We keep the data **your business runs on** correct, quick to access and recoverable.

Database design, administration, tuning and migration, pipelines, warehouses and dashboards — from Copenhagen.

Engineer company
Data

Databases and data pipelines that hold up

+45 50 10 61 84 · Copenhagen, Denmark

Carefully Crafted With Love by D.

We design, tune and migrate the databases a business runs on, and we build the pipelines that move data between them. The work records describe the starting situation, what was done and the result.

Bring one database or data flow and the questions it must answer. We can scope a review that leaves you with a written diagnosis and a prioritised implementation plan.

Database and pipeline work · Pricing and engagement options · Discuss your project

Start with the result you need

Tell us what your business needs to achieve. We connect software, data and infrastructure around the people who use them.

Explore our services Get in touch →

See how we approach the work

Find a service that matches your goal and see examples of what was delivered and how.

Cases Get in touch →

Read the work **in its context** — the decisions, delivery and reported outcomes.

See the whole catalogue

One partner for **9 services** from the schema to the dashboard.

Browse the catalogue

{ Data & Databases }

Somewhere sensible to keep your data

Database Design & Modeling

A structure that matches how the business actually works, so the questions you need answered later are still possible to ask.

Built a database-driven CV generator

Built a database-driven CV, references, portfolio and cover-letter generator in Python — 41 modules, 10,580 lines — rendering six output formats from one 23-table SQLite source assembled by a 19-step idempotent pipeline.

Built the hiring and seafarer career workspaces

Built hiring and seafarer career workspaces; the recorded implementation snapshot contains 151 stored functions across 48 tables for vacancies, applications, certificates, rank progression and sea time.

Built a loyalty wallet, leagues and badges

Built the loyalty and reputation system — 67 functions over a 31-table ledger, with leagues, badges and a redemption shop — taking a row lock on the balance to close the double-spend window.

See all 14 achievements

How we work

Understand first

We learn the problem before we propose a solution. The brief you arrive with is the start of that conversation, not the end of it.

Build it to last

Durable over quick, clarity over cleverness. The system still has to be workable by the people who inherit it.

Stand behind it

We agree any ongoing operation or support as part of the engagement. Handover makes responsibilities and the next contact route clear.

{ Data & Databases }

Looking after the database

Database Administration (DBA)

Somebody responsible for the thing the whole company rests on: it is up, it is backed up, and it is not quietly rotting.

1,022 migrations behind a schema drift gate

Kept the schema honest across 1,022 migrations with a CI gate that builds the database both ways — a fresh install, and an install plus every migration — and fails when the two disagree.

Implemented backups and disaster recovery as code

Configured database backup retention as infrastructure-as-code, then audited the recovery position and documented the restore procedure — naming the remaining gaps rather than leaving them to be found during an incident.

Engineered 600 PL/pgSQL ETL/ELT pipelines

Engineered 600 PL/pgSQL-based ETL/ELT pipelines to streamline complex data processing workflows across multiple PostgreSQL development and production environments.

See all 9 achievements

{ Data & Databases }

Making the slow parts fast

Database Performance Tuning

Finding the queries and the missing indexes behind a system that has got heavier every year, and fixing the ones that matter.

Time-ordered keys for a growing data model

Adopted UUID v7 time-ordered identifiers (PostgreSQL 18) as entity keys to reduce B-tree index fragmentation and speed up queries.

Accelerated PostgreSQL performance by 10x

Accelerated PostgreSQL performance by 10x via strategic indexing, partitioning, and query optimization, enhancing database efficiency for user, tenant, geospatial, and time-series electrical data.

Ran 100 PostgreSQL, MS SQL, BigQuery warehouses

Architected, created, and managed 100 PostgreSQL, MS SQL, and Google BigQuery data warehouse databases with primarily GIS and time-series data, optimizing performance and scalability.

See all 6 achievements

{ Data & Databases }

Getting off the old system

Database Migration & Modernization

Data taken out of something ageing and into something supported, with the business still running while it happens.

Moved every string into a content tree

Moved every user-facing string out of Python into a content tree of 874 files across three languages, after finding dead translations nobody could see were dead and a check silently grading a third of the achievements.

1,022 migrations behind a schema drift gate

Kept the schema honest across 1,022 migrations with a CI gate that builds the database both ways — a fresh install, and an install plus every migration — and fails when the two disagree.

Improved GIS map performance 10x on PostgreSQL

Improved geographical map application performance by 10x through strategic database transition from MSSQL to PostgreSQL, including the data cutover, optimizing processing and data security.

{ Data & Databases }

Getting your data from A to B on its own

Data Pipeline Development (ETL/ELT)

Numbers arriving where they are needed, on a schedule,
without anyone copying a spreadsheet at the end of the month.

Generated 495 pages from a SQLite export

Generated 495 achievement pages across three languages from a read-only SQLite export, with the page address authored as data so that correcting a sentence no longer moved the page and broke the link.

Built first-party analytics inside PostgreSQL

Built first-party product analytics in PostgreSQL — 47 functions over partitioned event tables that prune themselves — pseudonymised behind a rotating salt and gated on the visitor's consent.

A mapping product, carried end to end

Architected, developed and supported infrastructure, data processing and a map application over 2 years, covering both software delivery and operational maintenance.

See all 14 achievements

{ Data & Databases }

One place where all the numbers agree

Data Warehouse Design

Sales, operations and finance figures brought together, so two departments stop bringing two answers to the same meeting.

Ran 100 PostgreSQL, MS SQL, BigQuery warehouses

Architected, created, and managed 100 PostgreSQL, MS SQL, and Google BigQuery data warehouse databases with primarily GIS and time-series data, optimizing performance and scalability.

Designed and ran 6 ETL/ELT pipelines

Designed, implemented, and administered 6 ETL/ELT pipelines, utilizing Google BigQuery, MSSQL, PostgreSQL, Shell scripting, PL/pgSQL, and Transact-SQL, integrating data for efficient Python API processing.

Made a geographical data pipeline 50x faster

Recorded a 50x geographical data-pipeline improvement in the founder's Utiligize work through SQL and data-model changes across PostgreSQL, MS SQL and Google Cloud BigQuery.

{ Data & Databases }

Seeing what is going on

Data Analytics & BI Dashboards

Dashboards and reports showing the handful of figures a decision actually turns on, refreshed without being asked.

Built first-party analytics inside PostgreSQL

Built first-party product analytics in PostgreSQL — 47 functions over partitioned event tables that prune themselves — pseudonymised behind a rotating salt and gated on the visitor's consent.

Designed a time-management and reporting system

Designed a time-management and reporting system that remained in production use for years without significant modification.

Delivered analysis and regular progress reports

Delivered analysis and regular progress reports to the CEO, translating engineering KPIs and delivery status into decisions.

See all 11 achievements

{ Data & Databases }

Trusting the numbers

Data Governance & Quality

Rules and checks that catch bad data on the way in, and a clear answer to who owns which figure and where it came from.

Measured 13 addresses with no tracking

Measured thirteen addresses from the web server's own access log alone — no cookie, no tracker and no third-party service — masking the client network as each line is written, and folding the result into a permanent archive of 21 daily counts and 20 monthly facets.

Know the licence of everything you ship

Audited 644 Rust crates for licence compatibility on every build, and proved the check fires by rewriting one crate's licence and by moving the pinned core revision without regenerating.

Built a cross-language figure check

Built a cross-language content check that fails when a translation drops a figure the English states, and when a language uses notation it does not use — finding two Danish descriptions missing a metric and sixteen Ukrainian spans quoting in the English style.

See all 16 achievements

{ Data & Databases }

Data an AI can actually use

AI Data Readiness

Your records where AI agents can reach them, one version of each customer, nobody shown more than they should, and every answer traced to its source.

Replaced 963 metadata blocks with one graph

Replaced 963 anonymous structured-data blocks that restated the company 1,671 times with one linked graph of 16 types minted from stable origin identifiers.

Generated 495 pages from a SQLite export

Generated 495 achievement pages across three languages from a read-only SQLite export, with the page address authored as data so that correcting a sentence no longer moved the page and broke the link.

1,022 migrations behind a schema drift gate

Kept the schema honest across 1,022 migrations with a CI gate that builds the database both ways — a fresh install, and an install plus every migration — and fails when the two disagree.

See all 6 achievements

We work with **the platforms** your business already runs on.

Apple Microsoft DigitalOcean AWS Google

See the whole stack

Company data your AI agents can answer from

Gathered

Your CRM, ERP, shared drives and mail archives brought into one database an agent reaches through one described door.

Cleaned

Agreed identifiers and quality checks connect the selected data sources, with duplicate handling and validation rules documented.

Permissioned

Access checks and audit logging follow the permissions and data boundaries agreed for the integration.

See the data work Ask about your data →

Tell us what you are building.

No form, no funnel, no sales sequence — just one e-mail to a person who will read it.

Start the conversation

This address shows the data work only. The company also builds software, runs infrastructure and advises teams — that side of it is at engineer.company.
See everything Engineer ApS does →

Begin with a focused improvement

A workflow, a data flow or a system that needs improvement can be a useful first task. We can agree its scope together.

See what it costs Get in touch →

Plan the first conversation

Tell us the result you need, what is already in place and when you need it. We can use that to discuss the next step.

Explore our services

welcome@engineer.company +45 50 10 61 84 Copenhagen, Denmark

On the server, we work with **open-source Linux** — no licences to pay for, and nothing hidden.

Ubuntu Debian

How we run servers

Take this page with you

Contact card

Carefully Crafted With Love by D.

CVR 45969703 · Tagensvej 105C, 3. th, 2200 København N

Copyright © 2025 – present · Engineer ApS · All rights reserved.

Printed from https://data.engineer.company/