

Migrate Zabbix DB from Oracle to PostgreSQL with TimescaleDB



Thomas Oftring

Team Lead Monitoring
Zabbix Trainer/Database Architect
Rhein-Main Solutions GmbH




Why should you migrate?



Oracle support
is deprecated

Zabbix 7.0 LTS

- ▶ Last release
- ▶ 5 years to migrate



Decrease the costs

PostgreSQL

- ▶ Free license
- ▶ Support available if needed



Partitioning and
Compression

TimescaleDb

- ▶ Partitioning with Hypertables
- ▶ Compression

Available Migration Tools


- ▶ Pgora
<https://github.com/pgoracle/pgora-migrate>
11 years old no activity
- ▶ ora2pg
<https://ora2pg.darold.net>
Poor read performance on Oracle
- ▶ CYBERTEC Oracle Migrator
<https://www.cybertec-postgresql.com/en/products/cybertec-postgresql-migrator/>
No OpenSource
- ▶ More tools outside, not tested by our customers or ourself



PostgreSQL and TimescaleDB Toolset used for migration

- ▶ Foreign Data Wrapper for Oracle
https://github.com/laurenz/oracle_fdw
 - Stable, reliable and fast
 - Many configuration options
 - All PostgreSQL functions are available for data conversion and validation
- ▶ timescaledb-parallel-copy
<https://github.com/timescale/timescaledb-parallel-copy>
 - Really fast on importing CSV data to hypertables
could work parallel with one CSV file
 - Low memory usage
 - Parameter batch-error-output-dir
stores data that could not be imported
great to find messed up data and correct them/convert them





Migration Steps

Step by Step Database Migration

- ▶ Assessment and planning
- ▶ Database Logic and Schema migration
- ▶ Data migration
- ▶ Data transformation
- ▶ Testing and Validation
- ▶ Deployment and checks



Security and Access control

- ▶ Different models for security
- ▶ Custom policies



No big deal!

- ▶ Create new Zabbix PostgreSQL database
- ▶ Create the needed users and grant access rights
 - Zabbix Server account (should be DB owner)
 - Zabbix Web Frontend account
 - Zabbix Monitoring user account
 - Check for additional accounts




Partitioning challenges

► Oracle

- Oracle supports partitioning
Zabbix on Oracle does not!
- Check if there is manual partitioning by scripts
- Make sure to get all data including the partitioned one

► PostgreSQL

- Official Zabbix Support for TimescaleDB
- Transparent partitioning
- Transparent compression



Database logic migration

Stored Procedures and Triggers in Zabbix 7.0.12

- ▶ Oracle
 - 0 Stored Procedures
 - 0 Functions
 - 66 Trigger
- ▶ PostgreSQL
 - 0 Stored Procedures
 - 64 Functions
 - 65 Trigger



Nothing ToDo!

Zabbix Server and PostgreSQL database schema contains already the logic!



Schema migration

Table, indexes, constraints in Zabbix 7.0.12

- ▶ Oracle
 - 203 tables
 - 312 indexes
 - 272 constraints
- ▶ PostgreSQL
 - 203 tables
 - 312 indexes
 - 272 constraints



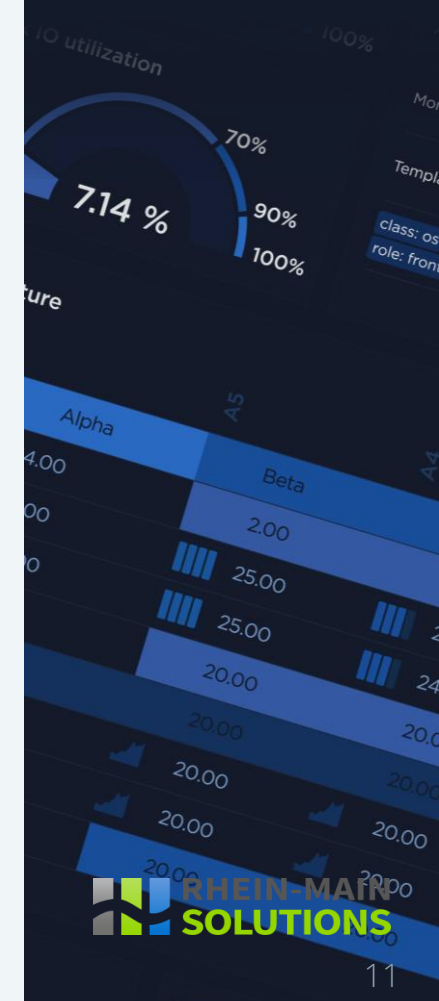
Nothing ToDo!

Zabbix Server and PostgreSQL database schema contains already the objects!



Pitfalls – Part I

- ▶ Column order will be different between Oracle and PostgreSQL
 - Especially if Oracle database has seen many Zabbix DB upgrades
- ▶ Datatypes mapping
- ▶ Differences in Date and Time
 - Oracle DATE vs PostgreSQL TIMESTAMP
 - Oracle TIMESTAMP WITH LOCAL TIME ZONE vs PostgreSQL TIMESTAMPTZ
- ▶ Encoding issues
 - nclob
 - nvarchar2



Pitfalls - Part II

▶ Different DBNULL handling

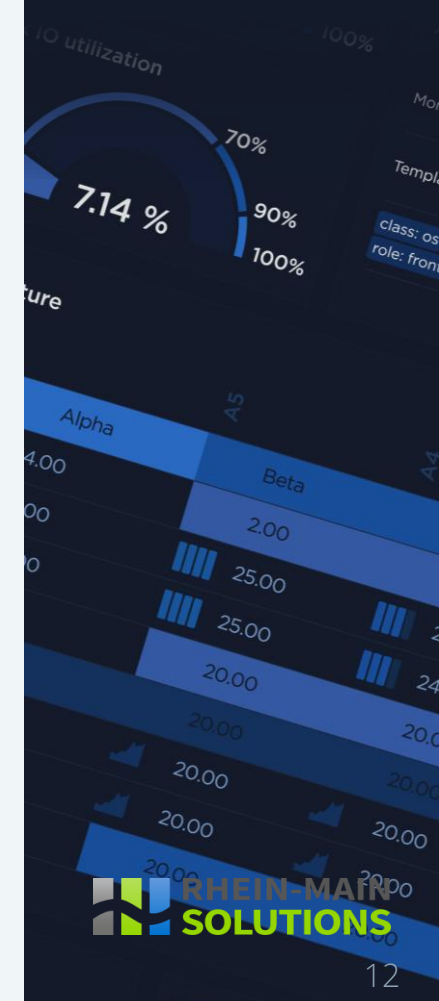
- Oracle treats empty strings as NULL
- PostgreSQL treats empty strings as NOT NULL
- Oracle concatenates NULL values with strings to strings
- PostgreSQL concatenates NULL values with strings to NULL

▶ Oracle SEQUENCE to PostgreSQL BIGSERIAL

- Only one table in Zabbix Database
changelog
- Just Insert the data to the table
- Reset the BIGSERIAL

```
SELECT max(changeloid) + 1 from changelog;
```

```
ALTER SEQUENCE changelog_changeloid_seq RESTART WITH xxxxxxxx;
```



Data migration – Step 1

Prepare Oracle Zabbix database

- ▶ Create Views
 - e.g extract data for nclob with DBMS_LOB.substr
 - Much faster on Oracle site than converting it in PostgreSQL view
- ▶ Check configured timeouts on Oracle
 - SQL command timeout
- ▶ Check Oracle log for errors
 - to know which problems already exist
 - Fix them
- ▶ Check Zabbix server logs
 - Database errors/warning
 - Slow queries (could get a problem during migration)



Data migration – Step 2

Prepare PostgreSQL

- ▶ Create Zabbix user and database
 - Create only the tables from Zabbix schema (CREATE TABLE statements)
 - No indexes and constraints!
- ▶ Install and configure PostgreSQL Foreign Data Wrapper
 - Create extension oracle_fdw
 - Create foreign server and user mapping
 - Create foreign tables for all tables (use an own schema e.g. or_import)
 - Use the Oracle views instead of tables if your created one
- ▶ Create views for data conversion and cleansing if necessary
 - DataType Mapping
 - Column order according to PostgreSQL Zabbix database schema
 - Use PostgreSQL logic for automating the view creation
 - Check every view and adjust if necessary



Data migration – Step 2

Prepare PostgreSQL

- ▶ Optimize PostgreSQL configuration for data loading
- ▶ Adjust to your DB Server configuration
- ▶ Good starting point
<https://www.pgconfig.org>
 - Application Profile:
DataWare house and BI Application
 - max_wal_size should be high for bulk load
Reduces the count of checkpoints
 - checkpoint_timeout should be increased
Reduces the amount of full-page image writes



Data migration – Step 3

TimescaleDB

- ▶ Install and configure TimescaleDB
- ▶ Enable timescaledb extension in the Zabbix database
- ▶ Run Zabbix TimescaleDB script against the Zabbix database
 - Checks TimescaleDB configuration
 - Creates hypertables
 - Enables compression



Data migration – Step 4

Transfer data – regular tables

- ▶ 203 tables in Zabbix Version 7.0.12
- ▶ 193 tables – very straight forward
- ▶ Create views on foreign tables if needed
- ▶ Export the data to CSV files and import it
Use table or view as source, check column order!

```
COPY (SELECT * FROM or_import.hosts ) TO '/tmpstorage/hosts.csv';
COPY public.hosts FROM '/tmpstorage/hosts.csv';
```

- ▶ Slower but also possible
Use table or view as source, check column order!

```
INSERT INTO public.tablename (column1, ..., columnN)
SELECT column1, ..., columnN FROM or import.tablename;
```

Data migration – Step 4

Transfer data – binary tables

- ▶ Oracle uses blob and PostgreSQL bytea
 - oracle_fdw supports casting blob column to bytea
 - create view with the according cast if Oracle Version older than 19
 - set lob_prefetch according to typical screenshot size for better performance

- ▶ Transfer the data

```
INSERT INTO public.images (imageid, imagetype, name, image)
SELECT imageid, imagetype, name, image::bytea FROM or_import.images;
```

- ▶ Infos according to binary tables
 - images (background images and icons)
normally fast
 - history_bin
definitely the table with the slowest migration speed



Data migration – Step 4

Transfer data – TimescaleDB tables

▶ How to handle the TimescaleDB Zabbix tables:

- auditlog
- history*
- trends*

▶ Create views on foreign tables if needed

▶ Export the data to CSV files
(Use table or view as source)

```
COPY (SELECT itemid, clock, value, ns FROM or_import.history)
TO '/tmpstorage/history.csv';
```

▶ It is possible to export the data partitioned by WHERE clause

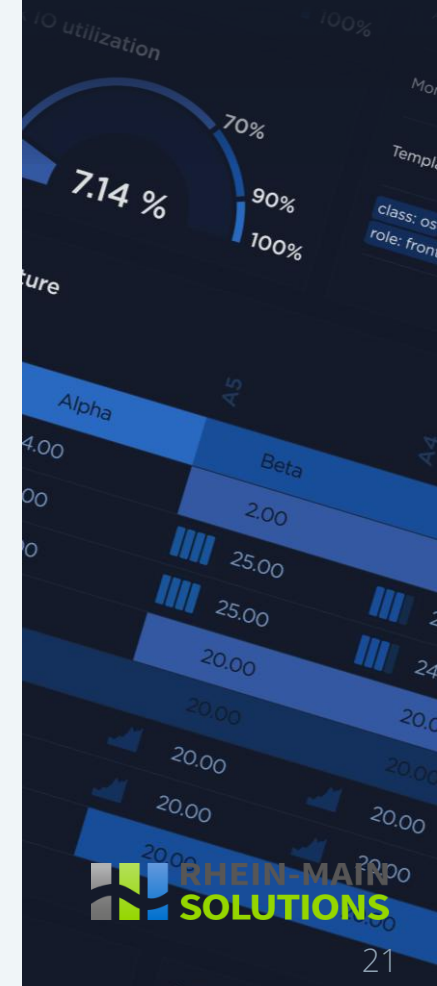
- Use interval that fits to your data days, weeks, months
- Can be done parallel to the production workload, no need for downtime
- Possibility to calculate time for migration by duration of testing interval



Data migration – Step 5


Indexes - Triggers - Constraints

- ▶ Create all unique indexes from Zabbix database schema file
 - May take some time
 - May cause troubles if data are not unique
 - Find source of the problem and fix it
- ▶ Create all non unique indexes from Zabbix database schema file
 - May take some time
- ▶ Create all triggers from Zabbix database schema file
 - Should be created without troubles
- ▶ Create all constraints from Zabbix database schema file
 - The point of truth, data consistence is checked!
 - Fix all problems that occurs
 - May take some time



Testing and validation

- ▶ Use DEV/QA systems
- ▶ Have all tables data?
- ▶ Verify the row count on the tables!
- ▶ Pay attention to NULL handling



Be careful and validate thoroughly!

- ▶ Check all errors and warning during data migration!
- ▶ Test the Zabbix Web Frontend against the migrated database
- ▶ Check Network Maps



Deployment, check and monitoring

- ▶ Stop old Zabbix Server and Web Frontend
- ▶ Replace Zabbix Oracle with PostgreSQL packages
- ▶ Migrate the latest data
- ▶ Start the Web Frontend and check
- ▶ Start the Zabbix Server and check

Checking and monitoring!!!

- ▶ Check log files for errors and warnings
 - PostgreSQL
 - Zabbix Server
- ▶ Monitor PostgreSQL with Zabbix



Summary/Lessons learned

- ▶ Migration Oracle to PostgreSQL with TimescaleDB is possible
- ▶ Needs preparation
- ▶ Needs time for testing
- ▶ Needs a Testing/QA Platform
- ▶ Possible with free toolchain
- ▶ It is possible to minimize the migration duration



ZABBIX '25

CONFERENCE

GERMANY



Thomas Oftring

Team Lead Monitoring
Zabbix Trainer/Database
Architect
Rhein-Main Solutions GmbH

