

emil.db

Emil Bialobrzewski

Senior Staff Database Reliability Engineer

DATABASE RELIABILITY · PLATFORM ARCHITECTURE · PRODUCTION ESCALATION

emildb.com · emil87b@gmail.com · linkedin.com/in/emilbialobrzewski · github.com/ebialobrzewski · United Kingdom · Remote

PROFILE

Senior Staff database engineer with 15+ years across relational and distributed data platforms, working at the intersection of reliability, architecture and operational automation. Acts as a senior escalation point for incidents that cross database, operating-system and infrastructure boundaries, and turns recurring operational problems into reusable platforms, standards and engineering patterns. Works across SQL Server, PostgreSQL, MariaDB and ScyllaDB, on-prem and in AWS.

SELECTED IMPACT

- **Replaced a 19-step daily DBA workflow with a shared multi-engine operations platform across 200+ on-prem hosts**, covering SQL Server, PostgreSQL, MariaDB and ScyllaDB and turning recurring health checks into a reusable read-only operating layer.
- **Reduced aggregate transaction-log footprint from ~5.2 TB to ~26 GB across ~1,970 databases** after tracing recurring growth to an estate-wide maintenance pattern; restored the intended RPO operating model.
- **Isolated ~20 GB of unexplained MySQL memory overhead** after engine-level diagnostics failed, tracing 53.7 GB resident against a ~33 GB buffer pool through Performance Schema, Linux process memory and allocator behaviour to glibc fragmentation amplified by Transparent Huge Pages.
- **Identified an AWS RDS point-in-time-recovery gap before an incident tested it** and designed the remediation model required to align platform capability with the documented RPO.

EXPERIENCE

Senior Staff Database Reliability Engineer — Diligent

Sep 2022 – Present

- Sets technical direction for database reliability across the estate: observability strategy, alerting model, operational standards and the automation roadmap that guides the DBA function.
- Owns reliability across a 400+ instance production estate spanning SQL Server, PostgreSQL, MariaDB, ScyllaDB, Cassandra and Couchbase, on-prem and in AWS, with the 200+ on-prem hosts under automated health checks.
- Senior escalation point for complex production incidents that cross database and operating-system layers — from wait statistics, execution plans and Extended Events to Linux process memory and allocator behaviour.
- Builds the operational leverage the team runs on: per-engine collectors, queue-based workers, Prometheus metrics, a read-only operations portal, self-service tooling and runbooks that reduce dependence on specialist intervention.
- Sets query-design and SQL standards with engineering teams, leads structured performance deep-dives, and mentors engineers on reliability and diagnostic practice.

Senior Database Engineer — Diligent

Feb 2018 — Sep 2022

- Broadened from SQL Server operations into ownership of a mixed-engine estate spanning PostgreSQL, MariaDB and AWS RDS / Aurora.
- Established SQL review and performance standards across development teams, reducing recurring production issues caused by poorly optimised queries.
- Resolved complex production escalations including Kerberos/SPN misconfiguration, execution plan regressions, replication failures and RDS / Aurora edge cases.

Earlier Database Engineering Roles

2011 — 2018

CashFlows · RateSetter · Cerulean Instrumentation · Regeneris

- Progressed from SQL development and data integration through development DBA work into production DBA and performance engineering.
- Delivered AlwaysOn Availability Groups as the HA/DR backbone, wait-statistics-driven tuning of high-traffic OLTP workloads, and database source control enabling versioned deployments.
- Authored the SQL coding standards and best-practice guides used by development teams — the start of the standards work that continues today.

CORE SCOPE

- PLATFORM DIRECTION** Database reliability patterns, shared operating models, observability strategy, operational standards.
- ARCHITECTURE & RISK** HA/DR, failover design, recoverability and RPO, capacity planning, cloud database architecture.
- TECHNICAL ESCALATION** Wait statistics, execution plans, Extended Events, Linux process memory, replication internals, engine diagnostics.
- ENGINEERING LEVERAGE** Python, PowerShell, T-SQL, Prometheus exporters, self-service tooling, runbooks, infrastructure as code, pragmatic LLM / agent workflows.
- DATA PLATFORMS** SQL Server, PostgreSQL, MariaDB / MySQL, ScyllaDB, Cassandra, Couchbase, AWS RDS / Aurora.

PERSONAL BUILD

GlidePlan · glideplan.org — flight-planning platform for glider pilots, built end to end with Flask, PostgreSQL, Docker, weather integrations, airspace validation and AI-assisted route planning.

CERTIFICATIONS & DEVELOPMENT

MariaDB Database Administration · PostgreSQL Administrator · AI_devs 4 Builders · production generative AI applications · 2026

EDUCATION

Engineer's Degree, Computer Science · Bialystok University of Technology · 2006–2011

Technical case studies and project details → emildb.com