication, compaction strategy, repair scheduling
Redis / Valkey	In-memory KV	Memory and eviction engineering, persistence trade-offs, Sentinel and Cluster, Redis-to-Valkey migration
Milvus	Vector search	HNSW, IVF and DiskANN index selection, partitioning, hybrid search and embedding pipelines for RAG

CLOUD DBAAS AND DATA PLATFORMS

Every major cloud, engineered rather than assumed

AWS
Amazon RDS, Aurora (PostgreSQL and MySQL), DynamoDB, ElastiCache, MemoryDB, Keyspaces, Redshift

Microsoft Azure
Azure SQL Database and Managed Instance, Azure Database for PostgreSQL and MySQL, Cosmos DB, Azure Cache for Redis

Google Cloud
Cloud SQL, AlloyDB, Spanner, Bigtable, Memorystore, BigQuery

Platforms and vendor clouds
Snowflake, Databricks, MongoDB Atlas, ClickHouse Cloud, Redis Cloud, Confluent, DataStax Astra

Kafka streaming

Debezium CDC

Kubernetes operators

Terraform IaC

Prometheus metrics

Grafana dashboards

pgBackRest backup

XtraBackup backup

Patroni HA

ProxySQL routing

dbt modelling

Iceberg / Delta table formats

DATABASE ENGINEERING

Engineering that is measured, rehearsed and reversible

Every recommendation cites the metric or system view behind it. Every change has a verification before, a confirmation gate, a validation after and a rollback written first.

Performance engineering

- Statements ranked by total database time
- Execution-plan and wait-event analysis
- Index strategy, statistics, memory and I/O tuning
- Connection pooling and concurrency control
- Measured before-and-after reports

Scalability architecture

- Read replicas and query routing
- Redis/Valkey cache tiers
- Analytical offload via CDC to ClickHouse
- Partitioning, then sharding (Citus, MongoDB, ClickHouse)
- Capacity forecasting from measured trends

High availability and DR

- Patroni, Group Replication, Galera, Always On, Data Guard, HADR
- RPO/RTO set in business terms
- Immutable, cross-region backups
- Quarterly restore and failover drills
- Fencing and split-brain prevention

Migration and modernization

- Oracle and Db2 to PostgreSQL or ClickHouse
- On-premises to DBaaS, and repatriation
- Bulk load plus CDC, row-level reconciliation
- Shadow period and gated cutover
- Reverse replication for rollback

Security and compliance

- CIS Benchmark and NIST CSF 2.0 hardening
- Least privilege, RLS, masking, TDE
- Audit logging outside the database host
- Evidence for SOC 2, ISO 27001, HIPAA, PCI DSS, GDPR, DPDP
- Just-in-time privileged access

Automation and observability

- Terraform and GitOps provisioning
- Reviewed, online schema migrations
- SLOs, error budgets, burn-rate alerts
- Engine-native telemetry dashboards
- Runbook automation and evidence capture

CONSULTING ENGAGEMENTS

ENGAGEMENT	TYPICAL SCOPE AND OUTPUT
Architecture review	Target architecture, capacity model, HA/DR design, security posture and a phased plan
Performance audit	Workload baseline, top statements by cost, remediation plan, tuned configuration, measured report
Health check	Configuration, replication, backup, security and capacity review, with risks ranked by impact and effort
Migration plan	Inventory, compatibility assessment, disposition per workload, cutover and rollback design
FinOps review	Billing joined to engine telemetry; right-sizing, storage class, reservations and query cost
Data strategy and governance	Metric tree, decision rights, classification and access enforced in the database, quality SLOs
Fractional CDO	Principal-level data leadership on a retained, part-time basis

HOW A PERFORMANCE AUDIT WORKS

1 · BASELINE

pg_stat_statements, Query Store, AWR, MON_GET



2 · DIAGNOSE

plans, waits, locks, I/O and memory



3 · REMEDIATE

index, query, config, pool, topology



4 · PROVE

p95/p99 and CPU per transaction, same window

Every finding names the metric or system view behind it.

LEGACY MODERNIZATION · ORACLE, EXADATA AND DB2

Modernize without a big-bang rewrite, and keep the legacy estate running under 24×7 support while you do

1 · Assess and dispose

Inventory schemas, PL/SQL, jobs and integrations; each workload is assigned migrate, refactor or retain.

2 · Convert

Schema and code conversion to PostgreSQL; analytical workloads to ClickHouse or the warehouse that fits.

3 · Synchronize

Log-based CDC keeps source and target in sync; z/OS workloads are offloaded incrementally.

4 · Cut over


Row-level reconciliation, rehearsed near-zero-downtime cutover and a tested rollback path.

24x7 CONSULTATIVE SUPPORT

Response targets that hold on every engine and every cloud




HOW EVERY INCIDENT IS HANDLED



REMOTE DBA SERVICES

Senior DBA capacity, fractional or full-time




ENGAGEMENT MODELS

MODEL	BEST FOR
24x7 Consultative Support	Teams that run their own estate and need expert escalation around the clock
Remote DBA	Lean teams that need senior DBA capacity without new hires
Managed Operations	End-to-end operation against SLOs with monthly reporting
Embedded Engineering	Sustained build-out, modernization and platform programmes
Fixed-scope projects	Audits, health checks, migrations and architecture reviews

ACCESS, ENGINEERED

Just-in-time, per-engagement credentials, hardware-backed MFA, sessions recorded through your bastion or PAM. Customer data never leaves your environment or enters a third-party AI tool.

WHAT EVERY MONTHLY SERVICE REPORT CONTAINS



DATA SCIENCE AND AI ON DATABASE-GRADE FOUNDATIONS

From trusted data to production AI

REAL-TIME ANALYTICS

ClickHouse and PostgreSQL analytics, streaming ingestion with Kafka and CDC, governed pipelines.

APPLIED ML AND MLOPS

Forecasting, pricing and fraud models in production; feature stores, registries, drift monitoring.

PRIVATE GENAI

In-VPC RAG on Milvus or pgvector, entitlement-aware retrieval, evaluation and guardrails.

GOVERNANCE

Metric trees, quality SLOs, lineage and access enforced in the database; GDPR, DPDP, HIPAA, EU AI Act.

INDUSTRIES

Banking and FinTech

Insurance

Healthcare and Life Sciences

Telecommunications

Retail and E-commerce

Consumer Packaged Goods

Manufacturing

SaaS and ISVs

Media and Gaming

Airlines and Airports

Supply Chain and Logistics

Public Sector and Defence

Global Capability Centres

WHY MINERVADB

Vendor neutrality

No licences sold, no referral fees; advice follows the workload.

Senior engineers only

Principal-level engineers design and operate your platform.

Full-stack accountability

One partner across engines and clouds; no finger-pointing at 3 a.m.

Knowledge transfer

Everything documented; your team gets stronger, not dependent.

HOW WE ENGAGE

01 · DISCOVER

workloads, telemetry, incidents, growth plans

**02 · DESIGN**

target architecture, risks, measurable outcomes

**03 · DELIVER**

build, migrate, tune, automate, tested change

**04 · OPERATE AND IMPROVE**

run to SLOs, report, optimize continuously

DATABASE SECURITY SERVICES
SUBSCRIPTION**Questionnaire-ready evidence across every engine**

- Hardening against CIS Benchmarks and NIST CSF 2.0, with a conformance score per instance
- Privileged-access reviews on a fixed cadence; just-in-time access, no standing credentials
- Evidence mapped to SOC 2, ISO 27001, GDPR, DPDP, HIPAA, PCI DSS and RBI programmes

TALK TO A PRINCIPAL ARCHITECT

The first conversation is always with an engineer, never a salesperson.**GLOBAL SALES, 24x7**

+1 (844) 588-7287 (toll free)
+1 (415) 212-6625 (USA)

EMAIL

contact@minervadb.com
support@minervadb.com
remotedba@minervadb.com

OFFICES**California**

MinervaDB Inc.
440 N Barranca Ave #9718
Covina, CA 91723

Delaware

MinervaDB Inc.
PO Box 2093
Philadelphia Pike #3339
Claymont, DE 19703

Book a consultation: minervadb.com/contact-minervadb-book-an-appointment

Engagement scope, response targets and deliverables are defined in each client agreement. Technical guidance should be validated in a non-production environment before it is applied to production, with verified backups and a tested disaster-recovery plan. Product names are trademarks of their respective owners.