



AI READINESS WITHOUT PLATFORM LOCK-IN

AI and real-time workloads stall when data is trapped in proprietary layers and governance cannot keep up.

Percona helps teams prepare their data layer for AI and real-time workloads without forcing a platform bet. We support open databases across MySQL, PostgreSQL, MongoDB, and Valkey/Redis, and help teams standardize governance, observability, and change execution across hybrid environments.

Percona helps teams prepare their data layer for AI and real-time workloads without forcing a platform bet. We support open databases across MySQL, PostgreSQL, MongoDB, and Valkey/Redis, and help teams standardize governance, observability, and change execution across hybrid environments.

Hybrid reality

75% of enterprise AI workloads will be deployed on hybrid infrastructure by 2027 (IDC).

- AI is not “one database” and “one cloud.” It is pipelines, staging, inference, and governance across environments.
- Teams need the ability to move workloads without rewriting operational controls or security posture.
- Portability preserves options when cost, regulation, or performance requirements change.

Governance friction

70% report data-governance difficulties in AI efforts (McKinsey).

- AI amplifies governance problems: access control, encryption, audit trails, retention, and lineage.
- Fragmented controls across engines and environments create audit gaps and operational drag.
- The goal is consistent, testable controls that travel with the workload.

Vector readiness

pgvector

Percona Distribution for PostgreSQL supports modern workloads through extensions like **pgvector** (for embeddings and vector search), where PostgreSQL is the right fit.

- “AI readiness” becomes expensive when it forces a new vendor platform or proprietary APIs.
- Start with practical steps: validate indexing, latency, concurrency, and observability for vector workloads.
- Keep the architecture open so you can evolve as ecosystems mature.



Learn more at percona.com



BUILD FOR AI, KEEP YOUR OPTIONS

AI pipelines demand mobility, security, and predictable operations. Some “AI-ready” platforms can trade those needs for proprietary APIs, fragmented governance, and unpredictable scaling cost.

Operate emerging workloads on your teams

► Software (Open source, upstream-compatible)

- Replace paid database subscriptions
 - **Percona Monitoring and Management** exposes I/O waste and oversizing
 - **Percona Operators** deliver features vendors usually charge for
-

► Expert Support

Expert Support is reactive, 24×7 help for incidents and performance issues.

- Fast resolution to minimize downtime
 - Deep performance tuning to prevent overscaling
 - Clear post-incident recommendations
-

► ExpertOps

ExpertOps Proactive Database Management is hands-on, ongoing ops support.

- Continuous tuning and rightsizing
 - Cleanup of replicas, storage, and configuration drift
 - Operational coverage for staffing gaps
-

► Expert Consulting & Services

Expert Consulting and Services offer project-based expertise for complex database work.

- Architecture reviews
- Migrations and EOL upgrades
- Targeted assessments for performance, sizing, and configuration



MySQL and MariaDB, PostgreSQL, MongoDB, Valkey, and Redis.
Operated with experts across on-prem, cloud, Kubernetes, and hybrid environments. No lock-in.