CLOUDSWITCHED

FREE RESOURCE — GUIDE

Database Scaling & Optimisation Guide

Practical guide to scaling your database infrastructure for growing UK businesses — capacity planning, vertical and horizontal scaling, read replicas, and cost optimisation strategies.

5	30+	Frameworks	Free
SCALING STRATEGIES	EXPERT TIPS	INCLUDED THROUGHOUT	PRINT & USE NO STRINGS

CAPACITY PLANNING · VERTICAL SCALING · HORIZONTAL SCALING · READ REPLICAS · SHARDING · COST OPTIMISATION

PREPARED FOR	PREPARED BY	VERSION	FORMAT
Cloudswitched Knowledge Library	Cloudswitched Ltd.	2026 Edition	Scaling Guide

01

Capacity Planning & Assessment

Effective scaling starts with understanding where you are today and where you need to be. Data-driven capacity planning prevents both over-provisioning waste and under-provisioning outages.

Before investing in scaling infrastructure, you need a clear picture of your **current resource utilisation, growth trajectory, and performance bottlenecks**. Many organisations scale prematurely or in the wrong dimension — adding CPU when the bottleneck is I/O, or adding read replicas when the problem is poorly optimised queries.

Key Metrics to Baseline

1. **CPU utilisation:** Track average and peak CPU usage over a full business cycle (at least 4 weeks). Sustained utilisation above 70% during business hours indicates scaling pressure.
2. **Memory utilisation:** Monitor buffer pool hit ratio, page life expectancy, and memory grants. If your working set exceeds available memory, query performance degrades significantly.
3. **Storage I/O:** Measure IOPS, throughput (MB/s), and latency for both read and write operations. Storage is frequently the first bottleneck to hit in growing databases.
4. **Connection count:** Track concurrent connections over time. Approaching the maximum configured limit causes connection timeouts and application errors.
5. **Database size:** Monitor total database size and growth rate. Project when you will reach storage limits, licensing thresholds, or backup window constraints.

01

Growth Projection Framework

Track each headline metric against a 6- and 12-month projection, and define the threshold that should trigger action well before you hit a wall.

METRIC	CURRENT VALUE	6-MONTH PROJECTION	12-MONTH PROJECTION	ACTION THRESHOLD
Database size (GB)				80% of storage limit
Peak concurrent connections				70% of max configured
Peak CPU utilisation (%)				70% sustained
Buffer pool hit ratio (%)				Below 95%
Average query response time (ms)				Exceeds SLA target
Peak IOPS				80% of storage tier limit
Daily transaction volume				Capacity planning input

OPTIMISE BEFORE YOU SCALE

Scaling hardware to compensate for inefficient queries is expensive and unsustainable. Always exhaust query optimisation, indexing improvements, and caching strategies before investing in infrastructure scaling. A single poorly written query can consume more resources than an entire application.

Document your findings in a **capacity planning report** that is reviewed quarterly. This report should drive budget requests, procurement timelines, and architecture decisions well in advance of hitting capacity limits.

02

Vertical Scaling Strategies

Scaling up (adding more resources to your existing server) is the simplest approach and should be your first consideration for most workloads.

Vertical scaling — increasing CPU, memory, or storage on a single server — is the **lowest-risk, lowest-complexity scaling approach**. It requires no application changes, no data redistribution, and no architectural redesign. For the majority of UK SMEs, vertical scaling is sufficient to handle growth for years.

When Vertical Scaling Works Best

1. **Memory-bound workloads:** If your buffer pool hit ratio is below 95% or page life expectancy is dropping, adding RAM delivers immediate and significant performance improvement.
2. **CPU-bound analytics:** Reporting queries, aggregations, and complex joins benefit directly from additional CPU cores and faster processors.
3. **I/O-bound workloads:** Migrating from spinning disks to SSD, or from SSD to NVMe, can deliver 10–100x improvements in I/O latency and throughput.
4. **Cloud-hosted databases:** Cloud providers allow vertical scaling with minimal downtime — Azure SQL and AWS RDS can scale compute tiers in minutes.

02

Cost Considerations for UK Businesses

Weigh one-off on-premises spend against incremental cloud cost — and never forget that per-core licensing often dwarfs the hardware itself.

RESOURCE	ON-PREMISES COST	CLOUD COST (APPROX. MONTHLY)	IMPACT
RAM (32 GB □ 64 GB)	£200–£400 one-off	£50–£150 increment	Buffer pool, caching
CPU (8 cores □ 16 cores)	£1,000–£3,000 one-off	£100–£400 increment	Query processing, concurrency
Storage (SSD □ NVMe)	£500–£2,000 one-off	£30–£100 increment	I/O latency, throughput
SQL Server licensing (per core)	£3,000–£15,000 per core p.a.	included in managed tier	Significant cost factor

LICENSING COSTS CAN DWARF HARDWARE COSTS

For databases licensed per core (SQL Server Enterprise, Oracle), adding CPU cores doubles your licensing cost. Always calculate the total cost including licensing before scaling CPU vertically. This is often the trigger that drives migration to open-source alternatives or cloud-managed services.

Vertical scaling has a ceiling — eventually you reach the maximum specification available for a single server. For most UK SMEs running business applications, that ceiling is **far higher than you think**. A modern server with 128 GB RAM, 32 cores, and NVMe storage can handle millions of transactions per day.

03

Horizontal Scaling & Sharding

Scaling out by distributing data across multiple servers is powerful but complex. Understand the trade-offs before committing to this path.

Horizontal scaling — distributing your database across multiple servers — is necessary when vertical scaling reaches its limits or when **geographic distribution, fault tolerance, or extreme throughput** demands require it. However, it introduces significant architectural complexity that must be justified by genuine requirements.

Sharding Strategies

1. **Range-based sharding:** Data is distributed by value ranges (e.g., customers A–M on shard 1, N–Z on shard 2). Simple to implement but prone to hotspots if data distribution is uneven.

2. **Hash-based sharding:** A hash function distributes data evenly across shards. Provides better load distribution but makes range queries across shards more complex.

3. **Geographic sharding:** Data is partitioned by region (e.g., UK customers on UK servers, EU customers on EU servers). Supports data sovereignty requirements and reduces latency for local users.

4. **Tenant-based sharding:** Each customer or tenant gets their own database or schema. Common in multi-tenant SaaS applications and simplifies data isolation.

03

Complexity Trade-Offs

Sharding solves extreme-scale problems, but every benefit comes with a cost in operational and development complexity.

CONSIDERATION	SINGLE SERVER	HORIZONTALLY SCALED
Query complexity	Standard SQL	Cross-shard queries require coordination
Transaction support	Full ACID	Distributed transactions are slow and complex
Data consistency	Immediate	Eventually consistent (in many designs)
Operational overhead	Single server to manage	Multiple servers, rebalancing, monitoring
Application changes	None required	Significant refactoring for shard-aware queries
Backup and recovery	Standard procedures	Per-shard backup coordination
Cost	Single server cost	Multiple servers plus orchestration tooling

MOST UK SMES DO NOT NEED SHARDING

Sharding is an engineering solution for extreme scale. If your database is under 500 GB and serves fewer than 1,000 concurrent users, vertical scaling and read replicas will almost certainly meet your needs at a fraction of the cost and complexity. Do not over-engineer.

CONSIDER MANAGED DISTRIBUTED DATABASES

If you genuinely need horizontal scaling, consider managed distributed databases (Azure Cosmos DB, Google Cloud Spanner, CockroachDB) that handle sharding, replication, and rebalancing automatically. The operational overhead of self-managed sharding is substantial and rarely justified.

04

Read Replicas & Load Distribution

Offloading read traffic to replicas is often the most cost-effective scaling strategy for read-heavy business applications.

Most business applications are **heavily read-biased** — reporting, dashboards, searches, and data lookups far outnumber write operations. Read replicas allow you to scale read capacity independently of write capacity without the complexity of full horizontal scaling.

Read Replica Architecture

1. **Primary handles all writes:** INSERT, UPDATE, and DELETE operations are directed exclusively to the primary database server, which maintains the authoritative copy of all data.
2. **Replicas handle read traffic:** SELECT queries for reporting, dashboards, search, and analytics are routed to one or more read replicas, offloading the primary server.
3. **Replication lag awareness:** Data on replicas may be seconds behind the primary. Applications must be designed to tolerate this — reading from a replica immediately after writing to the primary may return stale data.
4. **Load balancing:** A connection proxy or application logic distributes read queries across multiple replicas for additional throughput and resilience.

04

Use Cases & Implementation

A single read replica unlocks reporting, analytics, backup offloading, and geographic reach — with minimal application change.

Use Cases for Read Replicas

1. **Business reporting:** Dashboards, scheduled reports, and ad-hoc queries run against a replica without impacting transactional performance on the primary.
2. **Search and analytics:** Full-text search, aggregations, and data analysis workloads are directed to dedicated read replicas optimised for analytical queries.
3. **Backup offloading:** Backups are taken from a read replica rather than the primary, eliminating backup I/O impact on production transactions.
4. **Geographic distribution:** Replicas in different UK or European regions reduce read latency for distributed teams and improve disaster recovery posture.

DATABASE ENGINE	NATIVE REPLICATION	MANAGED REPLICA OPTION	TYPICAL LAG
PostgreSQL	Streaming replication	AWS RDS, Azure Flexible	Sub-second
MySQL / MariaDB	Binary log replication	AWS RDS, Azure MySQL	Sub-second to seconds
SQL Server	Always On Availability Groups	Azure SQL, AWS RDS	Sub-second (sync mode)
MongoDB	Replica sets (built-in)	Atlas (managed)	Sub-second

START WITH ONE REPLICA

A single read replica is sufficient for most UK SMEs. It offloads reporting and analytics from the primary, provides a warm standby for disaster recovery, and serves as a backup source. Add additional replicas only when monitoring shows the first replica is saturated.

05

Cost Optimisation & Right-Sizing

Database infrastructure is often over-provisioned. Right-sizing ensures you pay only for what you need without sacrificing performance.

UK businesses frequently over-provision database infrastructure "just in case", or because specifications were set during initial deployment and never revisited. A disciplined approach to **right-sizing, reserved capacity, and architecture optimisation** can reduce database costs by 30–50% without impacting performance.

Cloud Cost Optimisation

1. **Right-size compute tiers:** Review actual CPU and memory utilisation over 30 days. If average utilisation is below 40%, you are likely paying for capacity you do not use. Downsize to the next tier and monitor.

2. **Reserved instances:** For stable, predictable workloads, purchasing 1-year or 3-year reserved capacity on AWS or Azure saves 30–60% compared to on-demand pricing. Commit only for base-level requirements.

3. **Storage tier optimisation:** Not all data needs premium storage. Move cold and archival data to standard SSD or even HDD tiers. Use automated tiering policies where available.

4. **Dev/test environments:** Non-production databases should use the smallest viable tier and be shut down outside business hours. Automate start/stop schedules to eliminate overnight and weekend costs.

5. **Serverless options:** For variable or intermittent workloads, serverless database tiers (Azure SQL Serverless, Aurora Serverless) scale to zero when idle and charge only for actual usage.

05

On-Premises Cost & Tracking

On-premises estates have their own levers — licensing, consolidation, and refresh cycles — and a simple tracking framework keeps the savings honest.

On-Premises Cost Optimisation

1. **Licensing review:** Audit your database licensing to ensure you are not paying for Enterprise features when Standard edition meets your requirements. The cost difference is substantial — often £10,000+ per server.
2. **Consolidation:** Multiple under-utilised database servers can often be consolidated onto fewer, more powerful hosts. This reduces licensing, hardware, and operational costs simultaneously.
3. **Open-source migration:** Evaluate whether PostgreSQL or MariaDB can replace commercial databases for non-critical workloads. Licensing savings fund the migration effort many times over.
4. **Hardware refresh cycles:** Modern hardware delivers significantly more performance per pound than equipment from 3–5 years ago. A strategic refresh can improve performance while reducing power and cooling costs.

NEVER SACRIFICE RESILIENCE FOR COST

Cost optimisation must never compromise backup integrity, disaster recovery capability, or security controls. Saving £200 per month on a backup solution is meaningless if a data loss event costs your organisation £50,000 in recovery, downtime, and regulatory fines. Optimise intelligently, and schedule a quarterly cost review that compares actual spend against budget.

Your Database Scaling Action Plan

Turn this guide into a plan. Tick off each milestone as you complete it, then capture owners, dates and next steps below.

Scaling milestones

- ☐ **Capacity baselined** — CPU, memory, storage I/O, connections and database size measured over a full business cycle, with a growth projection and action thresholds recorded.
- ☐ **Optimise-before-scale exhausted** — query optimisation, indexing, caching and connection pooling reviewed before any infrastructure spend is committed.
- ☐ **Vertical scaling assessed** — right-size of CPU, RAM and storage costed (including per-core licensing) against the single-server ceiling.
- ☐ **Read replicas evaluated** — a single replica considered for reporting, analytics, backup offloading and DR before any sharding is contemplated.
- ☐ **Cost review scheduled** — right-sizing, reserved capacity and tiering opportunities identified, with a recurring quarterly cost review in place.

Top priority next steps

01 _____

02 _____

03 _____

04 _____

SCALING OWNER _____

TARGET REVIEW DATE _____

PRIMARY BOTTLENECK _____

CLOUDSWITCHED

IT SUPPORT & PROJECT SERVICES

info@cloudswitched.com

New London House, 6 London St
London EC3R 7LP · United Kingdom