

Open Source DBMS Platform Based on PostgreSQL

Tmax OpenSQL 3

Contents

01 / Tmax OpenSQL Product Introduction

02 / Cloud Optimized DBMS Platform, OpenSQL 3

03 / Tmax OpenSQL Key References

Annex / Tmax OpenSQL platform architecture and detailed key features

Chapter 1

Tmax OpenSQL Product Introduction

OpenSQL



Open Source DBMS Platform Based on PostgreSQL

Optimization with verified tools and extensions, and reliable maintenance supported by technical expertise

Reliable solution developed by DB experts

Monitoring, CDC, FDW, and other in-house tools developed by DB Specialists

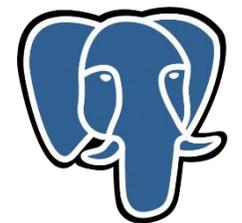
Extend and optimize features tailored to enterprise environment

Reliable product packaging and Optimized DB features such as redundancy, security, backup, and recovery

Stable technical support of 20 years of experience

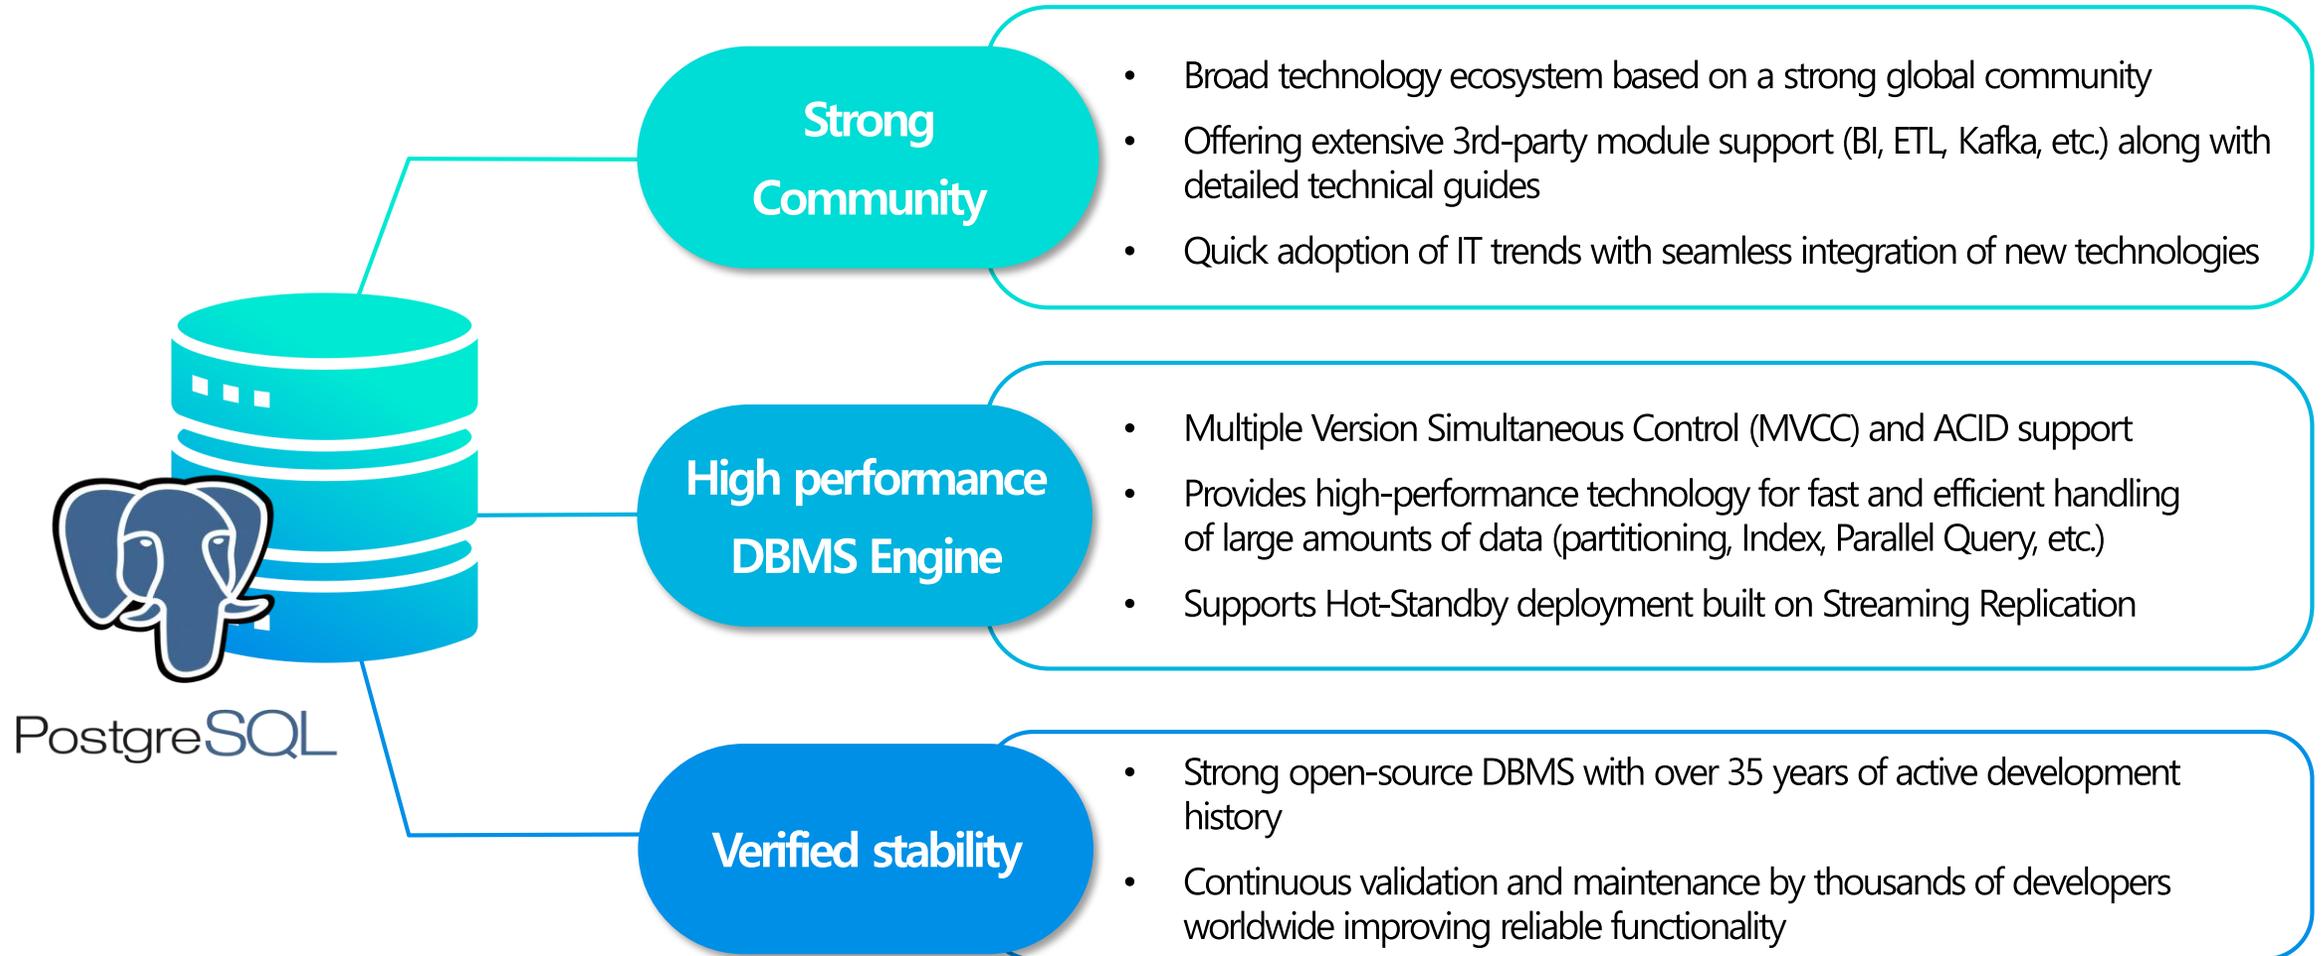
Best DB experts by industry
313 technical personnel, 72% of the total workforce

TMAXTibero

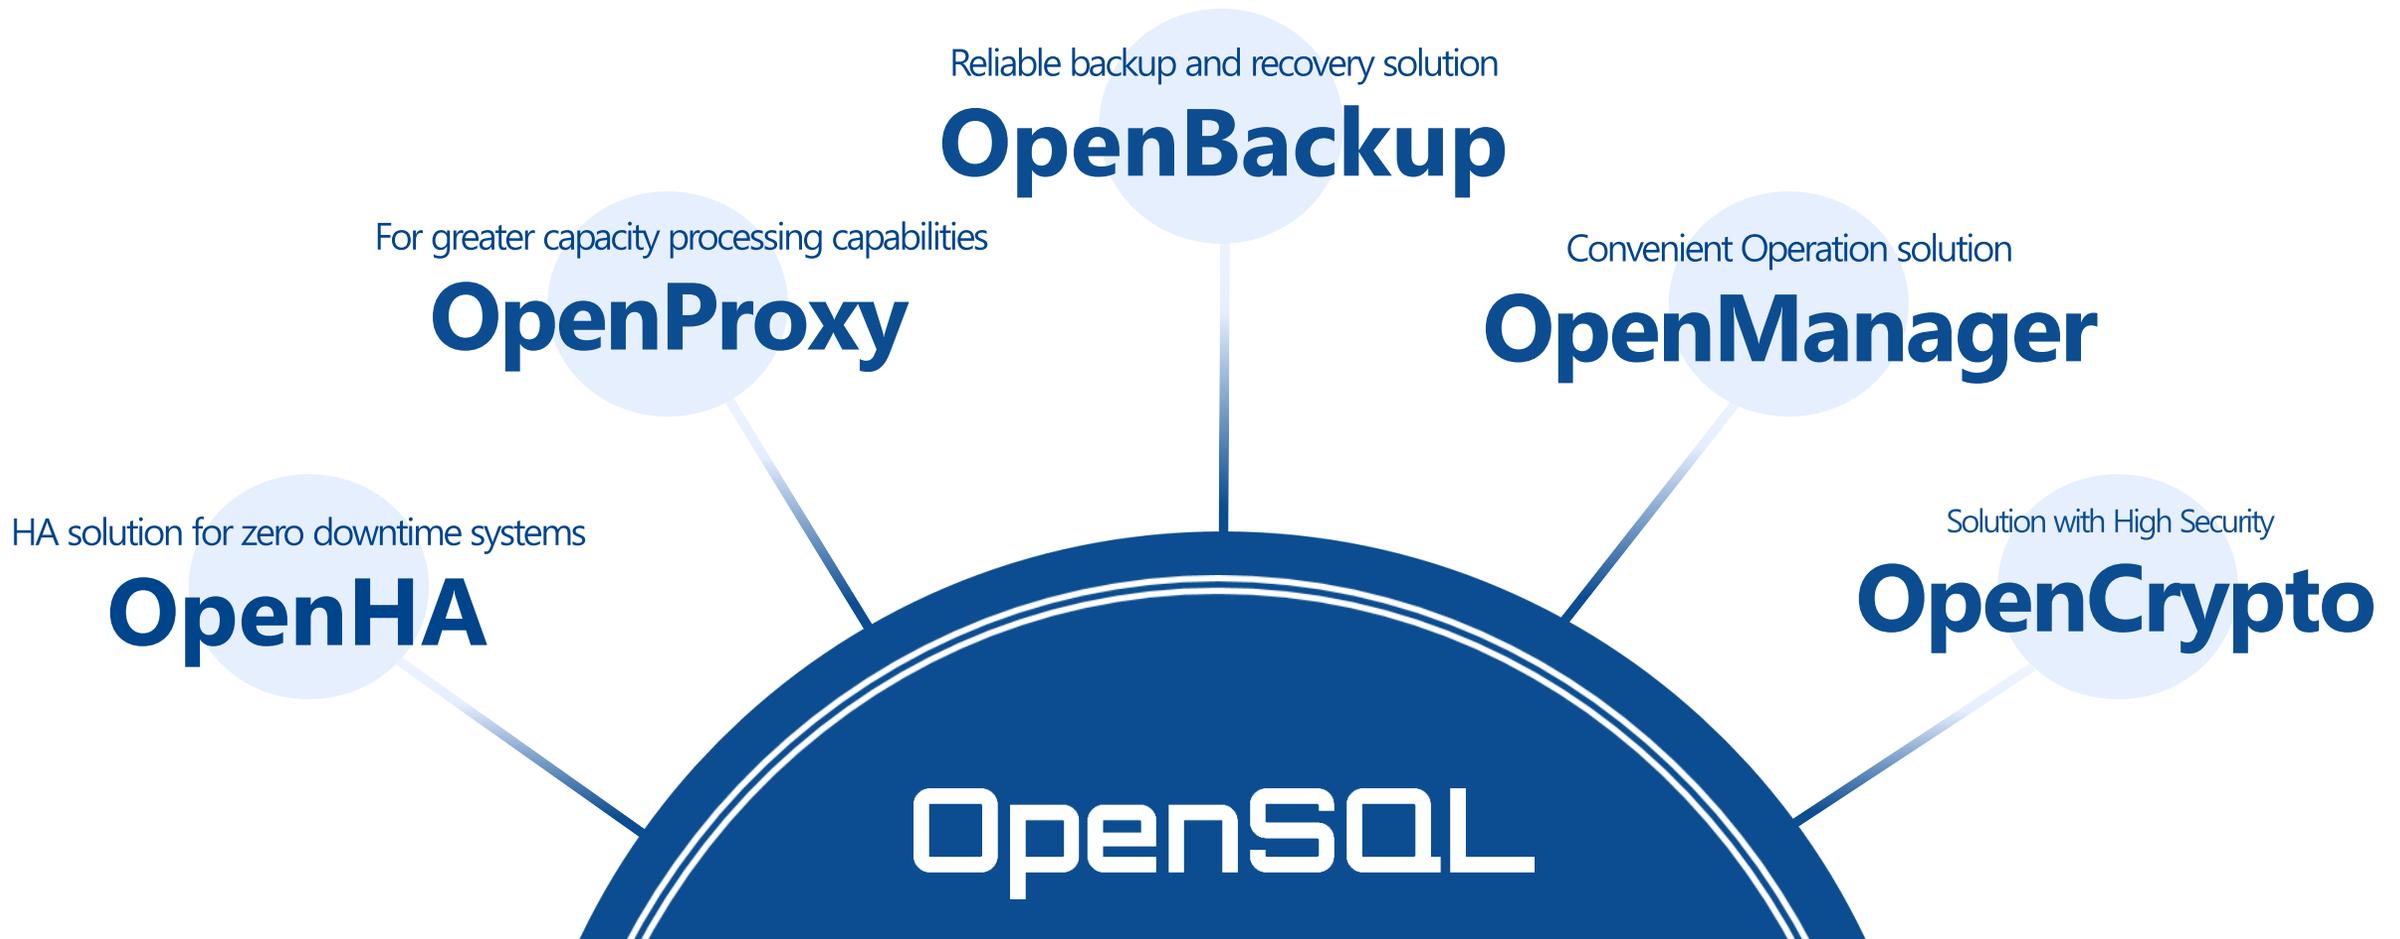


PostgreSQL

Built on PostgreSQL, OpenSQL leverages a robust community and exceptional performance



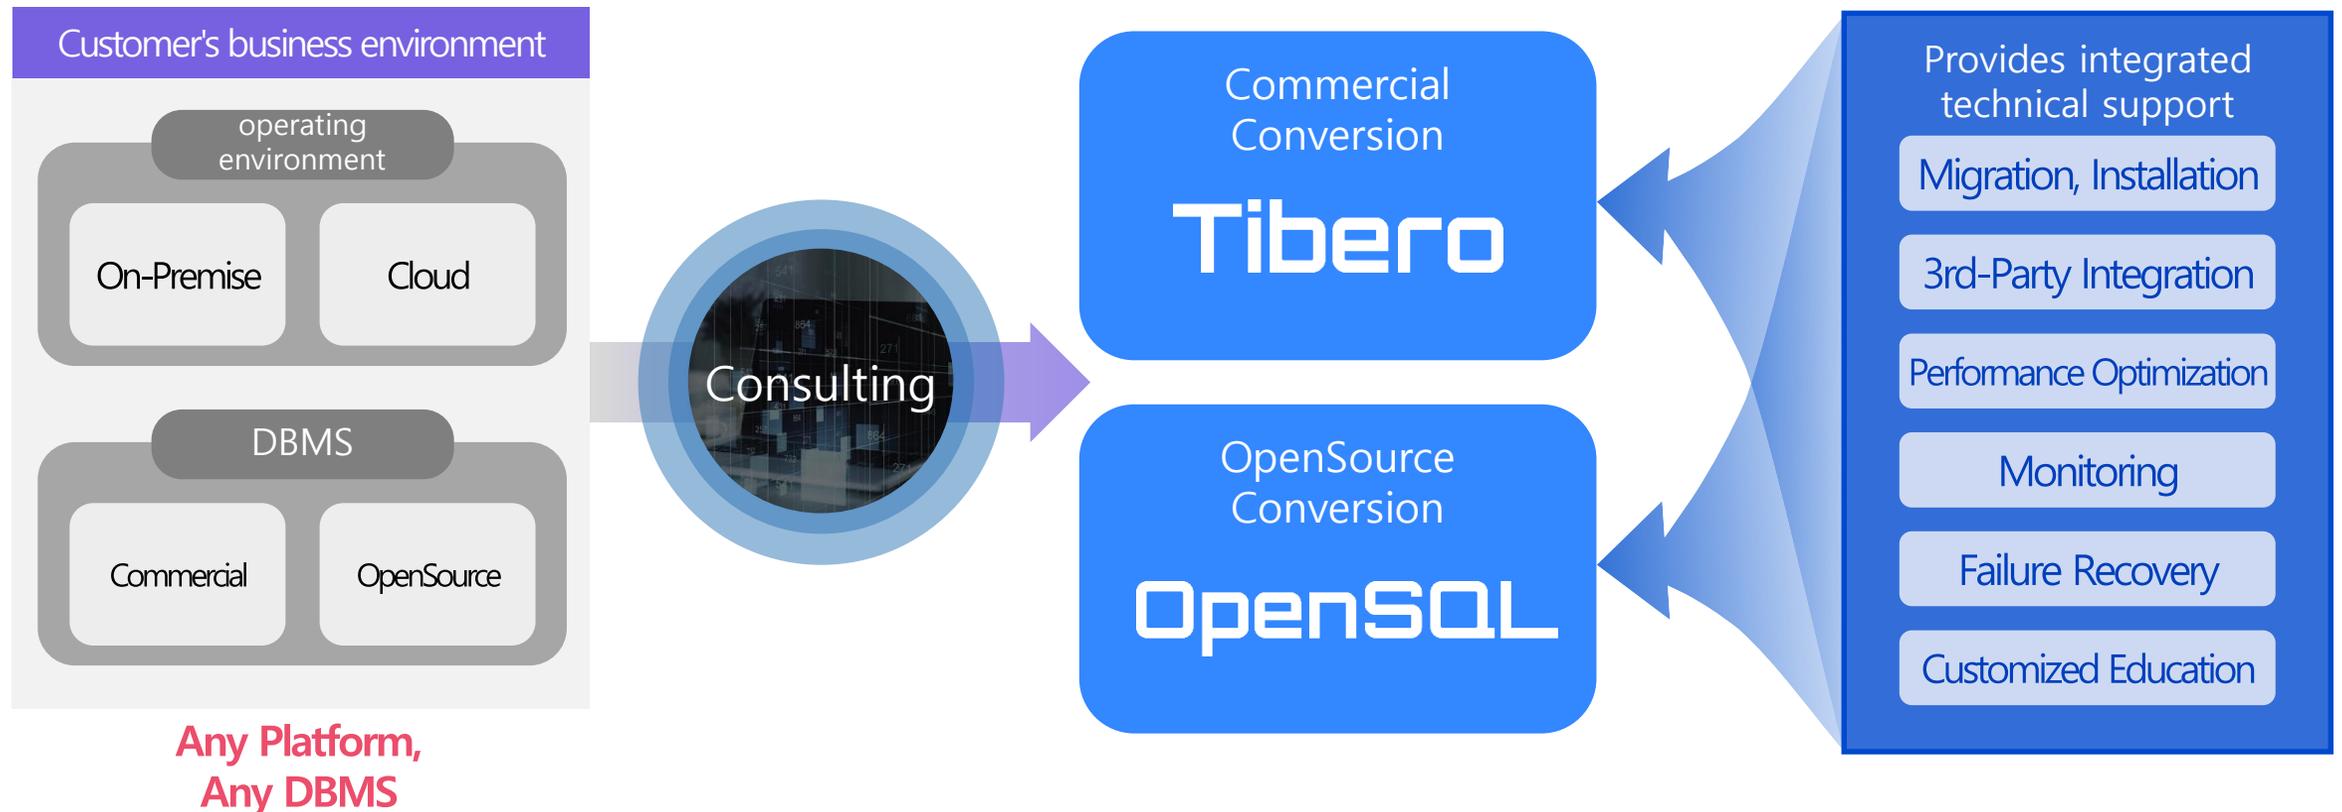
TmaxTibero, the expert in commercial DBMS environments, provides essential enterprise-ready features



Delivers reliable DBMS operations for clients with PostgreSQL-specialized technical support

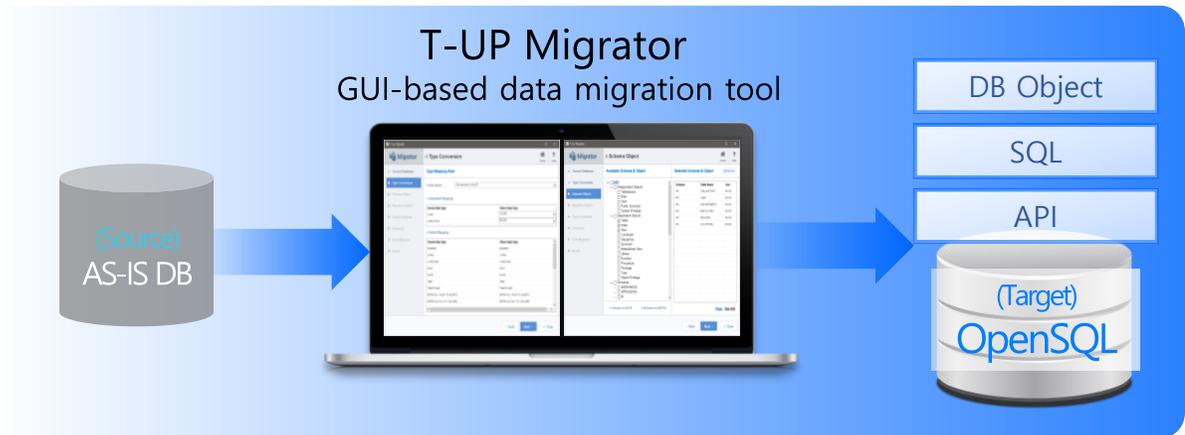
 PostgreSQL Community Version	vs.	OpenSQL
Difficulty selecting and validating extensions from diversity for commercial environments	Extension Verification	Extensions verified for production use, provided in a fully tested and optimized state
Limited technical support due to lack of systematization and manpower, resulting in delayed response	Issue Management	Provides systematic on/offline technical support for 24*7 Clear and immediate support in the event of a problem
Official documents exist, but they are highly technical with not enough practical guides	Technical Guides and Training Assistance	Provides manuals and operation/tuning guides, complemented by regular user training and seminars.
Only minimal default settings are configured during initial installation; all other options must be learned and adjusted	Tuning	Provides pre-optimized setups and professional guidance from performance analysis to tuning implementation

Comprehensive support for migrating customers' existing DBMS to TmaxTibero (commercial or open-source) and optimizing it for business use



Automated Migration Using T-UP Migrator

Provides an automated migration tool for stable integration and seamless transition with existing commercial environments.



Type Conversion

Schema Object

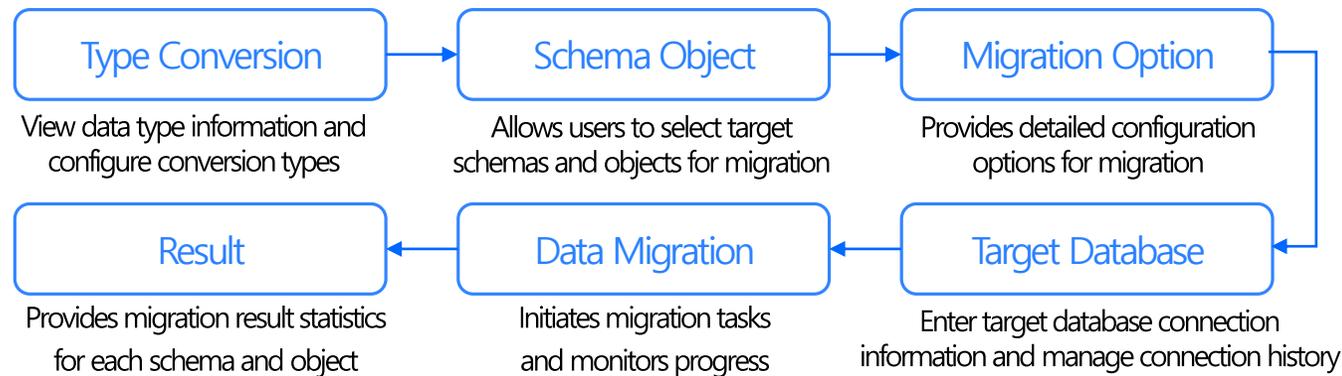
Migration Option

Target Database

Data Migration

Result

- DBMS Migration Execution (GUI based)
 - Automatically migrates user-selected schemas, objects, tables, and data



Chapter 2

OpenSQL 3, Cloud Optimized DBMS Platform

OpenSQL3

“PostgreSQL-based, cloud-optimized DBMS platform that enhances high availability, performance, security, and compatibility for enterprise environments”



Cloud Optimization

- ✓ Zero Downtime **Scale in / out**
- ✓ Real-time Cluster State Sharing and Automatic Failover
- ✓ Multi-Cloud Environment Support



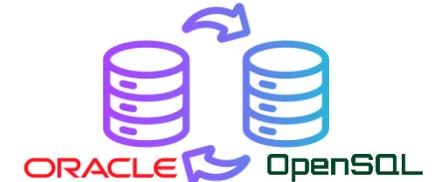
HA and Performance Improvement

- ✓ Improved high availability through architecture enhancements
- ✓ Enhanced connection pooler performance for large-scale processing
- ✓ Improved Cluster Manageability



Security Enhancement

- ✓ Support for Encryption Algorithms (ARIA, SEED)
- ✓ Rapid Patch Deployment in case of Security Vulnerabilities
- ✓ User Account and Profile Management

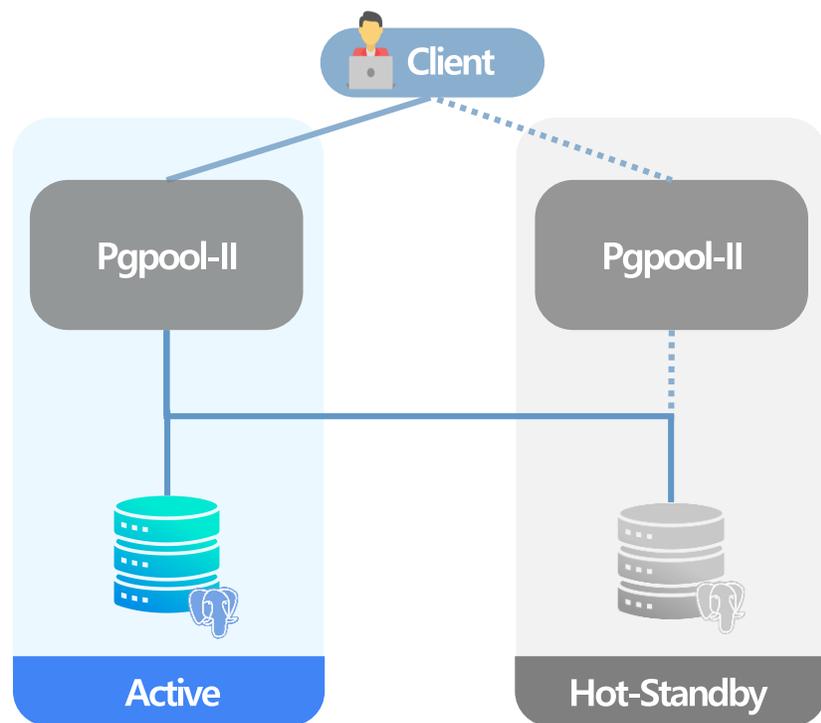


Oracle Compatibility

- ✓ Higher Compatibility Compared to Open Source
- ✓ Extensive Support for Oracle-Specific Functions and Packages → Quarterly Compatibility Improvements

The traditional centralized architecture commonly used for PostgreSQL replication, exhibits limitations in cloud environments

Previous Architecture



Centralized Architecture

Overloaded Single Layer (Pgpool-II) Leading to Inefficiency



Difficulty in Quickly Identifying Root Causes

Rapid Incident Response Could Be Challenging



Reduced High Availability and Split-Brain Occurrences

A Split-Brain Scenario Can Occur When Node-to-Node Communication Fails, Potentially Causing System-Wide Downtime

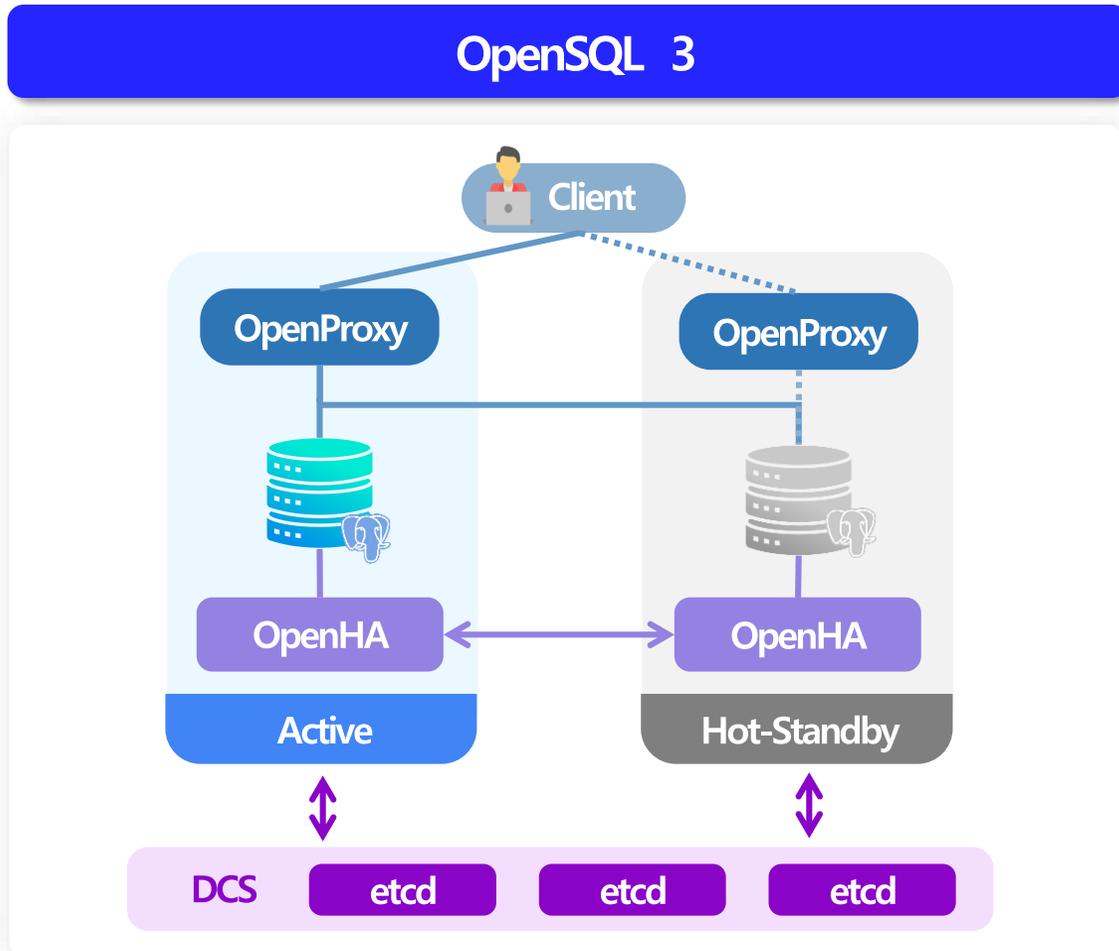


Restart Required During Scale-In/Out Operations

Frequent Scale-In/Out in Cloud Environments Leads to Operational Inefficiency

* Split Brain: A condition in which two or more nodes mistakenly consider each other to have failed and simultaneously assume the primary role.

OpenSQL 3 Adopts a Distributed High-Availability Architecture to Address Cloud Environment Limitations



Minimized Impact of Failures with Distributed Architecture

Distributed High-Availability Architecture Optimized for Highly Dynamic Cloud Environments



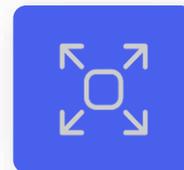
Quick Root-Cause Identification During Failures

Rapid Recovery Enabled by Quick Root-Cause Identification



High-Availability Architecture Without Split-Brain Risk

No Split-Brain Occurrence Even During Network Partitions Through Distributed State Store Consensus Mechanism



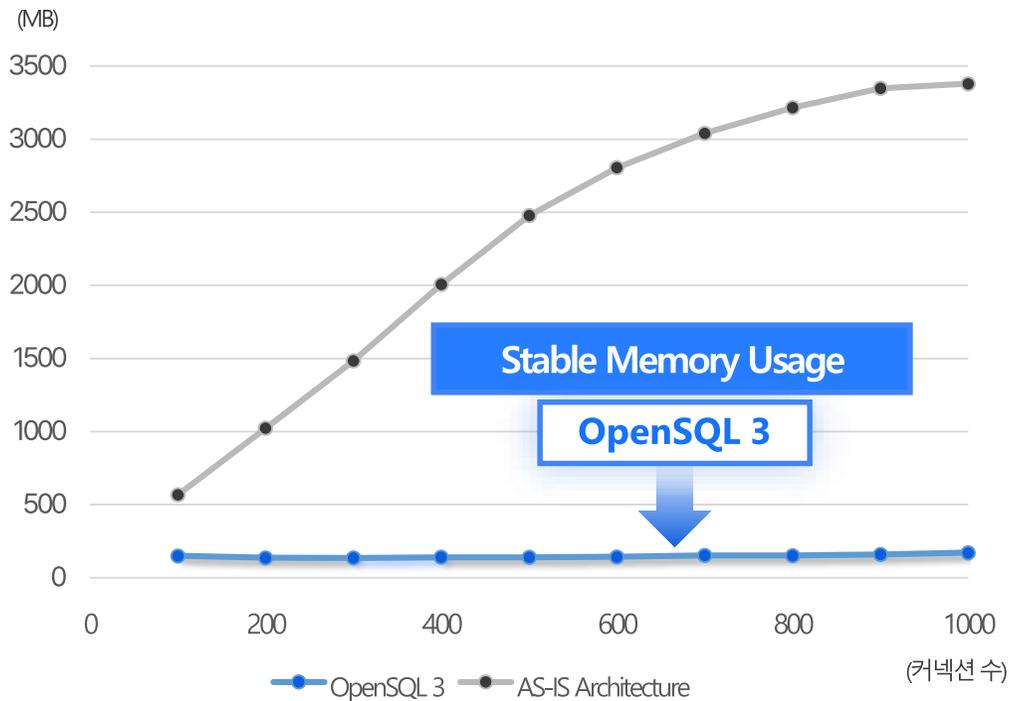
Non-Disruptive Scale-In/Scale-Out Capability

Flexible Scalability in Cloud Environments

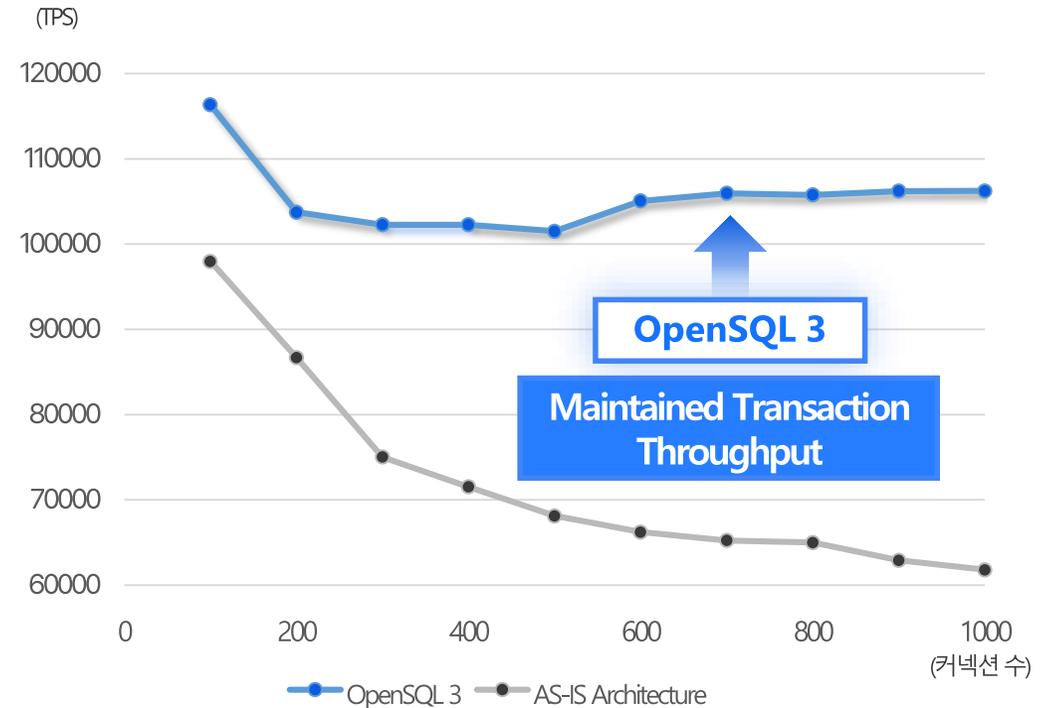
* Split Brain: A condition in which two or more nodes mistakenly consider each other to have failed and simultaneously assume the primary role.

Reduced memory usage and increased TPS leading to enhanced overall performance.

Memory Usage Comparison with Increasing Connections

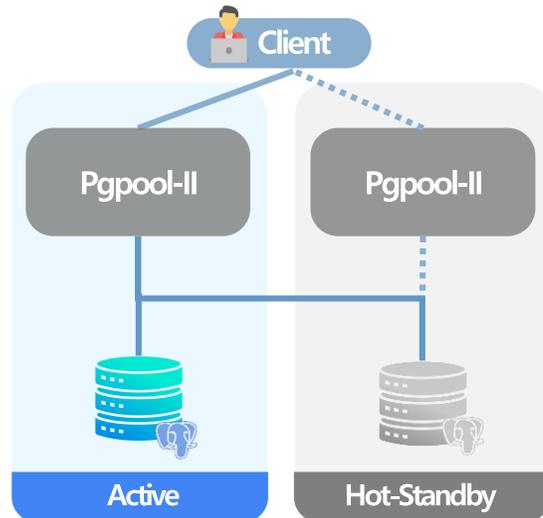


TPS (Transactions Per Second) Comparison with Increasing Connections



Provides Robust High Availability Through Advanced Cluster Management and Enhanced Stability

Previous Architecture

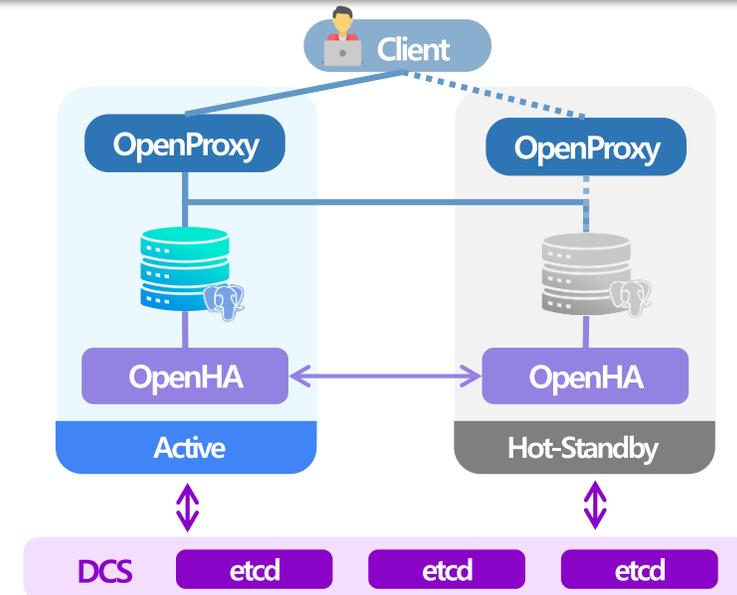


Upon PostgreSQL Failure Detection 🚨

- Script-Based Failover Execution (Separate Scripts Required)
- Cluster Nodes Do Not Share State (Risk of Split-Brain)

VS

OpenSQL 3



Upon PostgreSQL Failure Detection 🚨

- Update state in etcd and synchronize the overall cluster status
 - Elect a new leader in etcd if necessary
 - Reflect cluster state changes to manage the corresponding instances

Chapter 3

Tmax OpenSQL Key References

Proven stable operation across multiple domains and customers

OpenSQL

Win-back for commercial/open-source DBMS

Secured references for high-volume, high-performance DBMS

Cloud-optimized DBMS with elastic scale-in/out support



국가정보자원관리원

G-Cloud: Government Integrated IT Project

- OpenSQL deployed in Daejeon & Daegu for outsourced IT operations of 47 government agencies

Major Implementing Government Agencies



교육부



환경부



통계청



소방청
National Fire Agency



질병관리청



국립중자원



행정안전부



국토교통부



보건복지부



해양수산부



환경부
화학물질안전원



개인정보보호위원회



ELA Contract



e-Learning System
Deployment



Integrated IT System
Deployment



Implementation of Data
Analysis System

...



Thank You

TMAXTibero

별첨

OpenSQL Platform Architecture & Key Features



Product	PostgreSQL			
	Extension			
Function	High Performanace	High Availability	Compatibility	Data Interface
	OpenProxy	OpenHA	O2 Package	FDW
Support	Bakup/Recovery	Security	Management	CDC
	OpenBackup	OpenCrypto	OpenManager	ProSync
Support	Conversion Consulting	Migration	Technical Support	Education

Function	Specific functions	Component	Strength over competitors
DBMS	<ul style="list-style-type: none"> PostgreSQL DBMS 	OpenSQL DB	
Extensions	<ul style="list-style-type: none"> Additional features (GIS, Partitioning, Scheduler, etc.) 	Open-Source	Proven Extension Delivery
High Performance	<ul style="list-style-type: none"> Proxy server that supports load balancing, connection pooling Built with its own Query Parser, automatic load balancing of Read Only queries to Replica DB 	OpenProxy	Only available in OpenSQL High-Performance Connection Pooler
High Availability	<ul style="list-style-type: none"> Cluster Management, Fail-Over, Automatic Fail-Over Perform intra-cluster leader management with DCS, providing management flexibility and flexible scalability 	OpenHA	Only available in OpenSQL Distributed, high-cost architecture
Compatibility	<ul style="list-style-type: none"> Provides strong Oracle compatibility Can be used as is without Oracle query modification, reducing conversion costs 	O2 Package	Provides high Oracle compatibility over open source
Data Interface	<ul style="list-style-type: none"> Used to query or consolidate data in connection with external databases Oracle, Tiberio, Hadoop, MongoDB, MySQL etc 	FDW	Tiberio FDW Only Available
Backup / Recovery	<ul style="list-style-type: none"> Full, incremental, online backup / point in time, remote, parallel recovery Support for a variety of compression methods, enabling backup policies to be set 	OpenBackup	
Security	<ul style="list-style-type: none"> Providing domestic encryption algorithms (ARIA, SEED) Provides access control, audit, profile functionality 	OpenCrypto	Domestic encryption algorithm delivery X
Management	<ul style="list-style-type: none"> Monitoring Tool (SysmasterDB) Editing and Management Tool (pgAdmin) 	OpenManager	
CDC (Change Data Capture)	<ul style="list-style-type: none"> Replicate and synchronize real-time data 	ProSync	X

Function

High Performance

High Availability

Compatibility

Data Interface

Backup /Recovery

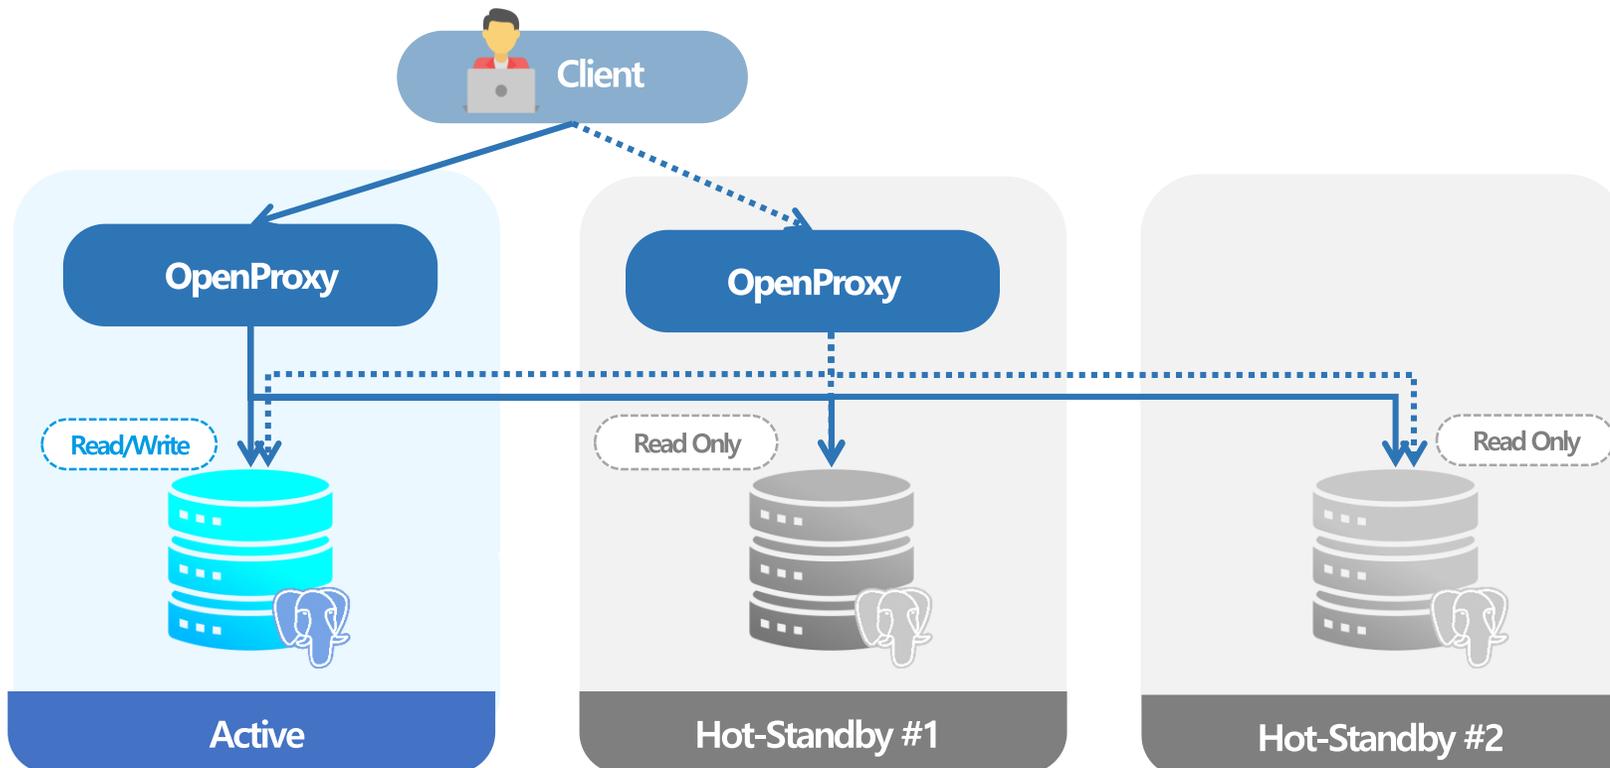
Security

Management

CDC

Improve system-wide response speed and processing efficiency with Connection Pool optimization

OpenSQL Architecture - OpenProxy



Characteristics

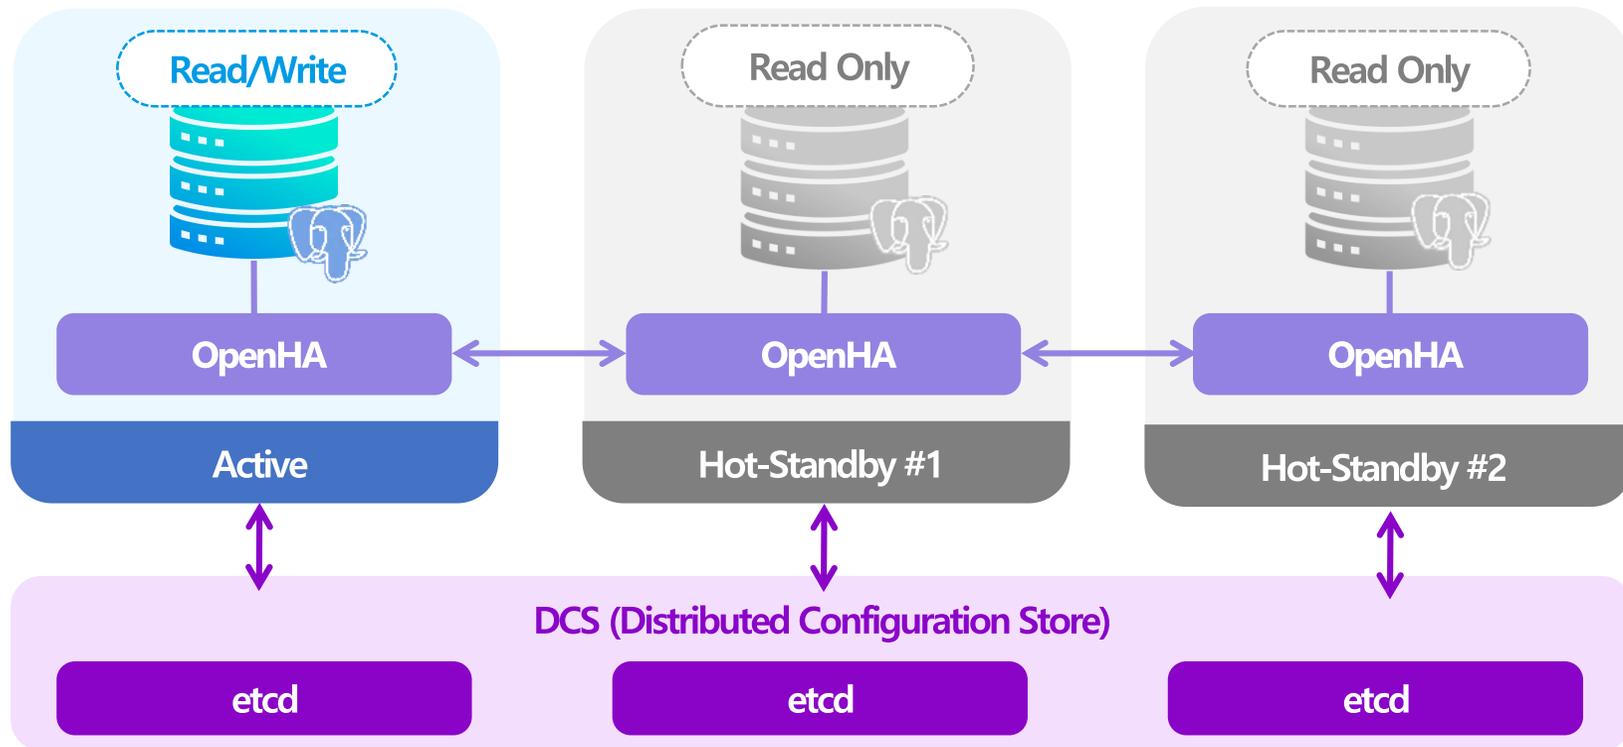
- ✓ **Enhanced large capacity processing capabilities**
Efficient processing with Connection sharing
- ✓ **Use resources efficiently**
Maximum connection processing with minimal resources
- ✓ **Provides convenience in development**
Provides automatic distribution of Read Only queries

Function	High Performance	High Availability	Compatibility	Data Interface	Backup /Recovery	Security	Management	CDC
-----------------	------------------	--------------------------	---------------	----------------	------------------	----------	------------	-----

Advanced cluster management and improved reliability provide strong high availability

OpenSQL Architecture – OpenHA

Characteristics



✓ **Auto Fail-Over**
Automatic fault detection and fast fail-over

✓ **Non-disruptive scale in/out**
Node-independent design enables uninterrupted operation

✓ **Minimize Split-brain***
Quorum*-based architecture enhances reliability

* Quorum : Minimum number of nodes for decision making in the cluster **23**

* Split brain : A state in which two or more nodes mistake each other for failure and simultaneously act as primary

Function

High Performance

High Availability

Compatibility

Data Interface

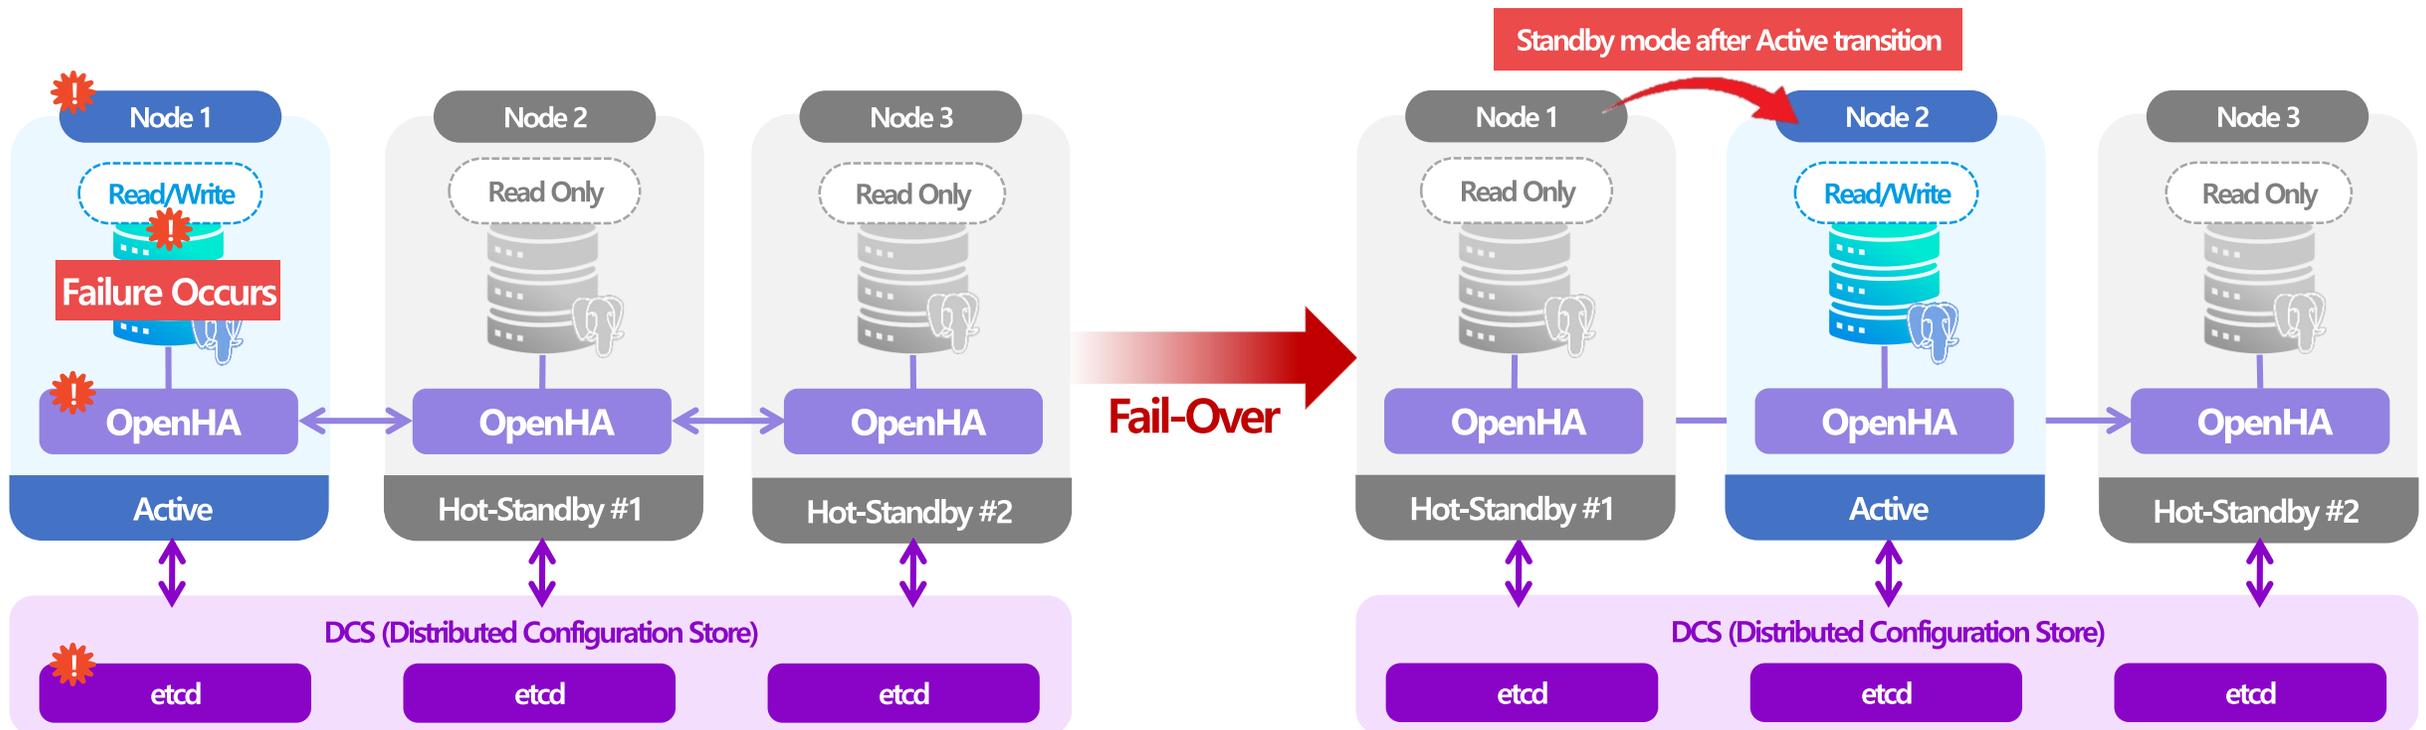
Backup /Recovery

Security

Management

CDC

Advanced cluster management and improved reliability provide strong high availability



Real-time cluster state sharing improves fast fault-tolerant → Cluster Management and High Availability

Function

High Performance

High Availability

Compatibility

Data Interface

Backup /Recovery

Security

Management

CDC

When switching from Oracle DB, use it as it is without query modification

O2 Package (OpenSQL for Oracle Package)



지원 범위	예시
Oracle compatible Data Types	• DATE, VARCHAR2, NVARCHAR2 etc
Oracle compatible View Types	• DUAL, DBA_OBJECTS etc
Oracle compatible Functions	• DECODE, TO_CHAR, NVL, TRUNC, SUBSTR, SYSDATE etc
Oracle compatible Packages	• DBMS_OUTPUT, DBMS_ALERT, DBMS_SQL etc

The effect of introduction

✓ High compatibility over open source

✓ Continuous Compatibility Improvements

✓ Reduce conversion costs

Function

High Performance

High Availability

Compatibility

Data Interface

Backup /Recovery

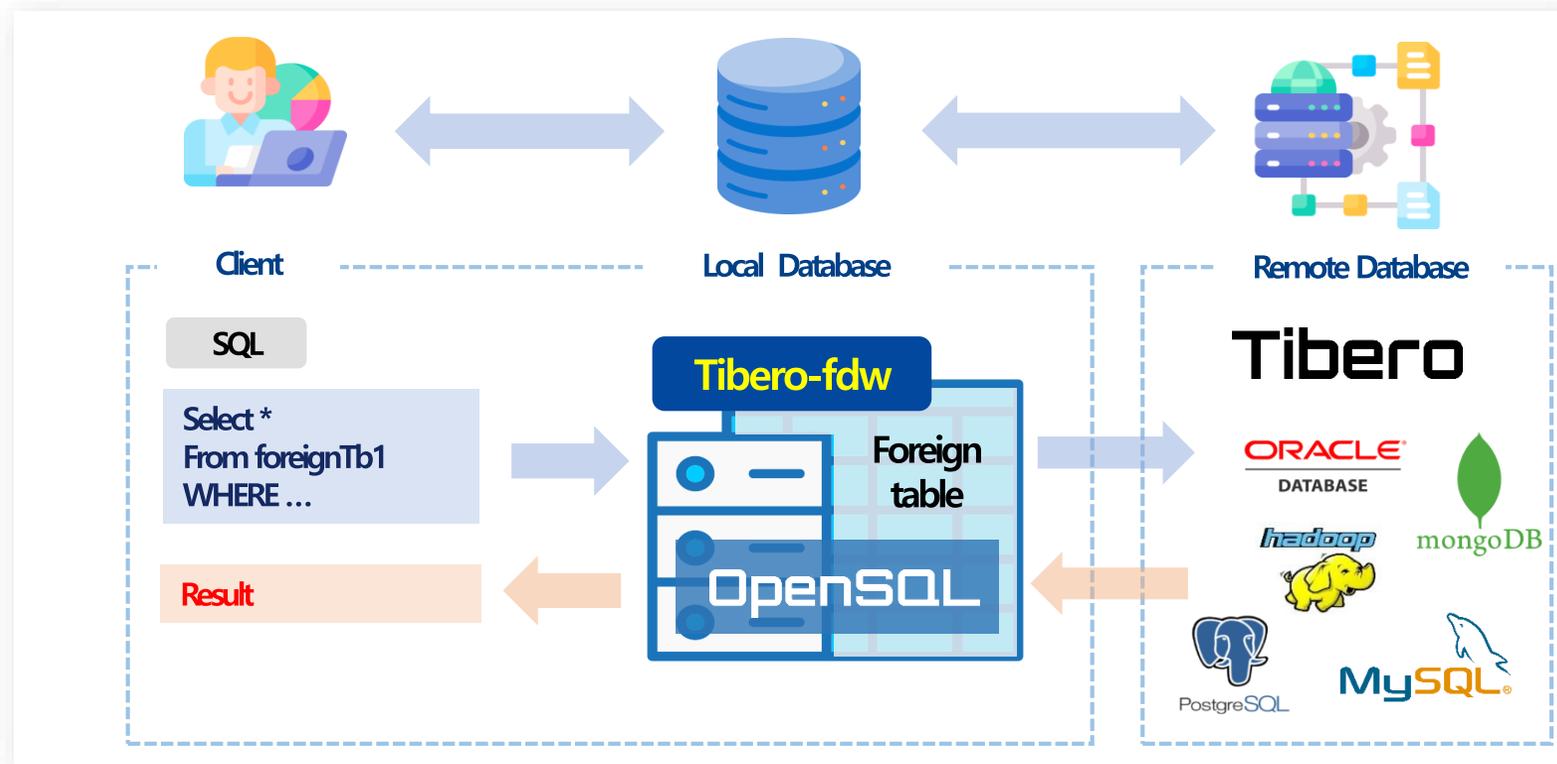
Security

Management

CDC

Used to query or consolidate data by connecting to databases outside OpenSQL

FDW(Foreign Data Wrapper) Architecture

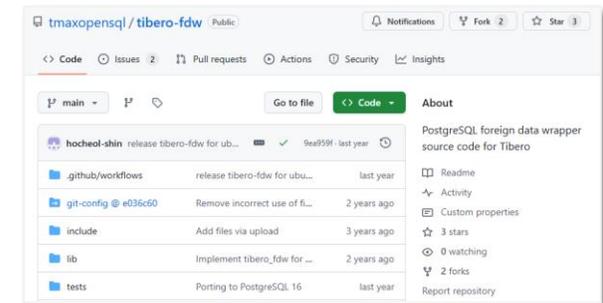


Characteristics

✓ Compatible with diverse database systems

✓ Designed for data migration

✓ Tibero_fdw



<https://github.com/tmaxopensql/tibero-fdw>

Function

High Performance

High Availability

Compatibility

Data Interface

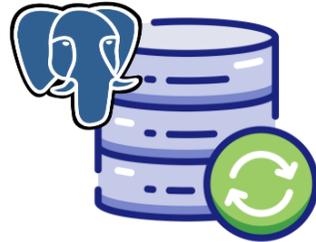
Backup
/Recovery

Security

Management

CDC

Provides reliable backup/recovery with a dedicated backup solution for OpenSQL



Reliable backup and recovery solutions

OpenBackup

Support for a variety of backup and recovery methods

- Backup : Full Backup , Incremental Backup, Online(Hot) Backup..
- Recovery : Point-In-Time Recovery, Parallel Recovery, Remote Recovery..

Works with cloud storage

- AWS S3, Google Cloud Platform, Azure, etc



Backup policy and archive cycle settings

- Effective management of backup retention and deletion policies

Support for a variety of compression methods

- Efficient management of backup space with support for gzip, lz4, ztsd, etc

Function

High Performance

High Availability

Compatibility

Data Interface

Backup
/Recovery

Security

Management

CDC

Provides local encryption extension module for PostgreSQL



optimized for the domestic information security environment

OpenCrypto

Provides a variety of encryption algorithms

- Support for domestic NIS encryption algorithms (ARIA, SEED)
- Provides global encryption algorithms (AES, DES, etc.)

Strengthen user authentication and access control

- Configuring granular certification policies tailored to the enterprise environment

Fast-response framework for security vulnerabilities

- Provides quick patches and updates for the latest security issues

Provide security guides and regular inspection services

- Provides practical, operations-focused guidance while minimizing administrator workload

Function

High Performance

High Availability

Compatibility

Data Interface

Backup /Recovery

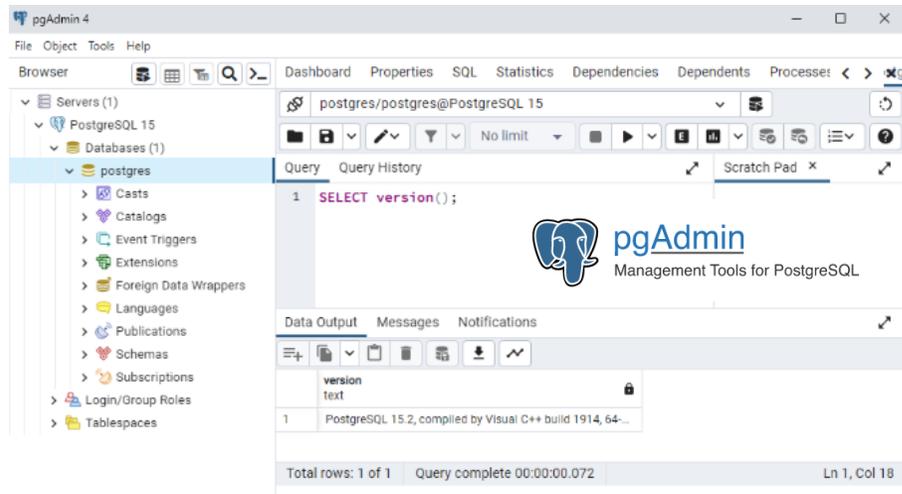
Security

Management

CDC

Provides tools for easy and convenient DB operation management

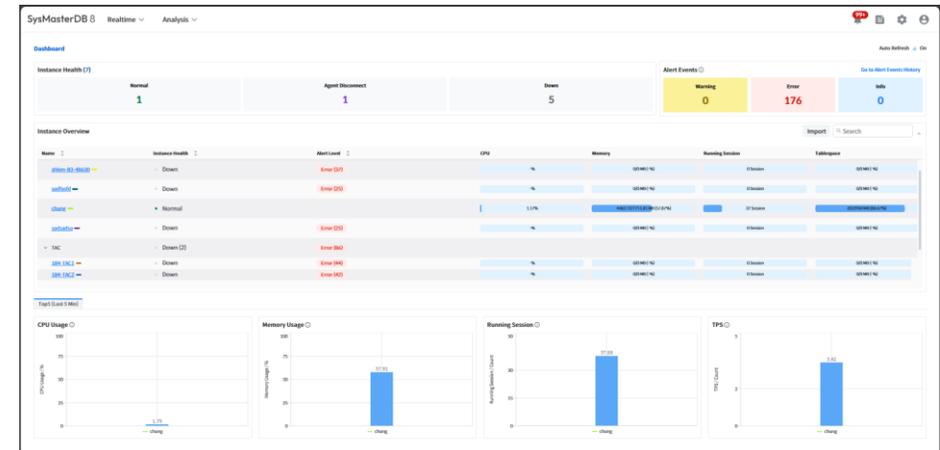
Development/Operation



pgAdmin

- Official Management Tools for PostgreSQL
- DBMS Monitoring and Administrator Features
- Provides SQL developer capabilities

Monitoring



SysmasterDB

- Provides real-time/integrated monitoring
- DBMS, OS area indicators can be monitored
- User-friendly GUI for ease of management

Function

High Performance

High Availability

Compatibility

Data Interface

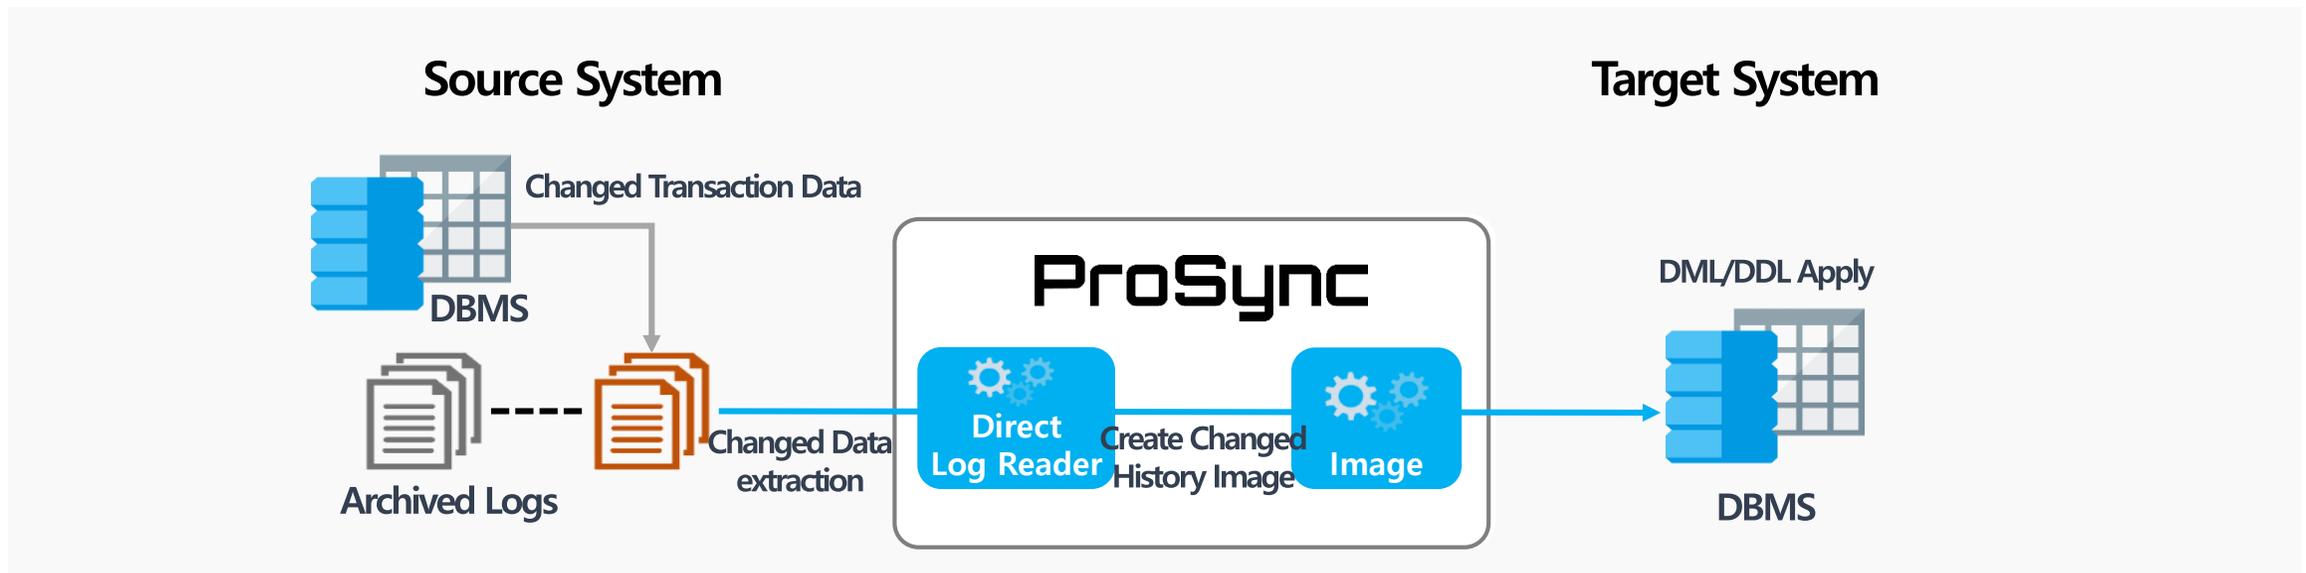
Backup /Recovery

Security

Management

CDC

CDC solutions that provide real-time data replication and synchronization



Real-time extraction of changed data

- **OpenSQL**, Tiberio, Oracle
- DB log file variation detection
- Only detect and send changes

Processing Extracted Data

- Extracted Data Simple Variations, Aggregation
- Changed data queuing, ensuring order

Real-time reflection of target systems

- **OpenSQL**, Tiberio, Oracle
- Ensures consistency of applied data

Standard architecture designed for high availability

