CYBERTEC
POSTGRESQL SERVICES & SUPPORT

CYBERTEC SCALEFIELD

TECHNICAL GUIDE

Document version: 1.0
Last change: 2024-12-09

TABLE OF CONTENTS

SCALEFIELD: COMPREHENSIVE POSTGRESQL

Scalefield is a unified platform built on Kubernetes / OpenShift that integrates all our solutions into a single, easy-to-use system.

This document provides an overview of the general architecture and highlights use cases to help customers utilize the platform as efficiently as possible.

SYSTEM ARCHITECTURE

This section describes the general system architecture of CYBERTEC Scalefield and outlines the deployment options available to customers.

TECHNICAL REQUIREMENTS

Scalefield is capable of operating on various platforms, which include:

- Bare metal servers
- Virtual machines (VMWare, Promox, etc.)
  - On-premise
  - Cloud deployments
- Kubernetes / OpenShift / Rancher
  - On-premise
  - Cloud deployments

The following system architectures are supported:

- Linux on x86_64 (Intel, AMD)
  - Ubuntu 22.04 or higher
  - Debian 12 or higher
  - RedHat compatible Linux (RHEL 8/9 or higher)
- Linux on ARM

> **NOTE:** For local test deployments, we offer a standalone OS X version (with some restrictions), mostly intended for demonstration purposes.

Scalefield requires a minimum of **16 GB or RAM per Kubernetes node**, with significantly more recommended for large deployments.

Support versions of Kubernetes: v1.21 - 1.28.3

SCALEFIELD CLUSTER LAYOUT

Kubernetes provides us with all the scalability needed to operate Enterprise database deployments at scale. The general layout of a cluster running Scalefield is as follows:

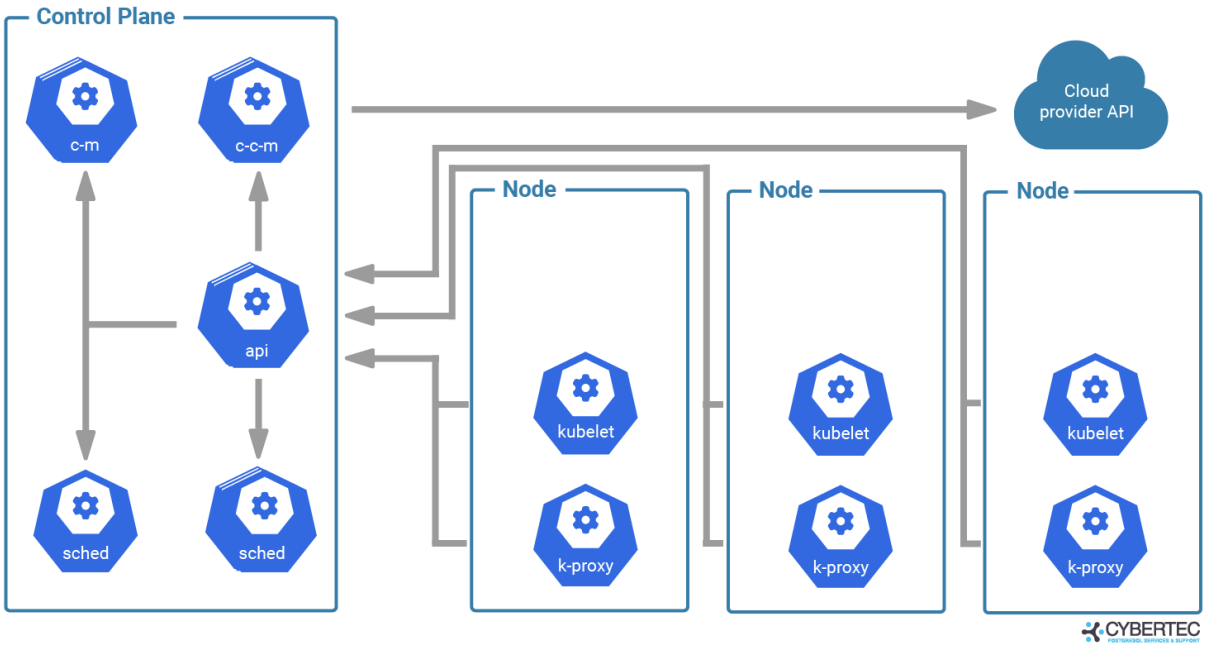

On top of Kubernetes / OpenShift container orchestration, Scalefield will run all services essential to a cutting edge user experience. This includes various components, including:

- Graphical user interface
- Backup and recovery automation
- Inventory database
- CYBERTEC Services
  - PostgreSQL
  - CYBERTEC Enterprise PostgreSQL
  - CYBERTEC Migrator
  - CYPEX low code
  - Babelfish MS SQL compatibility
- Monitoring and log management
- Compliance and security modules

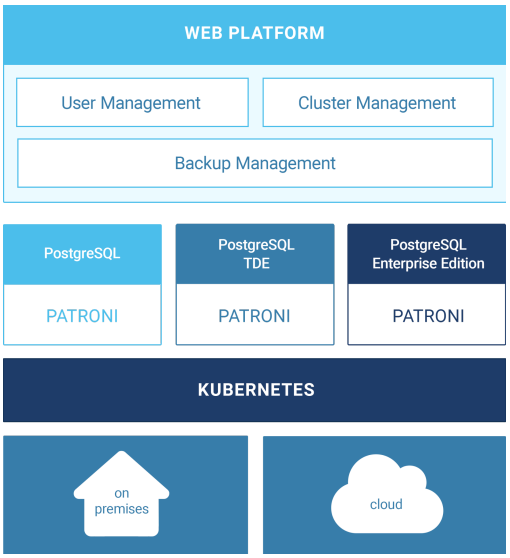

SCALEFIELD STORAGE OPTIONS

Scalefield supports various storage options provided by Kubernetes.
The following options are available:

Storage type	Performance	Scalability
Local storage	Extreme performance	Limited by local disk
Distributed block storage	High performance	High-Scalability
SAN storage	Extreme performance	High-Scalability
Cloud storage	-	High-Scalability

Scalefield is based on Kubernetes-compatible services, including OpenShift, SUSE Rancher, or publicly available cloud services, such as Amazon EKS.

As far as storage is concerned, Scalefield needs a "PVC" (= Persistent Volume Claim) provided by Kubernetes.

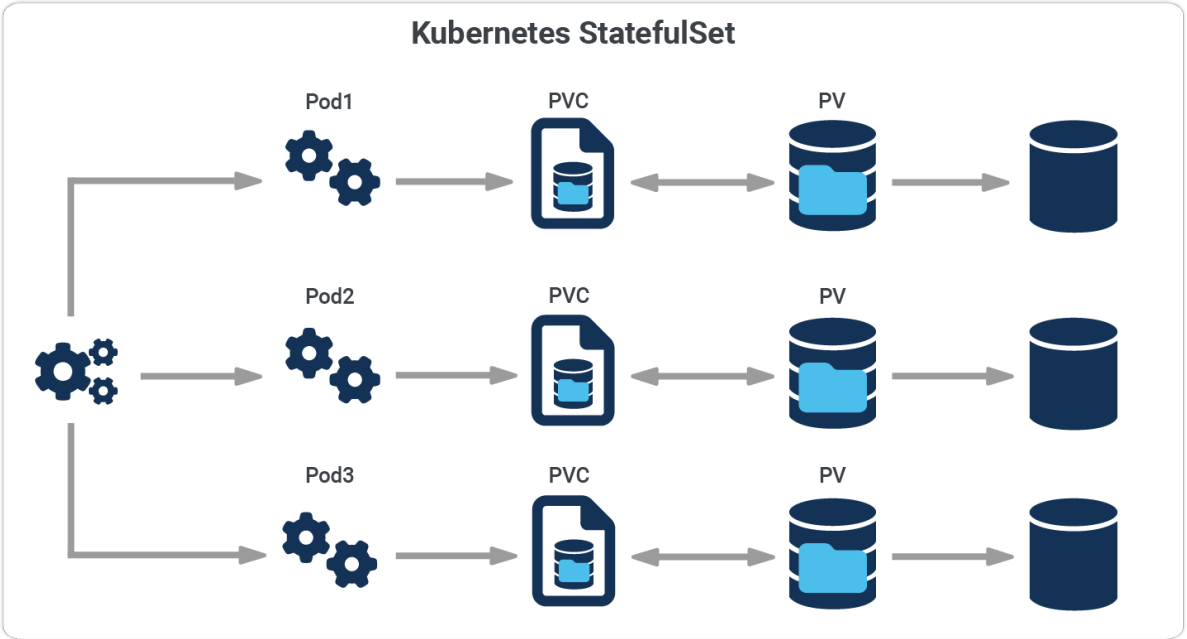

Scalefield can operate with storage classes of your choice to run small, as well as scalable database services.

Kubernetes / OpenShift provide all the abstraction needed to integrate Scalefield with Enterprise storage solutions:

- Dell EMC PowerMax and PowerStore
- Hewlett Packard Enterprise (HPE) Primera and Nimble
- Hitachi Vantara Virtual Storage Platform (VSP)
- IBM FlashSystem
- Infinidat InfiniBox
- NetApp All Flash FAS (AFF)
- Pure Storage FlashArray

This capability enables Scalefield to operate in diverse environments with scalability while adhering to industry standards.

HIGH-AVAILABILITY BY DESIGN

Database work is by definition, critical, and downtime is not acceptable. To reflect this reality, Scalefield provides High-Availability by design and out of the box.

> "High-Availability by design from the start"

By default, Scalefield allows you to configure the desired level of redundancy. Simply move the slider and make the system handle as many replicas as needed:

Behind the scenes, Scalefield will deploy a **Patroni cluster** which provides us with

- Automatic failover
- Fully automated recovery
- Shared nothing architecture

Scalefield will deploy various PostgreSQL nodes inside the Kubernetes cluster, ensuring 100% redundancy in case of a node failure. By careful orchestration, Scalefield makes sure that all members of the cluster reside on different host systems to guarantee fault tolerance and scalability at every level:

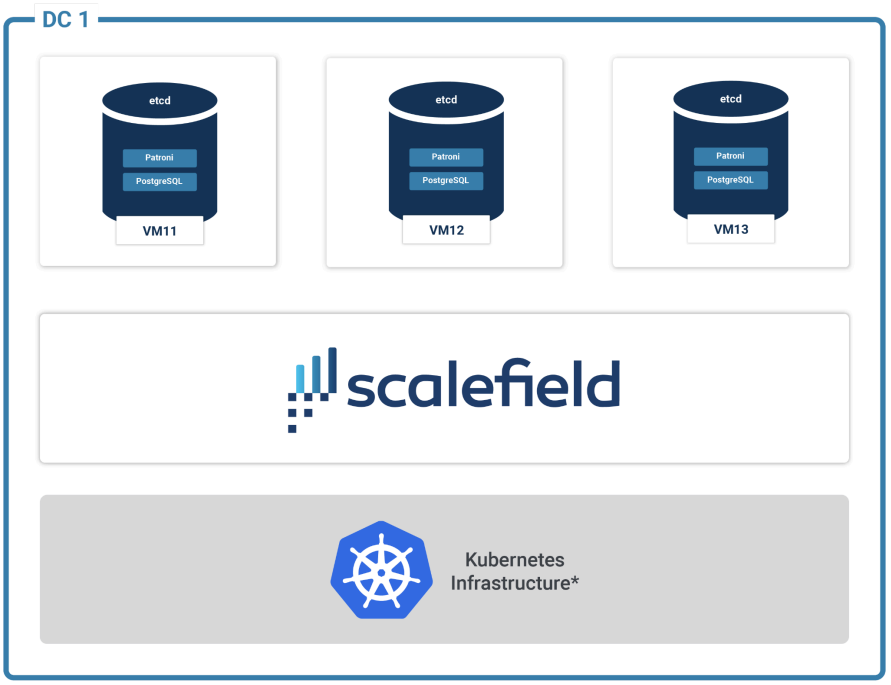

*Can be Kubernetes or a Kubernetes distribution such as OpenShift or Rancher.

To external applications, Scalefield exposes the cluster using different IPs and ports, simplifying the use of replicas for purposes such as:

- Read scalability and load balancing
- Analytics and statistics

ORCHESTRATION AND API ACCESS

Scalefield not only provides an intuitive user interface but also enables seamless integration through ready-to-use API endpoints, simplifying connectivity with other systems and applications.

The Scalefield REST API offers a JSON-based interface, exposing all key aspects of the infrastructure for advanced orchestration.

SCALEFIELD PRODUCTS

Scalefield incorporates the main product stack, including:

- Scalefield PostgreSQL Automation:
  - PostgreSQL as a Service for Kubernetes / OpenShift
  - PGEE: PostgreSQL Enterprise Edition
- CYBERTEC Migrator: Oracle to PostgreSQL Migration
- CYPEX: Low Code Application Development
- Babelfish for PostgreSQL: MS SQL compatibility for PostgreSQL
- pg_deep_thinker: "Machine driven PostgreSQL consulting"
- Scalefield Secure: Compliance Monitoring for Oracle, PGEE and PostgreSQL

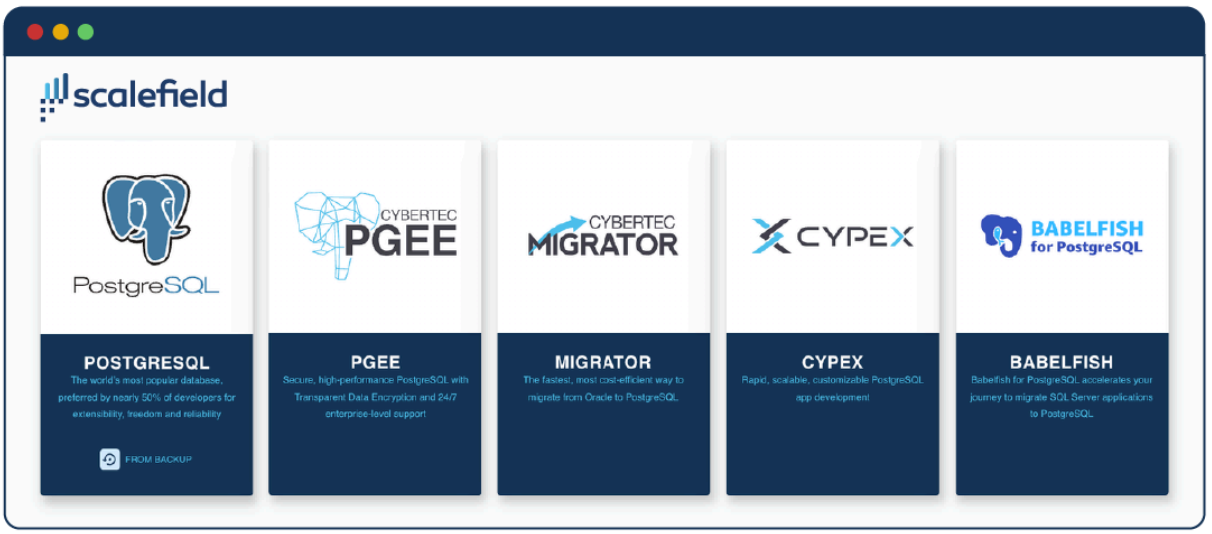

Scalefield is a 100% Kubernetes-based solution that allows customers to directly deploy their own private cloud and integrate with public clouds. It is 100% visualized.

All products are also available separately. However, **full integration** offers customers numerous benefits that are not available when deployed separately:

- A single **compliance report**
- **Automatic backups** for all products
- Integrated automation
- Faster **security updates**
- **Better** user **experience**

SCALEFIELD POSTGRESQL AUTOMATION

Scalefield PostgreSQL **Automation** allows you to quickly deploy various flavors of PostgreSQL, such as PostgreSQL and CYBERTEC PostgreSQL Enterprise Edition (**PGEE**). Quickly **deploy** and automate your database platform in a **user friendly** environment:

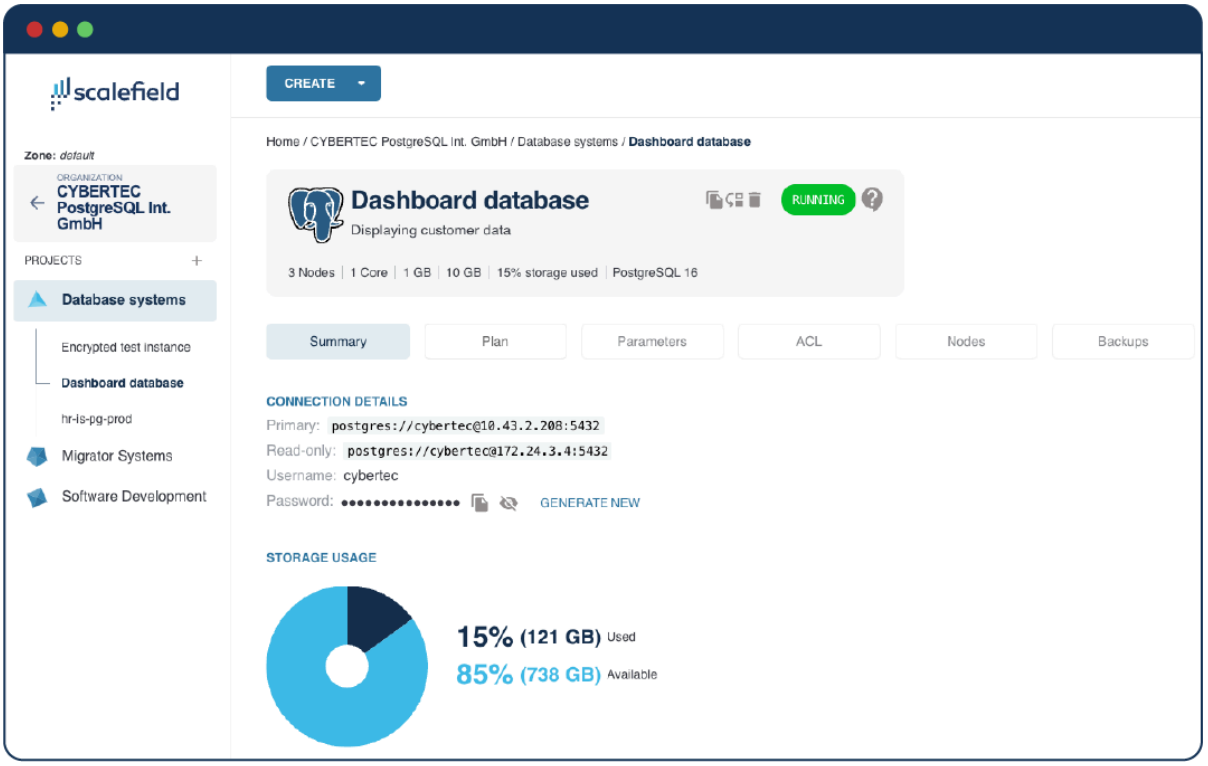

Scalefield PostgreSQL Automation provides ready-to-use services, including:

- High-Availability out of the box
- Easy scalability and dynamic resizing
- Integration with backup and recovery tooling
- Automated monitoring
- Compliance enabled

kubernetes OPENSHIFT

PGEE: CYBERTEC ENTERPRISE POSTGRESQL

PGEE is a commercial version of PostgreSQL that allows customers to make use of advanced features, such as

> "Transparent Data Encryption"

It provides countless features which make it **attractive to clients** across various industries, including but not limited to:

- **Finance** and **Banking**
- High-**Security** computing
- Medical and health
- **Government** and law enforcement agencies

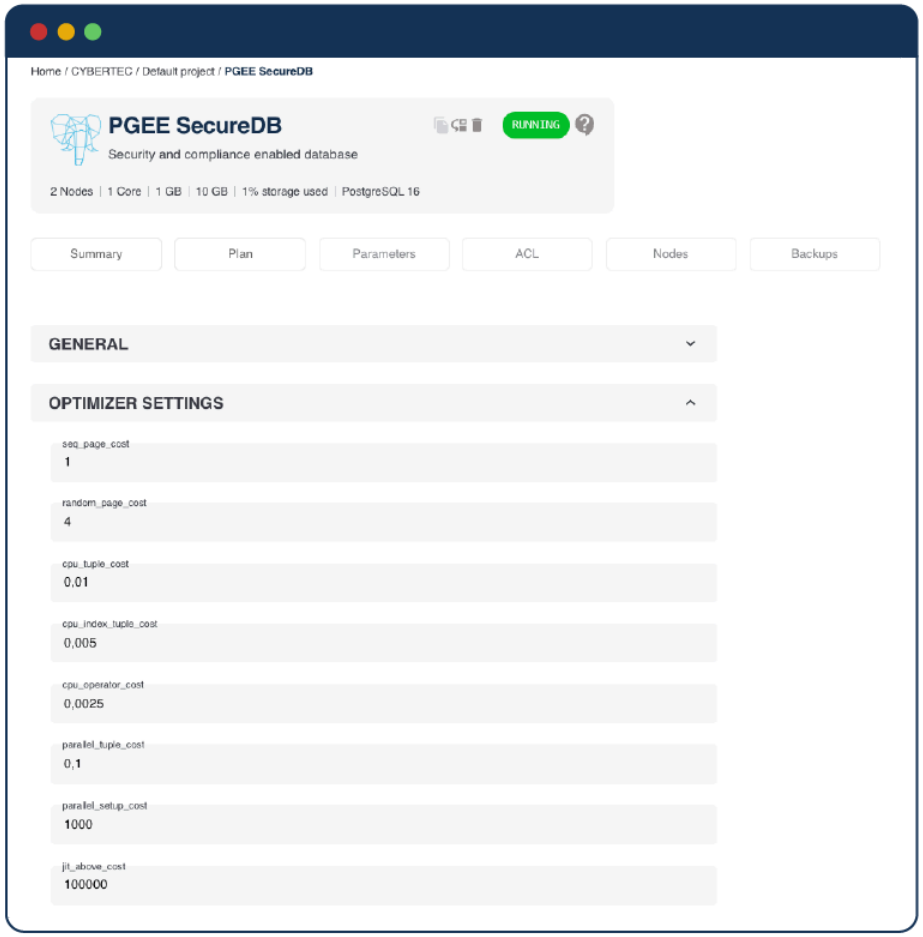

We offer **additional features** on top of PostgreSQL, such as:

- Transparent Data **Encryption** (TDE)
- Stored Procedure **Encryption**
- Enterprise-level **auditing**
- Anomaly detection
- Data masking and **obfuscation**
- **Advanced optimizer** improvements

CYBERTEC PostgreSQL Enterprise Edition (PGEE) is important to all clients needing extra security as well as **compliance** with modern IT standards:

- ISO 27001
- TISAX
- SFCA
- POPI-A
- GDPR
- HIPAA
- … and many more …

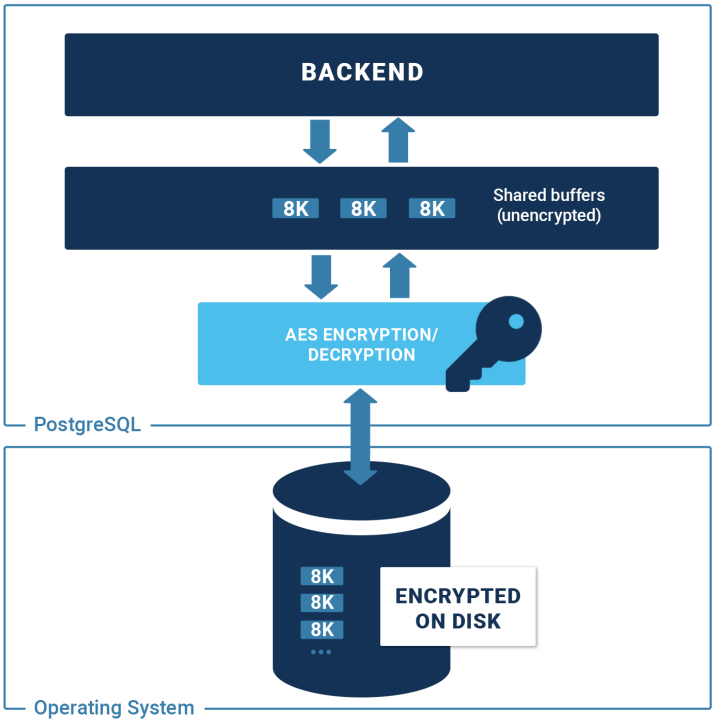

"Professional database compliance for PostgreSQL"

Fixing real world problems

CYBERTEC MIGRATOR FOR ORACLE

The **CYBERTEC Migrator** is the entry ticket into our ecosystem, enabling customers to migrate their legacy databases to a new infrastructure.

We offer expert level tooling for high-performance migration, designed to serve various purposes:

- **Bulk assessment** for large **Oracle** customers
  - Get a handle on **big infrastructures** in no time
  - Provide a **migration roadmap** to clients
- **Enable partners** and **customers** to migrate
  - Provide expert tooling
  - Help with quick migration
- **Migrate** to PostgreSQL **at scale**
  - Cutting edge transfer speeds (> 1.4 GB / sec)
  - Optional zero downtime migration

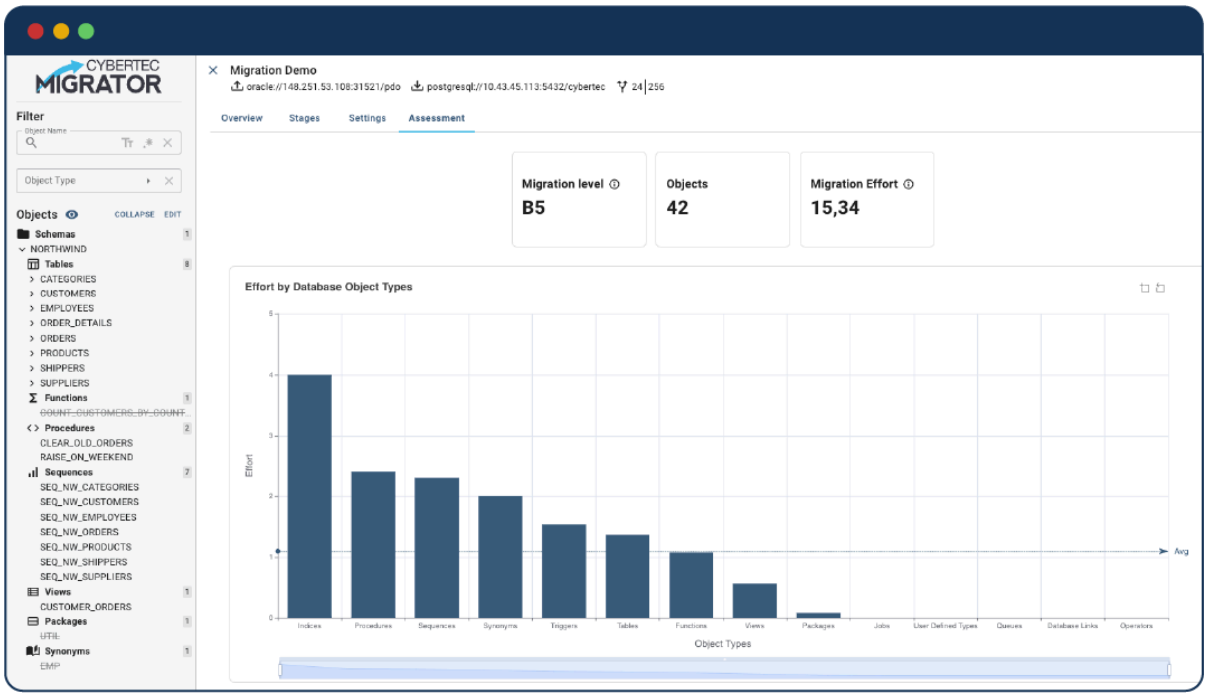

Migration from Oracle to PostgreSQL has never been easier:

- Move Oracle to cloud solutions (Azure, GCP, AWS, etc.)
- Migrate to on-premise solutions

Stored procedures have always been a major hurdle when migrating away from Oracle. CYBERTEC Migrator simplifies this process with modern **AI technology**:

"AI assisted code explanation"

With AI-assisted code explanation, developers can understand complex code faster and more efficiently:

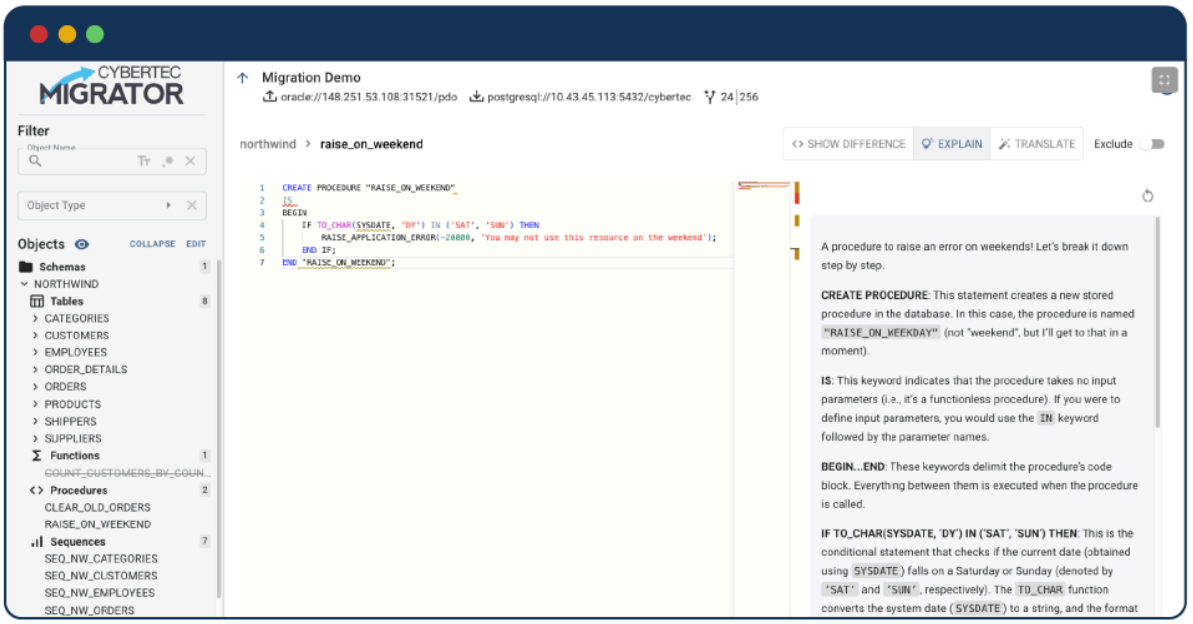

The CYBERTEC Migrator assists with **code** pre-**translation** by using modern technology tailored to the problem.

We fully comply with EU AI regulations. All code processing happens within Scalefield, ensuring that no customer source code is sent to public APIs. This means:

- Maximum **compliance** with regulations
- Maximum **digital independence**
- **Zero** external **dependencies**

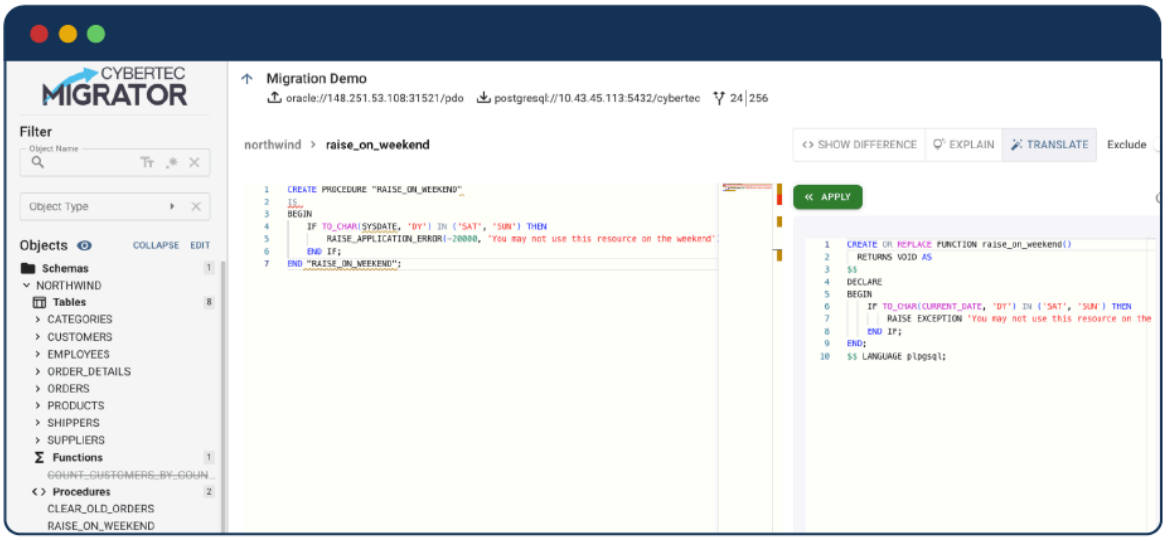

Code translation has many **advantages**:

- **Migrate** code up to 80% **faster**
- Improve code **quality**
- **Finish** migration projects faster
- **Enable developers** to understand code more quickly

CYPEX: LOW CODE DEVELOPMENT FOR POSTGRESQL

With **CYPEX**, you can build customer specific applications in remarkably short timeframes. Workflows, dashboards, maps, forms – CYPEX makes it all possible within minutes.

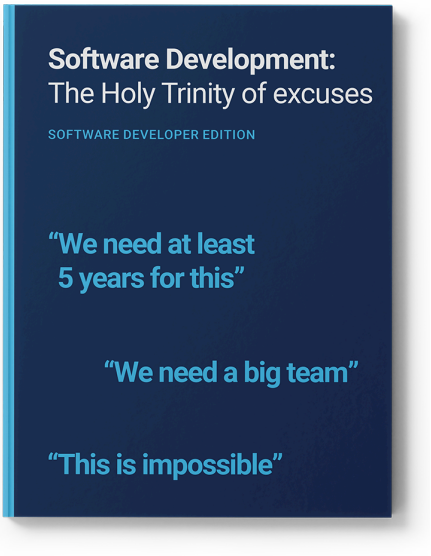

What **problem** are we **solving**?

Modern software development is often hindered by

- Information loss
- Repetitive tasks
- Inefficient requirements gathering
- Slow feedback

CYPEX is addressing these challenges. We enable the creation of data models in hours rather than days, using those data models to predict applications quickly. These applications can then be modified as needed.

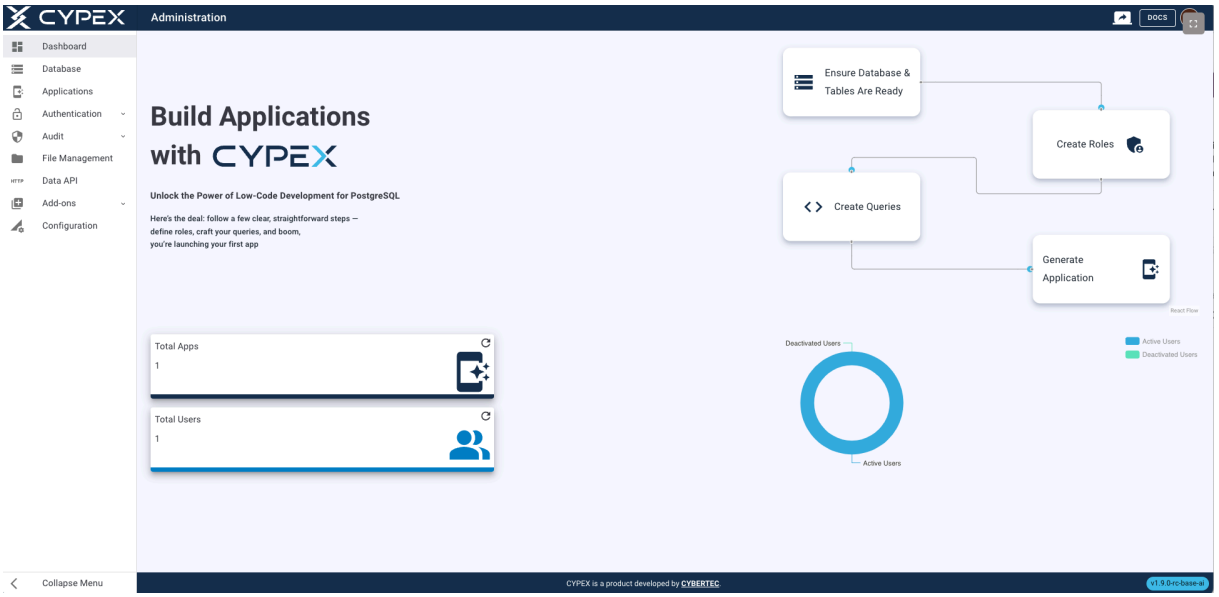

Our approach offers numerous benefits::

- Gather specifications **quickly**
- Quick **prototyping** and efficient demos
- Minimal time to market
- **Build applications cheaper** than ever before

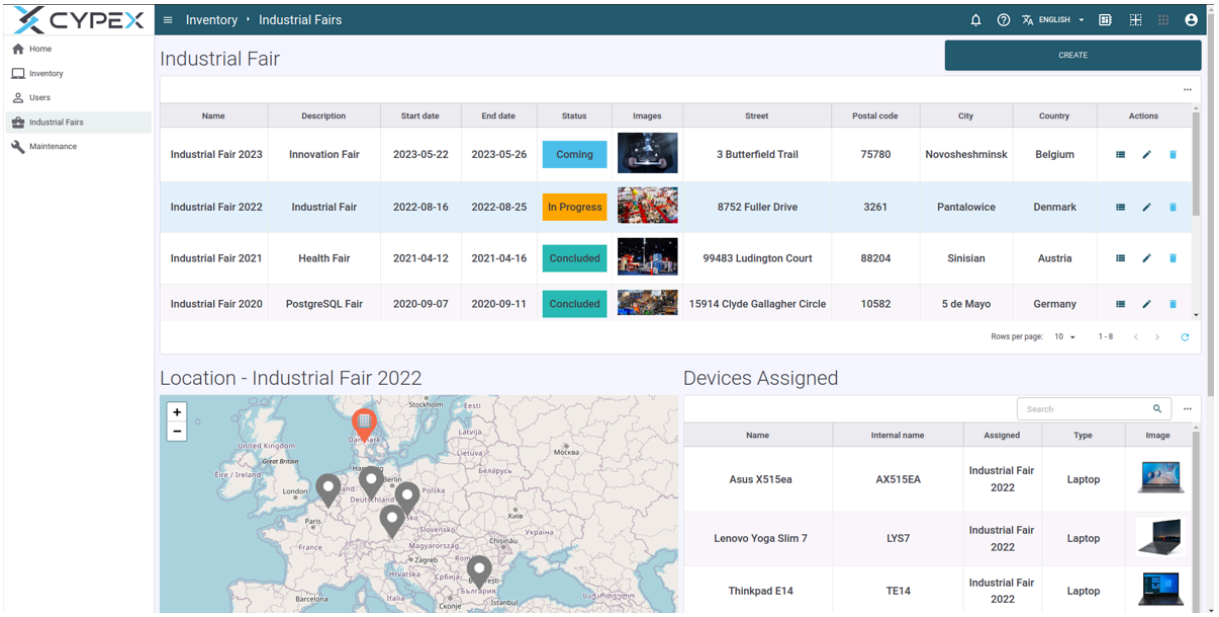

BABELFISH FOR POSTGRESQL: BRINGING AN END TO MS SQL LICENSING

What are the challenges of Microsoft SQL Server?

- Escalating license costs
- High support cost
- Technically outdated legacy solution
    - Poor locking and concurrency behavior
    - Subpar transactional behavior

These issues drive many customers away from MS SQL. But what does this process look like in real life? Consider a customer running 1.000 MS SQL Server who is eager to migrate:

- They must talk to 200 application owners
- They must launch 200 migration projects
- The migration takes years to complete
    - Meanwhile, license cost keeps piling up

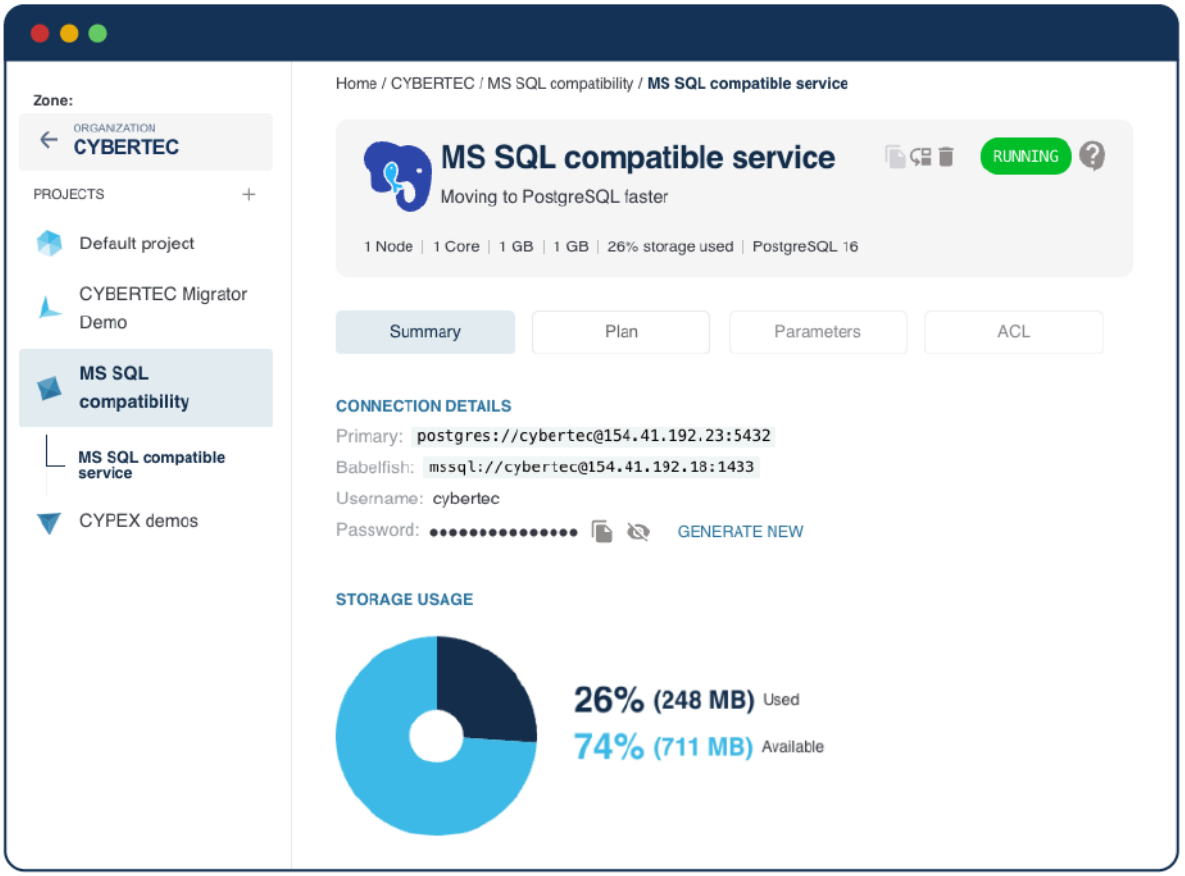

Babelfish for PostgreSQL comes to the rescue. It can run inside Scalefield providing countless advantages to the customer:

A cheaper solution WITHOUT touching the application

This means we can **move customers** to an MS SQL compatible solution (in 60-80% of all cases) **without launching a migration project**.

Key benefits include:

- Instant cost saving
- No migration project
- Improved automation
- Maintaining compliance

Babelfish behaves just like MS SQL making it an ideal solution for all clients aiming to cut expenses quickly:

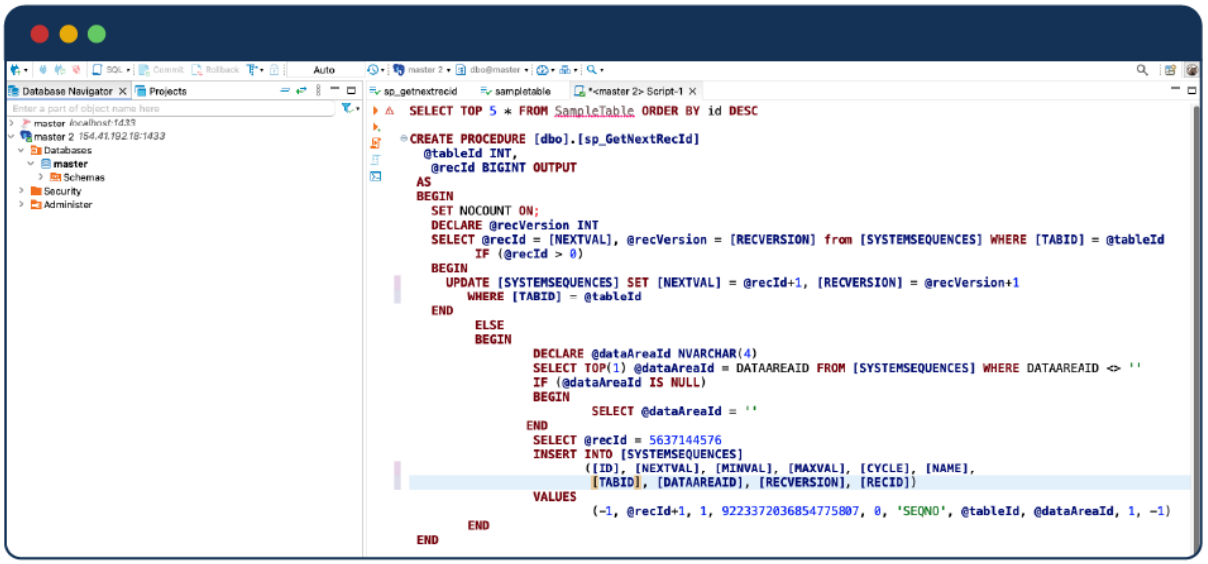

PROVING VALUE TO CUSTOMERS

Each product can be marketed individually; however, the power of integration enables us to offer a **single package** that includes everything:

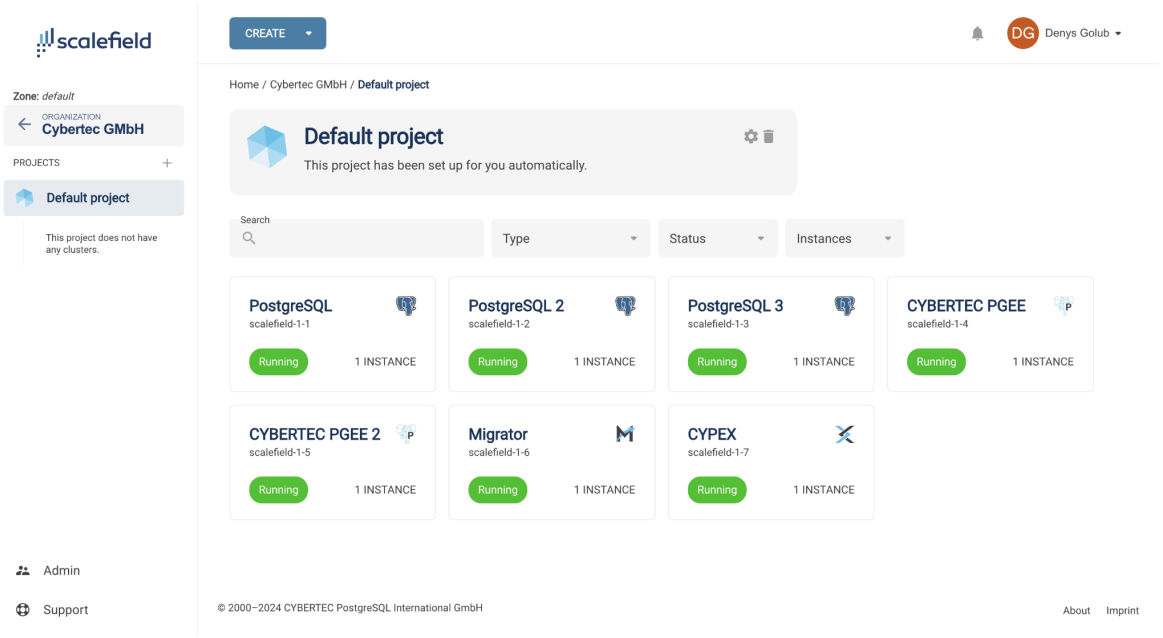

By marketing **Scalefield** as a product that allows a customer to deploy every single solution as part of a larger product, we can fully benefit from tight integration:

- Integrated **support** and consulting
- More oversight over **resource** usage
- More streamlined **customer experience**

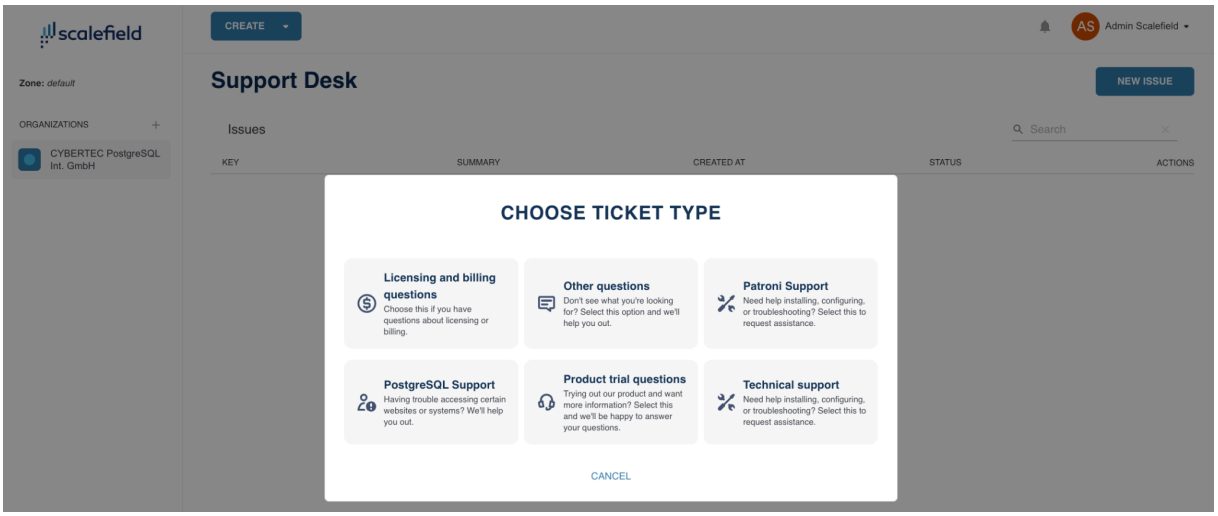

Finally, we fully integrate with **CYBERTEC support services** to ensure that customers have a single point of contact for all potential issues.

SERVICE STRATEGY: EXCELLENCE MATTERS

We at CYBERTEC focus heavily on **professionalism** at **every level**. This is reflected in our feedback, **SLA** statistics (no misses) and in overall customer **satisfaction** which helps us to achieve:

- **Long term** customer relationships
- Excellent customer **satisfaction**

We never have, and never will, sacrifice short-term profits over long-term **customer satisfaction**, as we consider unhappy customers to be the worst investment known to us.

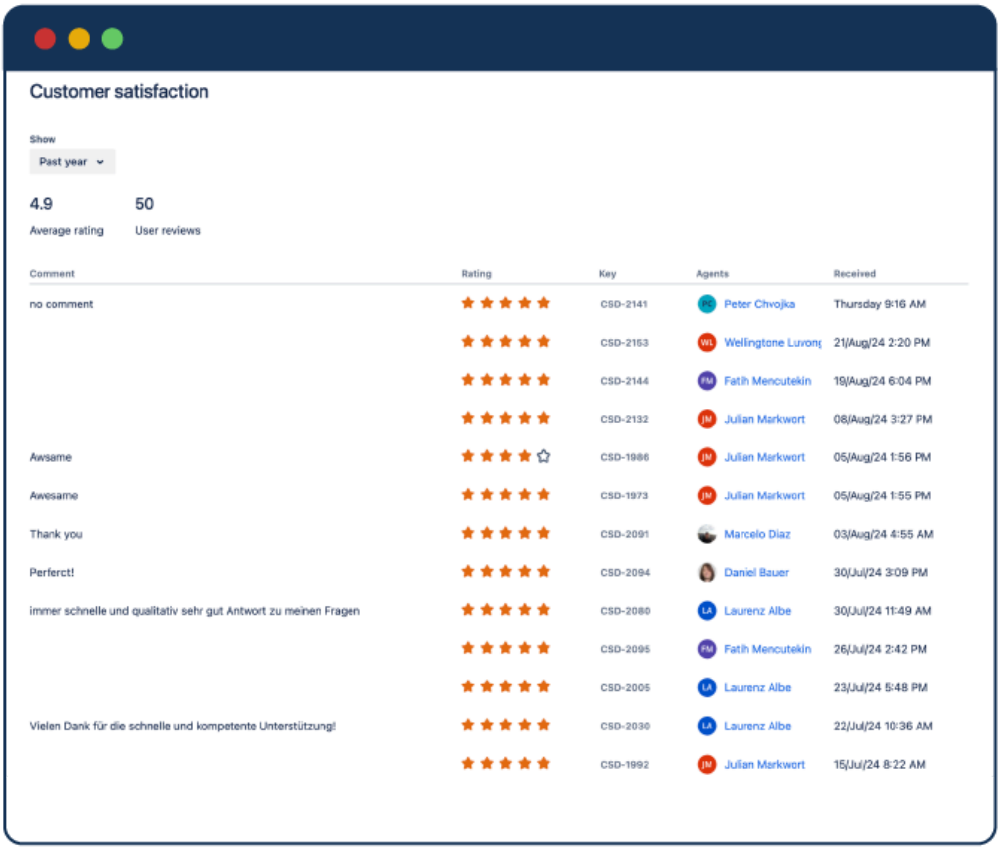

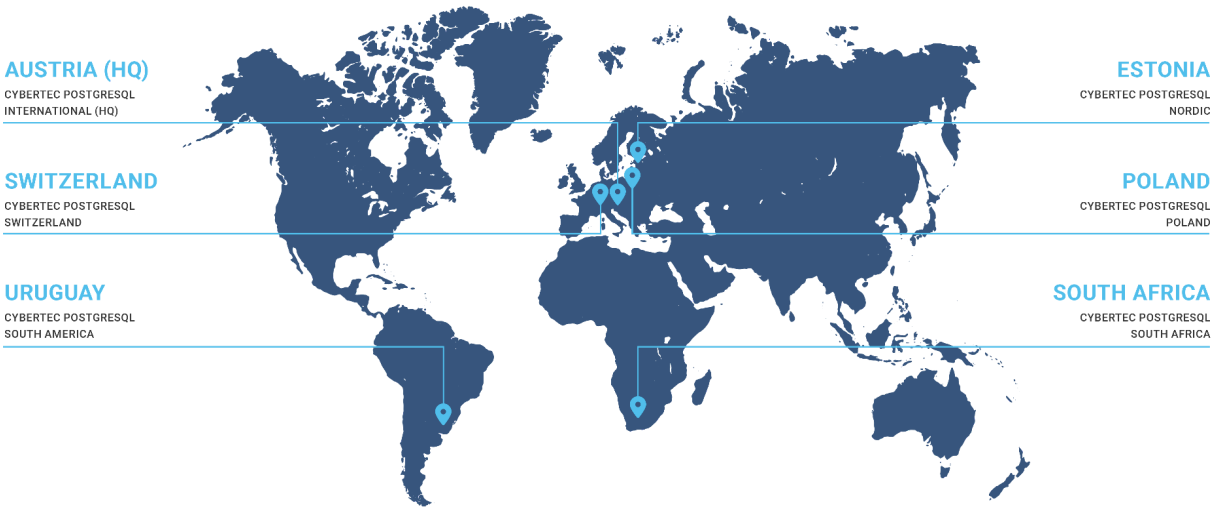

CYBERTEC PostgreSQL International (HQ)
Römerstraße 19
2752 Wöllersdorf
Austria
Phone: +43 (0)2622 93022-0
E-Mail:
sales@cybertec-postgresql.com

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.

ISO 27001 CERTIFIED

Contact

CYBERTEC PostgreSQL International GmbH
Römerstraße 19
2752 Wöllersdorf
AUSTRIA

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