



MIGRATING FROM MS SQL TO POSTGRES ENTERPRISE EDITION (PGEE) AT SCALE

CYBERTEC DATABASE CONCEPTS

Technical paper

Hans-Jürgen Schöning
2025-01-01

TABLE OF CONTENTS

- TABLE OF CONTENTS..... 2**
- CURRENT STATUS AND GOAL..... 3**
- REDUCING THE MS SQL BY THOUSANDS
OF SERVERS..... 3**
 - HOW TDS EMULATION WORKS..... 4**
- A ROADMAP TO SUCCESS..... 5**
 - ADDRESSING COMPLIANCE
AND SECURITY ISSUES..... 5**
 - ADDRESSING AUTOMATION AND HARDWARE COSTS..... 6**
 - TRANSITIONING MADE EASY..... 7**
 - BABELFISH IN ACTION..... 8**
 - ONGOING SUPPORT AND MAINTENANCE..... 9**
 - SUPPORT SERVICES..... 9**
 - TRAINING AND KNOWLEDGE TRANSFER..... 10**
 - ONGOING DEVELOPMENT..... 10**
- CONTACT INFORMATION..... 11**
- VERSION HISTORY..... 12**

CURRENT STATUS AND GOAL

Currently, many institutions operate thousands of MS SQL deployments that should be migrated as soon as possible to **reduce long-term** support and licensing **costs**.

To facilitate the transition, CYBERTEC proposes a **comprehensive solution** based on a multi-stage approach yielding almost instant success. In the current installation, various types of systems are to be taken into consideration:

- Development environments
- Test and stage systems
- Production deployments

Overall, this involves thousands of systems, potentially requiring hundreds or even thousands of systems to be addressed in order to **remove MS SQL** from the infrastructure. Managing this volume of systems introduces several implications:

- Accelerated migration timelines
- Increased pressure
- Multiple concurrent projects
- Strain on human resources

The goal of this proposal is to **mitigate these concerns by introducing a new concept** aiming to reduce the exposure to MS SQL, while reducing the cost of the deployment considerably.

REDUCING THE MS SQL BY THOUSANDS OF SERVERS

What if we could **remove MS SQL without modifying the application side**?

This approach would offer several advantages:

- Nearly immediate cost reduction for MS SQL
- Additional time for a real migration
- Reduced pressure on manpower
- A clear path to PostgreSQL

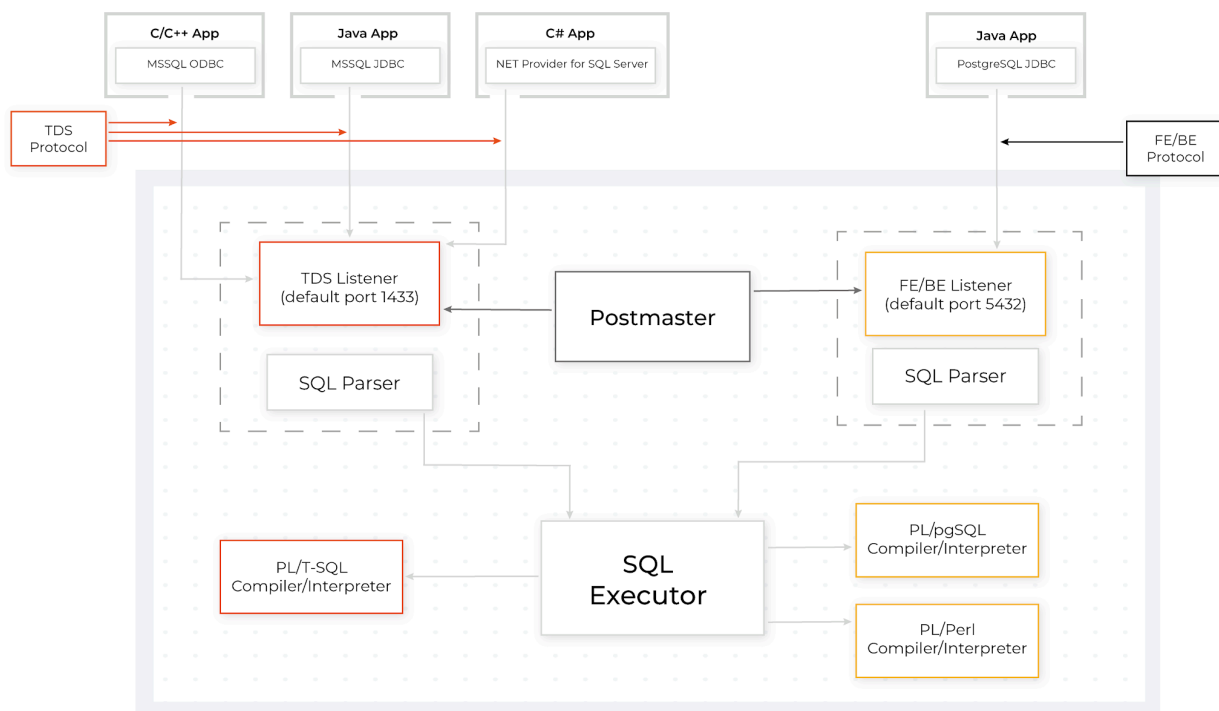
The solution to the problem is a concept in PostgreSQL called “protocol hooks”. Using this facility allows to **expose PostgreSQL as a different database**, in case of MS SQL, PostgreSQL will expose itself using the TDS protocol spoken by the closed source counterpart.

How does this work? In the Open Source world, the tool Babelfish offers an approach **adjusted** to fit customers needs on the compliance side.

HOW TDS EMULATION WORKS

From a technical perspective, Babelfish works by parsing incoming MS SQL requests which is then translated into commands that can be handled by the PostgreSQL optimizer and executor.

The diagram below shows how the process works:



This allows us to communicate with this specialized database as if it were MS SQL. For a significant number of applications this is entirely sufficient and it offers a cost effective solution, enabling quick savings.

The following types of databases can be moved to Babelfish quickly:

- Test databases and development
- Small databases with little T-SQL code
- Databases with not too many MS SQL specific features

These systems can be dumped and reloaded into Babelfish. After a quality inspection by the app team, the applications are moved off MS SQL without requiring a full migration.

Once the Babelfish dump is complete, the MS SQL estate will have been reduced dramatically (by an expected 40-70%). Ultimately, these applications will be fully migrated to PostgreSQL. However, the Babelfish solution provides a significant reduction in MS SQL license requirements while enabling full migrations to PostgreSQL to be completed without the commercial pressures of license reduction.

A ROADMAP TO SUCCESS

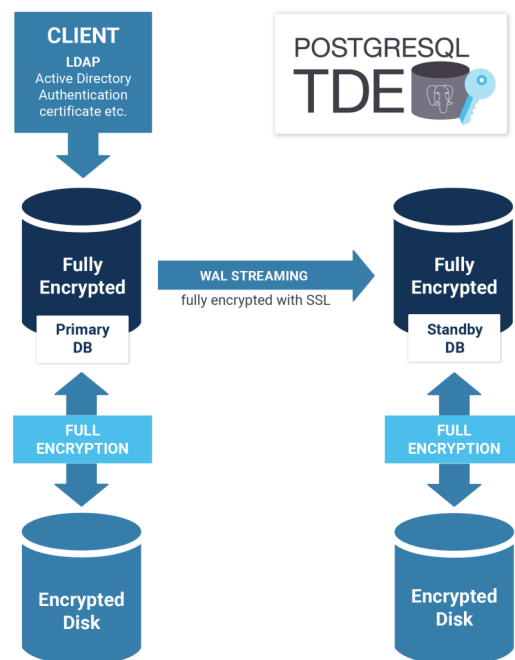
To achieve the goal of reducing the MS SQL by 40-70% we suggest the following roadmap and process, implementing it step by step to ensure a quick reduction of the Microsoft solution:

ADDRESSING COMPLIANCE AND SECURITY ISSUES

The first and most crucial issues that have to be addressed are related to:

- **Security and compliance**
- **Encryption and key management**
- **Security monitoring**

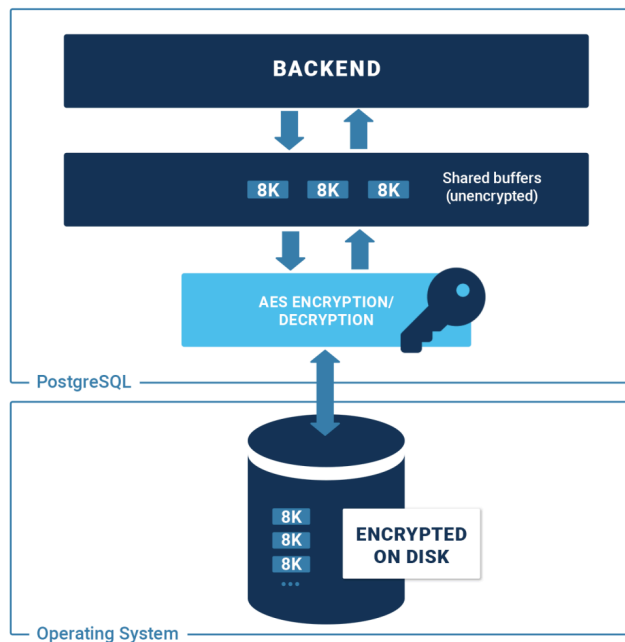
The solution involves porting the encryption code (TDE = **Transparent Data Encryption**) from PGEE to Babelfish. This approach ensures we meet the **compliance** requirements satisfied by PGEE and other CYBERTEC tools.



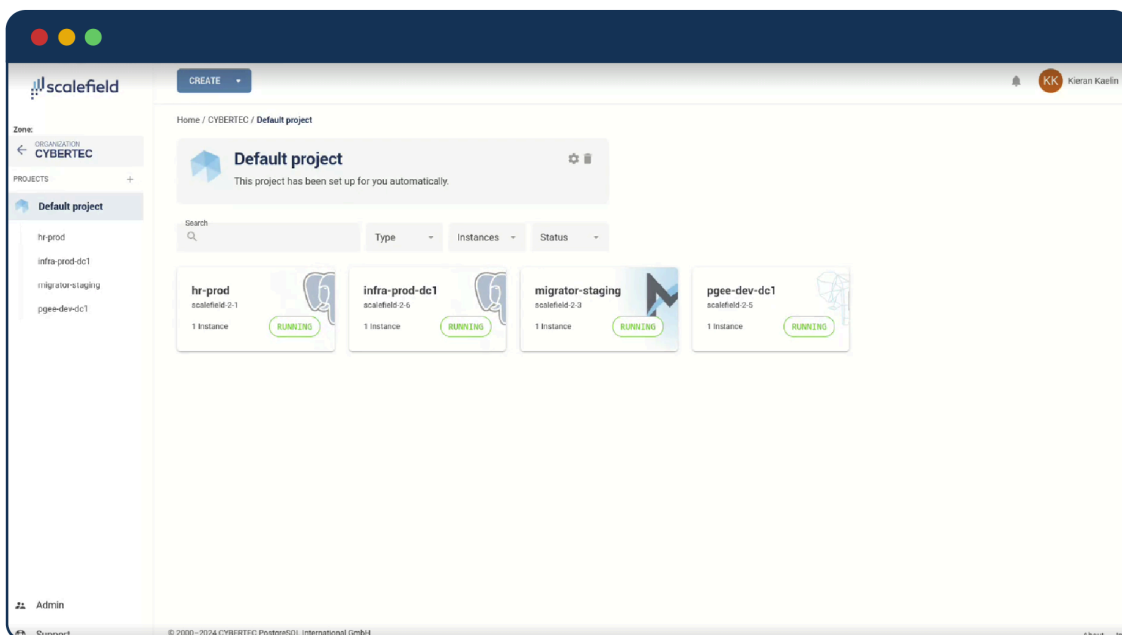
After the integration of the security code, we will be able to deploy a secure version of Babelfish ready for field testing and auditing.

ADDRESSING AUTOMATION AND HARDWARE COSTS

Managing thousands of MS SQL deployments can be cumbersome. Therefore; it is important to use the transition to the new platform for automation and streamline deployment processes.



The solution we propose is to integrate Babelfish into the CYBERTEC Scalefield:



New instances including **High-Availability** can then be created visually. **Scalefield** is fully Kubernetes / OpenShift based and allows for **rapid deployment** of single instances or entire clusters. This is **ideal** for development, testing, migration, and agile development following the dev-ops approach.

The following features are part of the platform:

- **Integrated backup** and recovery
- Support for **full encryption**
- Integrated support for “data sets”
- Out-of-the-box **clustering** and **HA**
- Support for SSO (**Single-Sign-On**)
- Visual support (**GUI**) for
 - Instance creation
 - Management of configuration
 - Backup and replay
 - Scaling of deployments
- Full **integration** into the CYBERTEC **support platform**

The entire solution is easy to deploy using a **single Helm chart**.

TRANSITIONING MADE EASY

The transition from MS SQL to Babelfish works as follows and is really simple to facilitate.

Preparation phase:

- Creation of a secure Babelfish **cluster in Scalefield**
- **Move** the the **database**
 - **Dump** MS SQL database
 - **Reload** into Babelfish
- Testing and **quality**
 - **Connect the application** to the new IP
 - Run extensive **quality management**
- Make a decision
 - Stick with MS SQL
 - **Migrate to Babelfish** (as many as possible)

Real migration:

- Deploy on **Scalefield** or as **RPMs / Debs / Docker / K8S operator**
- Dump and reload data
- Connect the application to the new system
- **Run in production**

Servers can be added to Scalefield easily. A major advantage here is that Kubernetes allocates capacity efficiently. This transition will likely reduce not only the cost of MS SQL but also the overall hardware requirements by consolidating workloads onto existing servers.

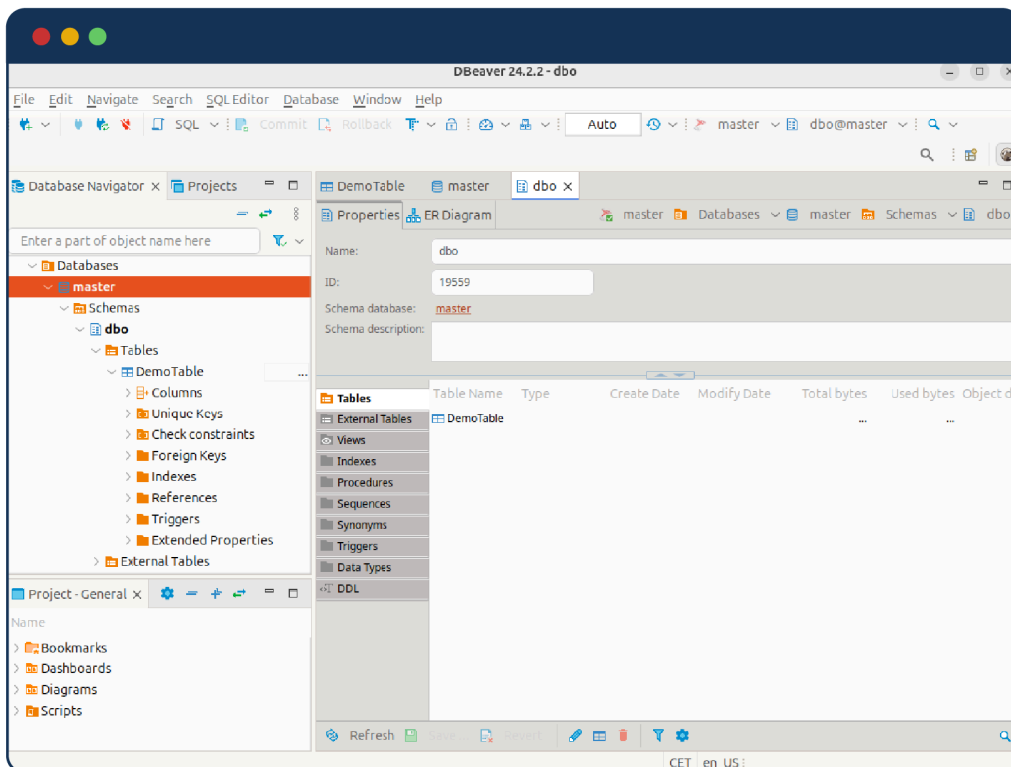
BABELFISH IN ACTION

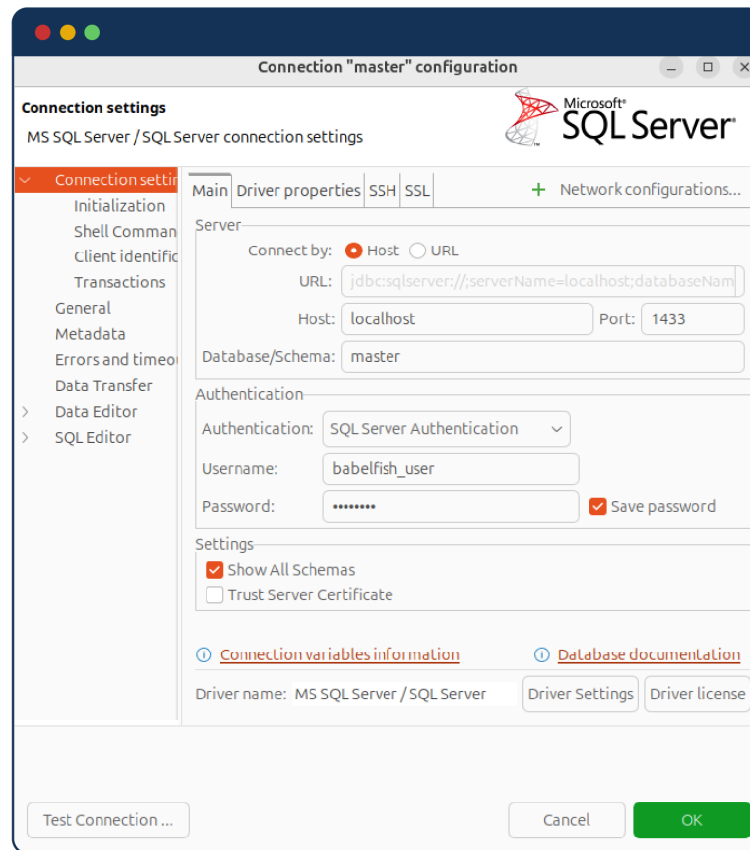
To see this concept in action, let's take a look at the UNIX process table:

```

40436 ?      Ss      0:00  \_ /bin/sh /start.sh
40518 ?      S       0:00  \_ ./postgres -D /data/babelfish/ -i
40519 ?      Ss      0:00  \_ postgres: checkpointer
40520 ?      Ss      0:00  \_ postgres: background writer
40522 ?      Ss      0:00  \_ postgres: walwriter
40523 ?      Ss      0:00  \_ postgres: autovacuum launcher
40524 ?      Ss      0:00  \_ postgres: logical replication launcher
41001 ?      Ss      0:00  \_ postgres: babelfish_user babelfish_db 172.17.0.1(41820) idle
41002 ?      Ss      0:06  \_ postgres: babelfish_user babelfish_db 172.17.0.1(41830) idle
    
```

It reveals the existence of PostgreSQL processes and real connections inside a Docker container (process isolation). In addition, we can connect to this database using DBeaver via the standard JDBC driver for MS SQL:





ONGOING SUPPORT AND MAINTENANCE

Our services include several critical aspects and ensure smooth 24/7 operations of the infrastructure and deployment.

SUPPORT SERVICES

We offer **CYBERTEC PostgreSQL 24/7 Enterprise Support**, which includes key features, such as:

- High-priority tickets: **30 minutes** response time
- Medium priority tickets: 4 hours response time
- Low priority tickets: Next business day

For more information, refer to our “CYBERTEC Support Policies” (version 1.1) which will be amended to include Babelfish.

We provide **support 24/7** in English.

TRAINING AND KNOWLEDGE TRANSFER

Everything outlined in this document comes with:

- **Knowledge transfer** to local, South African staff
- Documentation and **training**
- Access to skilled resources

We strongly believe in **upskilling local people** as it greatly **increases efficiency** for both our clients and ourselves (“**win-win**”).

ONGOING DEVELOPMENT

We are continuously developing our platform and committed to a long term, fruitful cooperation built on mutual trust and constructive feedback.

Although Babelfish is not a 100% replacement for MS SQL, it supports a vast majority of the commonly used features in the most compatible way possible. This makes it an alternative for potentially thousands of servers at low cost, minimal maintenance, and a smooth transition to PostgreSQL on a flexible timeline.

As the Scalefield platform already allows for full automation, transitioning entirely to PostgreSQL does not require additional changes in architecture.

As part of this agreement, we will work with the Babelfish **community to contribute** to the project and to push important changes forward.

CONTACT INFORMATION

AUSTRIA (HQ)

CYBERTEC POSTGRESQL
INTERNATIONAL (HQ)

SWITZERLAND

CYBERTEC POSTGRESQL
SWITZERLAND

URUGUAY

CYBERTEC POSTGRESQL
SOUTH AMERICA

ESTONIA

CYBERTEC POSTGRESQL
NORDIC

POLAND

CYBERTEC POSTGRESQL
POLAND

SOUTH AFRICA

CYBERTEC POSTGRESQL
SOUTH AFRICA

CYBERTEC PostgreSQL International (HQ)

Römerstraße 19
2752 Wöllersdorf
Austria
Phone: +43 (0)2622 93022-0
E-Mail:
sales@cybertec-postgresql.com

CYBERTEC PostgreSQL Nordic

Fahle Office
Tartu mnt 84a-M302
10112 Tallinn
Estonia
Phone: +372 712 3013
E-Mail:
sales@cybertec-postgresql.com

CYBERTEC PG Database Services South America S.A.

Misiones 1486, Piso 3
11000 Montevideo
Uruguay
E-Mail:
sales@cybertec-postgresql.com

CYBERTEC PostgreSQL Switzerland

Bahnhofstraße 10
8001 Zürich
Switzerland
Phone: +41 43 456 2684
E-Mail:
sales@cybertec-postgresql.com

CYBERTEC PostgreSQL Poland

Pl. Inwalidów 10
01-552 Warsaw
Poland
E-Mail:
sales@cybertec-postgresql.com

CYBERTEC PostgreSQL South Africa

No. 26, Cambridge Office Park
5 Bauhinia Street, Highveld Techno
Park
0046 Centurion
South Africa
Phone: +27(0)012 881 1911
E-Mail:
sales@cybertec-postgresql.com



If you need further information

For more information, or if you have any questions about our range of products, tools and services, contact us. There's no obligation—send us an inquiry via email or give us a call.



Contact

 **CYBERTEC PostgreSQL International GmbH**
Römerstraße 19
2752 Wöllersdorf
AUSTRIA

 + 43 (0) 2622 93022-0
 sales@cybertec-postgresql.com