

Improve PostgreSQL replication with Oracle GoldenGate 19c

2021

About RheoData


A global systems integrator headquartered in Atlanta, Georgia. Providing solutions across the enterprise with architecture design/assessments, implementation, data replication, and professional services to ensure high-quality enterprise systems leveraging Oracle and Microsoft technologies. RheoData works hand-in-hand with our strategic partners and their sales teams to build and deploy high performance and scalable enterprise solutions in addition to providing clients with cost savings professional services for turnkey enterprise management.



Adoption over time

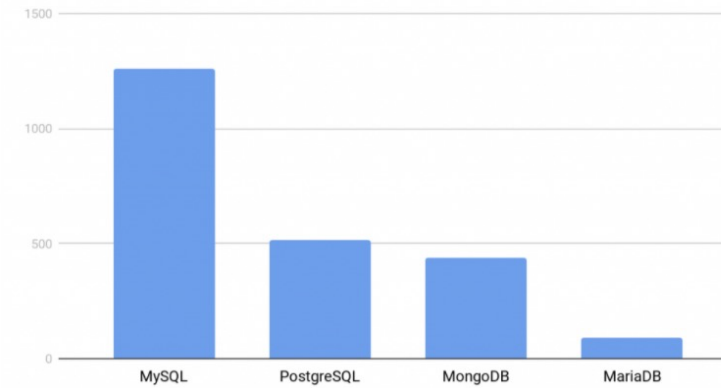
Open Source vs. Commercial

Popularity trend



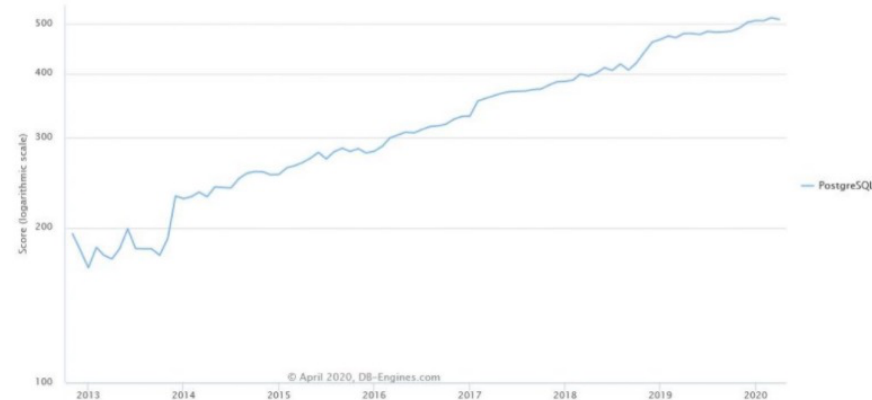
Ranked #2

March 2020 DB-Engines Rankings

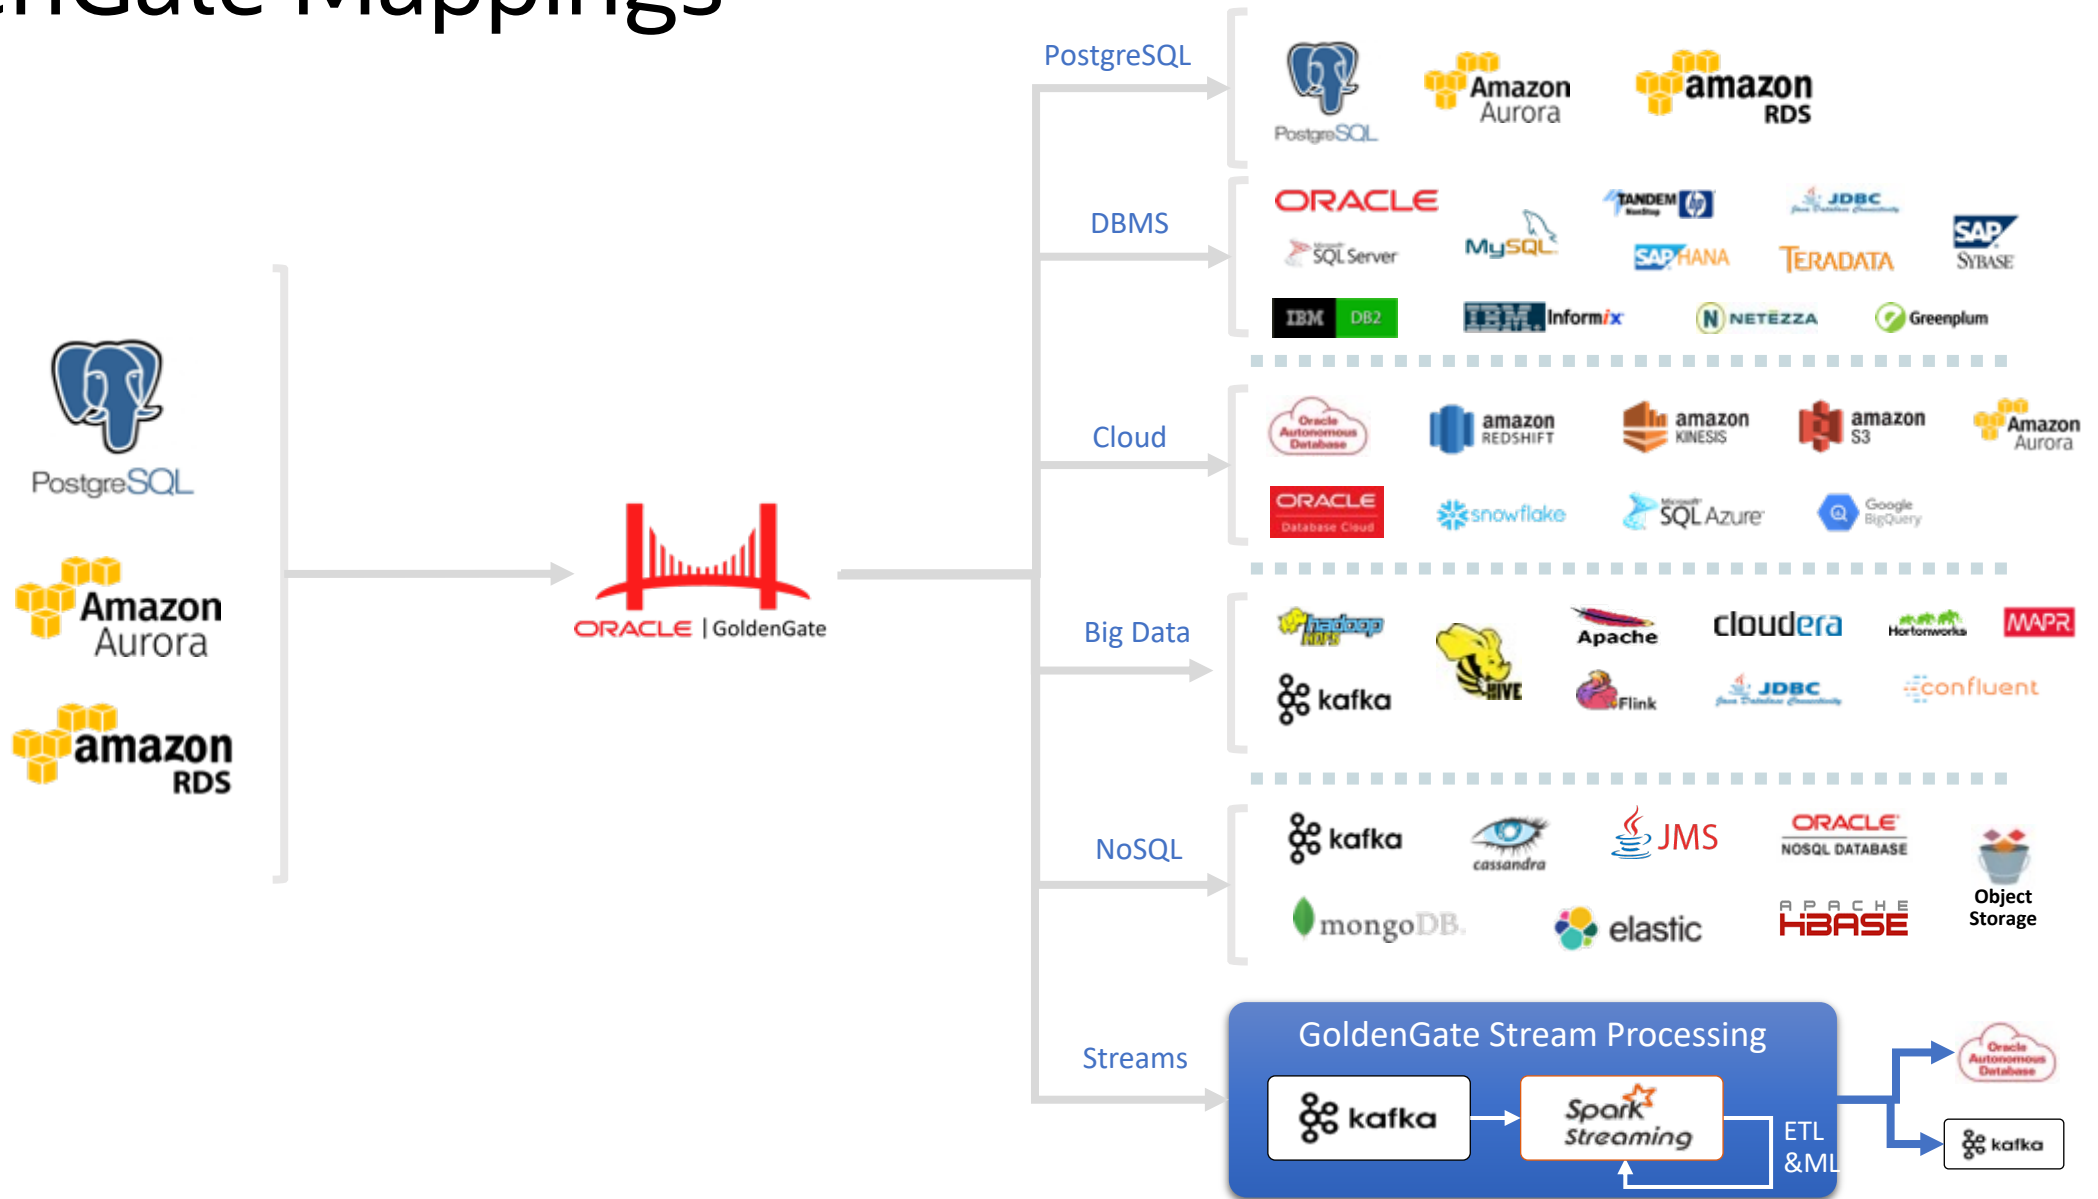


Steady Growth

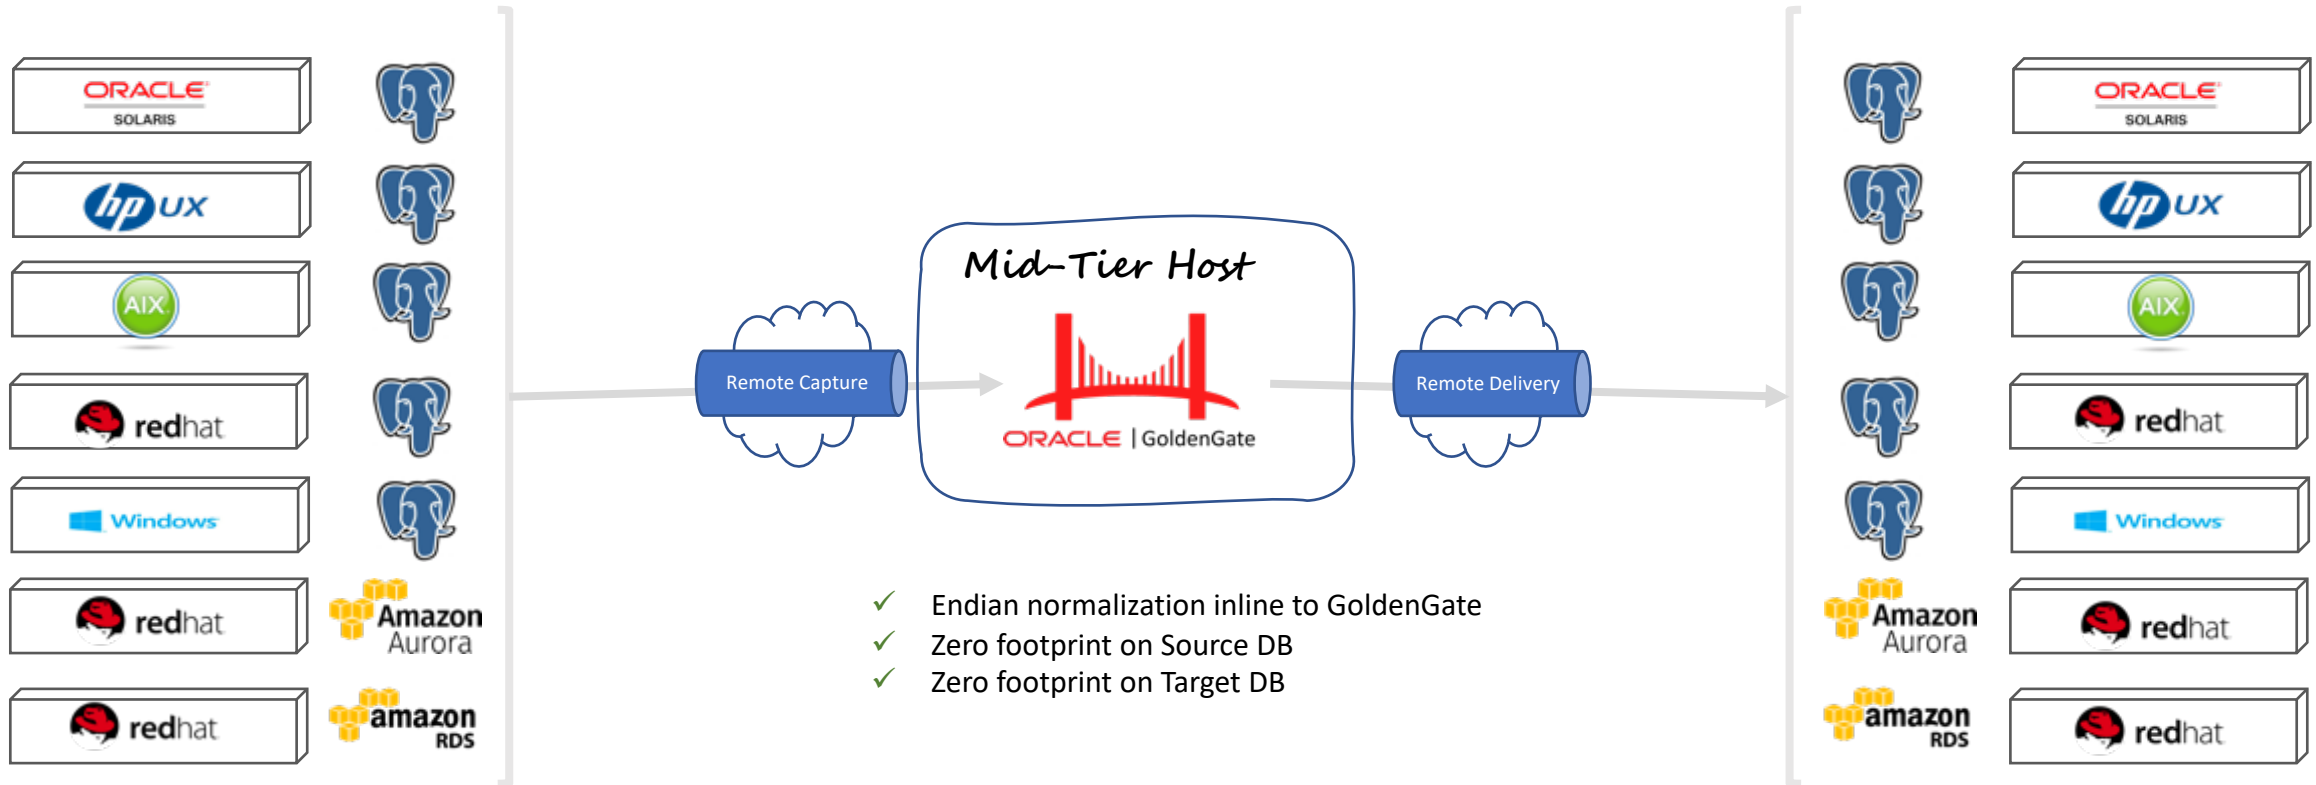
DB-Engines Ranking of PostgreSQL



GoldenGate Mappings



Middle-Tier Support



Benefits of GoldenGate PostgreSQL Capture

- ✓ Initial Load – to/from PostgreSQL Database
 - ✓ PostgreSQL v10 & above - Capture support
 - ✓ Native build on Linux platform
 - ✓ Majority of data types, plus CLOB/BLOB supported
 - ✓ SSL enabled database supported – capture & delivery
 - ✓ Globalization Support
 - ✓ Multiple replication slots to increase capture throughput
 - ✓ Co-Exist with native replication tool
 - ✓ Heartbeat Table support
 - ✓ Coordinated replicat & BatchSQL support
 - ✓ Conflict, Detection & Resolution (CDR) Support



PostgreSQL



GoldenGate for PostgreSQL

Prerequisites

- Required RPMs
- `PostgreSQL.conf` file
- Replication Users, i.e. GoldenGate User (`ggate`)
- Environment Variables
- Open Database Connector file (`ODBC.ini`)



PostgreSQL



GoldenGate for PostgreSQL

Prerequisites - RPMS

- Ensure all the required RPMs are installed
 - `PostgreSQL##-contrib`
 - `oracle-database-preinstall-19c`



PostgreSQL



GoldenGate for PostgreSQL

Prerequisites – PostgreSQL.conf

- Location will vary depending on PostgreSQL version
- Updates required:

| Parameter | Value | Description |
|------------------------------|---------|-------------|
| wal_level | logical | Default |
| max_replication_slots | 10 | Default |
| max_wal_sender | 10 | Default |
| wal_receiver_status_interval | | Optional |
| wal_sender_timeout | | Optional |
| track_commit_timestamp | False | Optional |



PostgreSQL



GoldenGate for PostgreSQL


Prerequisites – GoldenGate User

```
psql# create user {user} with password '{passwd}' login;
```

```
psql# alter user {user} with replication;
```



PostgreSQL



ORACLE | GoldenGate

GoldenGate for PostgreSQL

Prerequisites – Environment Variables

| Environment Variable | Value (change as needed) |
|----------------------|--|
| PG_HOME | /usr/pgsql-12 |
| OGG_HOME | /opt/app/oracle/19.1.0/oggcore_1 |
| LD_LIBRARY_PATH | \$PG_HOME/lib:\$OGG_HOME/lib:\$LD_LIBRARY_PATH |
| ODBCINI | {path to odbc.ini}/odbc.ini |
| PATH | \$ODBCINIT:\$PATH |



PostgreSQL



GoldenGate for PostgreSQL

Prerequisites – ODBC.ini

- Location of ODBC.ini file can be anywhere, just needing to be referenced by environment variables
- Sets up default connection items for PostgreSQL

```
[ODBC Data Sources]
postgre=DataDirect 7.1 PostgreSQL Wire Protocol
#pgdsn=DataDirect 7.1 PostgreSQL Wire Protocol

[ODBC]
IANAAppCodePage=4
InstallDir=/opt/app/oracle/19.1.0/oggcore_1

[pgdsn]
Driver=/opt/app/oracle/19.1.0/oggcore_1/lib/GGpsql25.so
Description=DataDirect 7.1 PostgreSQL Wire Protocol
Database=postgres
HostName=10.0.2.15
#HostName=localhost
PortNumber=5432
LogonID=postgres
Password=Welcome1
```



PostgreSQL



GoldenGate for PostgreSQL

GoldenGate Installation



- Same as any other Oracle GoldenGate Classic installation
 - `unzip/tar -xvf` into `$OGG_HOME`
 - `$OGG_HOME/GGSCI`
 - `CREATE SUBDIRS`

```
drwxr-x---. 2 oracle oinstall      6 Aug 26 00:51 dirwlt
drwxr-x---. 2 oracle oinstall      6 Aug 26 00:51 dirdef
drwxr-x---. 2 oracle oinstall      6 Aug 26 00:51 dircrd
drwxr-x---. 2 oracle oinstall     31 Aug 26 01:39 dirdmp
drwxr-x---. 2 oracle oinstall     29 Aug 26 02:43 dirsql
drwxr-x---. 2 oracle oinstall     63 Aug 26 10:48 dirdat
drwxr-x---. 2 oracle oinstall     59 Aug 26 10:53 dirprm
drwxr-x---. 2 oracle oinstall     44 Aug 26 11:00 dirchk
drwxr-x---. 2 oracle oinstall    4096 Aug 26 11:00 dirrpt
drwxr-x---. 2 oracle oinstall     59 Aug 26 11:00 dirpcs
drwxr-x---. 2 oracle oinstall     86 Aug 26 11:00 dirtmp
```

GoldenGate for PostgreSQL Manager



- Configure and start Manager (MGR)
 - `$OGG_HOME/GGSCI`
 - `EDIT PARAMS MGR`
 - Port 16000 (default 7809)

```
drwxr-x---. 2 oracle oinstall      59 Aug 26 10:53 dirprm
|
-rw-r-----. 1 oracle oinstall    11 Aug 26 00:52 mgr.prm
```

GoldenGate for PostgreSQL

Register Extract



```
GGSCI> dblogin sourcedb {dsn} userid {userid} password {password}
GGSCI> register extract {extract_name}
```



Make sure that the RPM for *PostgreSQL##-contrib* is installed (Linux platforms)

If RPM is not installed, then registering will not happen:

2020-08-26 02:57:43 WARNING OGG-00552 Database operation failed: Executing statement to create replication slot. ODBC error: SQLSTATE S1000 native database error 8803729. [Oracle][ODBC PostgreSQL Wire Protocol driver][PostgreSQL]ERROR: VERROR; could not access file "**test_decoding**": No such file or directory(File dfmgr.c; Line 210; Routine internal_load_library;.

GoldenGate for PostgreSQL

Add Trandata



```
GGSCI> dblogin sourcedb {dsn} userid {userid} password {password}  
GGSCI> add trandata {schema}.{table} [ALLCOLS] | [KEYCOLSONLY]
```


GoldenGate for PostgreSQL

Add Extract



```
GGSCI> add extract {extract_name}  
GGSCI> add exttrail ./dirdat/aa, extract {extract_name}  
GGSCI> edit params {extract_name}
```

- Extract Data Pumps can be created *IF* trail file shipping is required

GoldenGate for PostgreSQL

Add Replicat



```
psql# create schema {schema}

GGSCI> dblogin sourcedb {dsn} userid {userid} password {passwd}
GGSCI> add checkpointtable {schema}.{checkpointtable_name}
GGSCI> add replicat {replicat_name}, exttrail ./dirdat/aa, checkpointtable {schema}.{checkpointtable_name}
GGSCI> edit params {extract_name}
```

PostgreSQL Replication Slots



PostgreSQL

```
psql# select * from pg_catalog.pg_replication_slots;
```

| slot_name | plugin | slot_type | datoid | database | temporary | active | active_pid | xmin | catalog_xmin | restart_lsn | confirmed_flush_lsn |
|--------------------------|---------------|-----------|--------|----------|--------------------------|-------------------------------------|------------|--------|--------------|-------------|---------------------|
| postext_91ab904197741863 | test_decoding | logical | 14,188 | postgres | <input type="checkbox"/> | <input checked="" type="checkbox"/> | 13,187 | [NULL] | 3313 | 0/186FF00 | 0/18710F0 |
| salext_927300d810b1ff66 | test_decoding | logical | 16,458 | sales | <input type="checkbox"/> | <input checked="" type="checkbox"/> | 16,206 | [NULL] | 5829 | 0/19A6F38 | 0/19A6F38 |

- Provides information on the physical slots

PostgreSQL Replication Stats




PostgreSQL

```
psql# select * from pg_catalog.pg_stat_replication;
```

| | pid | usesysid | username | application_name | client_addr | client_hostname | client_port | backend_start | backend_xmin | state | sent_lsn | write_lsn |
|---|--------|----------|----------|-------------------|-------------|-----------------|-------------|---------------------|--------------|-----------|-----------|-----------|
| 1 | 17,195 | 16,412 | ggate | GoldenGateCapture | 10.0.2.15 | [NULL] | 65,328 | 2020-08-27 00:50:30 | [NULL] | streaming | 0/1D6F738 | 0/1D6F618 |
| 2 | 16,206 | 16,412 | ggate | GoldenGateCapture | 10.0.2.15 | [NULL] | 65,320 | 2020-08-27 00:34:34 | [NULL] | streaming | 0/1D6F738 | 0/19A6F38 |

- Information on replay lag
- Information on outgoing/sent LSN

On-Premises to Cloud Remote Apply




Key Item: Connection needs to be $\geq 100\text{ms}$

Note: Used AWS Free Tier during testing

On-Premises to Cloud

Configure Remote Apply




AWS Items:

- Open up ACL for Security Groups
 - Make sure that your host can access in/out of the AWS network
 - Gather connection details
- Endpoint Address - DNS style
 - Port (5432)

Note: Used AWS Free Tier during testing

On-Premises to Cloud Configure Remote Apply



Host Items:

- Create/Update ODBC DSN
 - Make sure that your host can access in/out of the AWS network
- Endpoint Address - DNS style
 - Use same parameters to make connection to a database


SOURCEDB DSNNAME USERID, userid, PASSWORD password

```
[aws]
Driver=/opt/app/oracle/19.1.0/oggcore_1/lib/GGpsql25.
so
Description=DataDirect 7.1 PostgreSQL Wire Protocol
Database=ggtest
HostName=ggtesting.cc36pqhw59w8.us-east-
2.rds.amazonaws.com
PortNumber=5432
LogonID=postgres
Password=<password>
```

Note: Used AWS Free Tier during testing

On-Premises to Cloud

Configure PostgreSQL (RDS)



RDS Setup:

- Create Tablespaces

```
Create tablespace <ts_name> owner <schema> location '<path>';
```

- Create Schemas

```
Create schema <schema>;
```

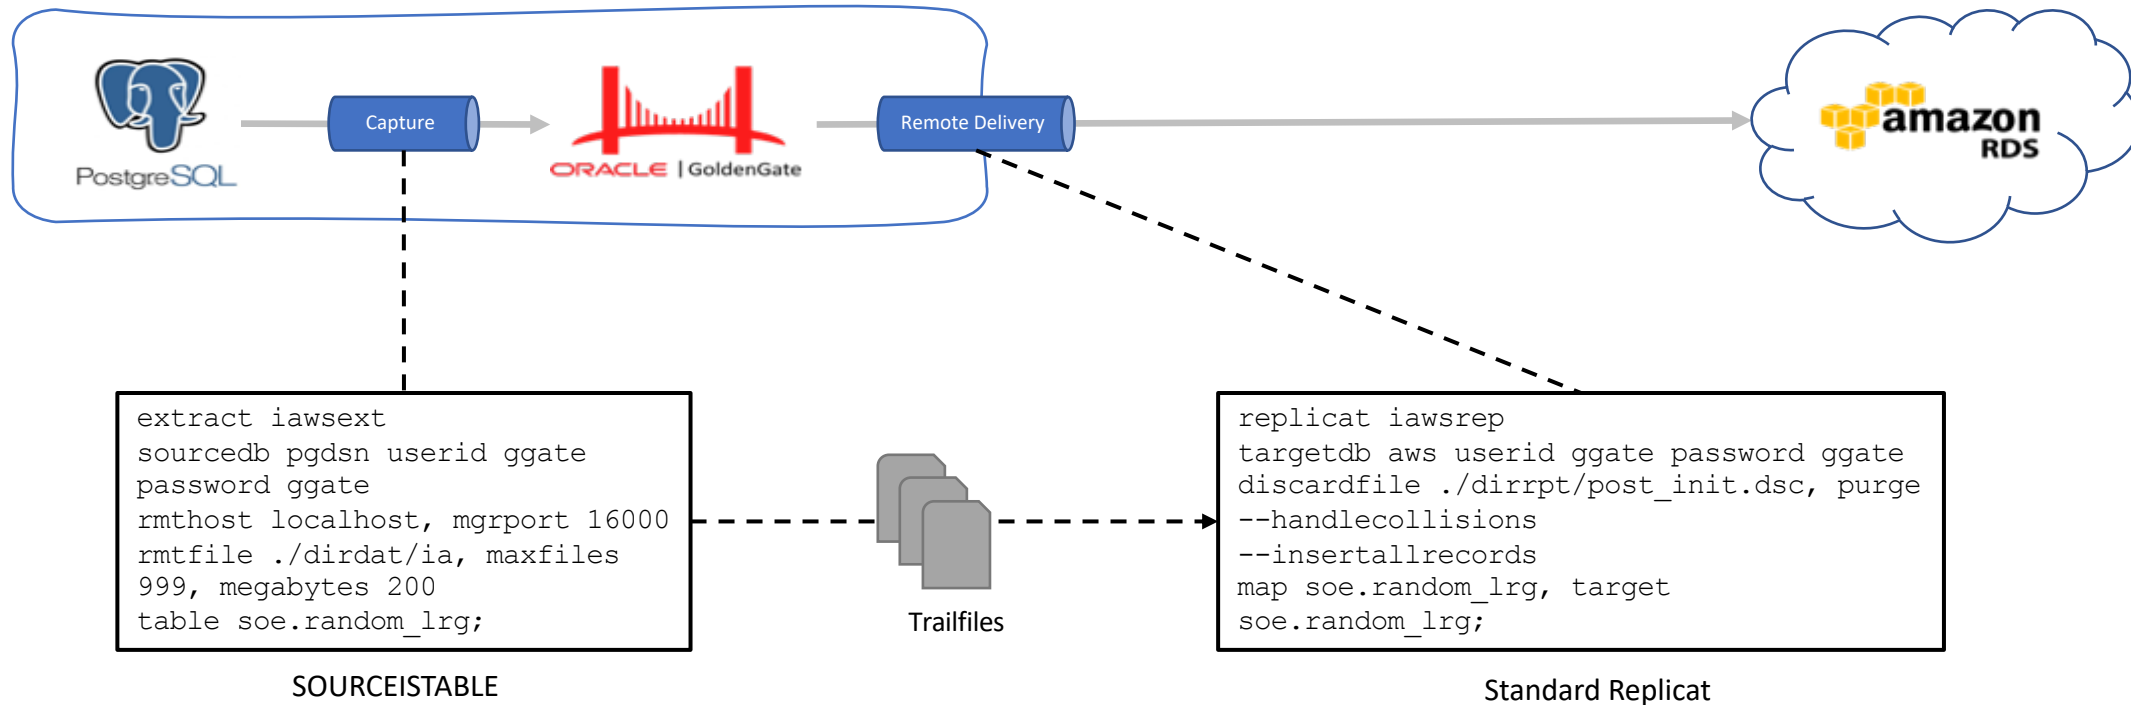
- Create Users/Role

```
Create user <role> with password '<passwd>' login;  
Grant all on tablespace <ts_name> to <user>;  
Grant all on schema <schema> to <user>;
```

Note: Used AWS Free Tier during testing

On-Premises to Cloud


Initial Load - Tables



Note: Used AWS Free Tier during testing

On-Premises to Cloud

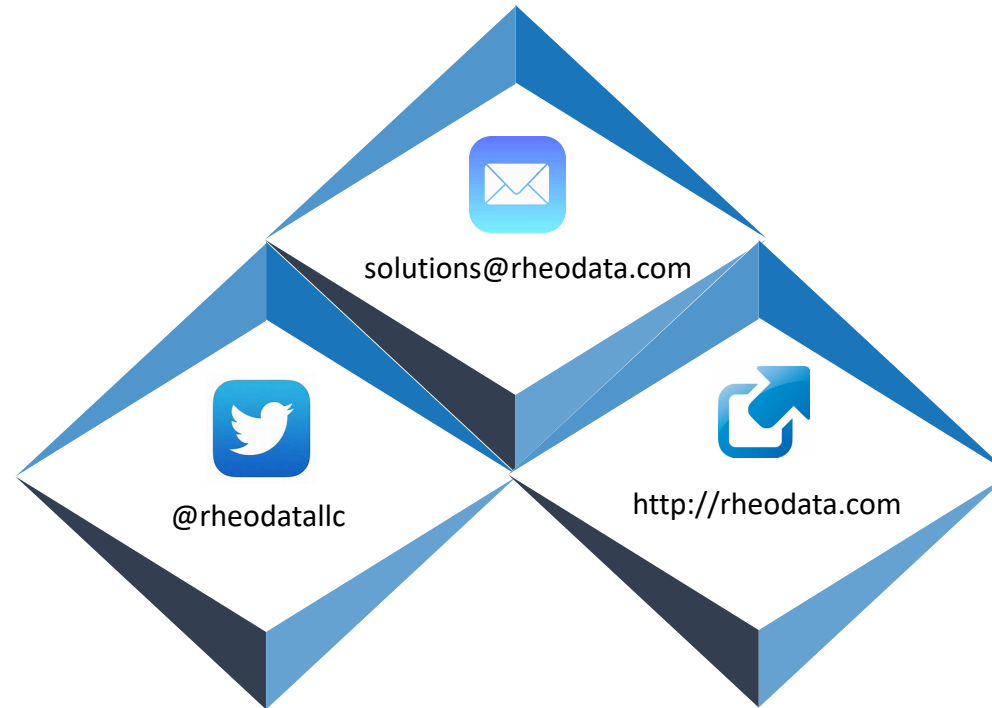
Thinks to consider



- Network bandwidth could be a bottleneck
 - Connection should be $\geq 100\text{ms}$; anything greater will cause latency
- Size of the RDS instance determine processing speed
 - EC2 and other cloud-based PostgreSQL environment – may provide more direct access and options
- All connections work over ODBC (Data Direct drivers)
 - Remote Capture
 - Remote Delivery

Note: Used AWS Free Tier during testing

Contact RheoData



Why RheoData

Advantage:

RheoData services are specifically designed and optimized for customers who are looking for data migration and infrastructure solutions to maximize their investment on-premises and in the cloud. By leveraging the right combination of experience, expertise, and up-to-date product knowledge, Rheodata can customize a solution to fit your needs. Our staff offers top-notch service that is backed by a proven track record of excellence, weather performed on-site or remote. RheoData provides services from the beginning of the architecture design process through the implementation phase, ensuring an aligned and complete project based on overall customer goals.

Dynamic Teaming:

- Supplemental IT Services for Customers
- Onshore Resources
- Complete Infrastructure Services/Management

Consulting Services:

- Cloud Platforms (OCI, Azure, AWS)
- Database Migrations and Upgrades
- Streaming Data Platforms (Apache Kafka)
- Engineered Systems
- Architecture Assessments/Implementations
- Enterprise Monitoring

Strategic Services:

- Software Selection Assistance
- Solution Architecture
- Proof of Concepts (POC)
- Implementation, Upgrade, Migration and Conversion Estimates

Tools & Platforms:

Database:

- Oracle
- Microsoft SQL Server
- PostgreSQL
- MySQL

Data Integration:

- Oracle GoldenGate
- Apache Kafka
- HVR

Engineered Systems:

- Exadata
- Oracle Database Appliance

Cloud Platforms:

- AWS
- Microsoft Azure
- Oracle Cloud Infrastructure

Infrastructure as Code:

- Terraform
- Vault

Staff:

RheoData is founded and owned by a former Oracle Product Manager that brings 25+ years of industry experience to any project. This background and expertise is critical for customer project success and allows RheoData exceptional access to support channels as well as product development to ensure success of any project.

Contact:

RheoData, LLC
8410 Fairthorn Way
Douglasville, GA 30135

solutions@rheodata.com

+1 (770) 876-5498

www.rheodata.com

RheoData makes everything easy



Implementation made easy!!

Data movement made easy!!

Management made easy!!