

Scaling AWS Redshift Concurrency with Postgres

By Elliott Cordo, Will Liu, Paul Singman

EQUINOX

Integrated luxury and lifestyle company with offerings centered on movement, nutrition, and regeneration

blink FITNESS

PURE YOGA

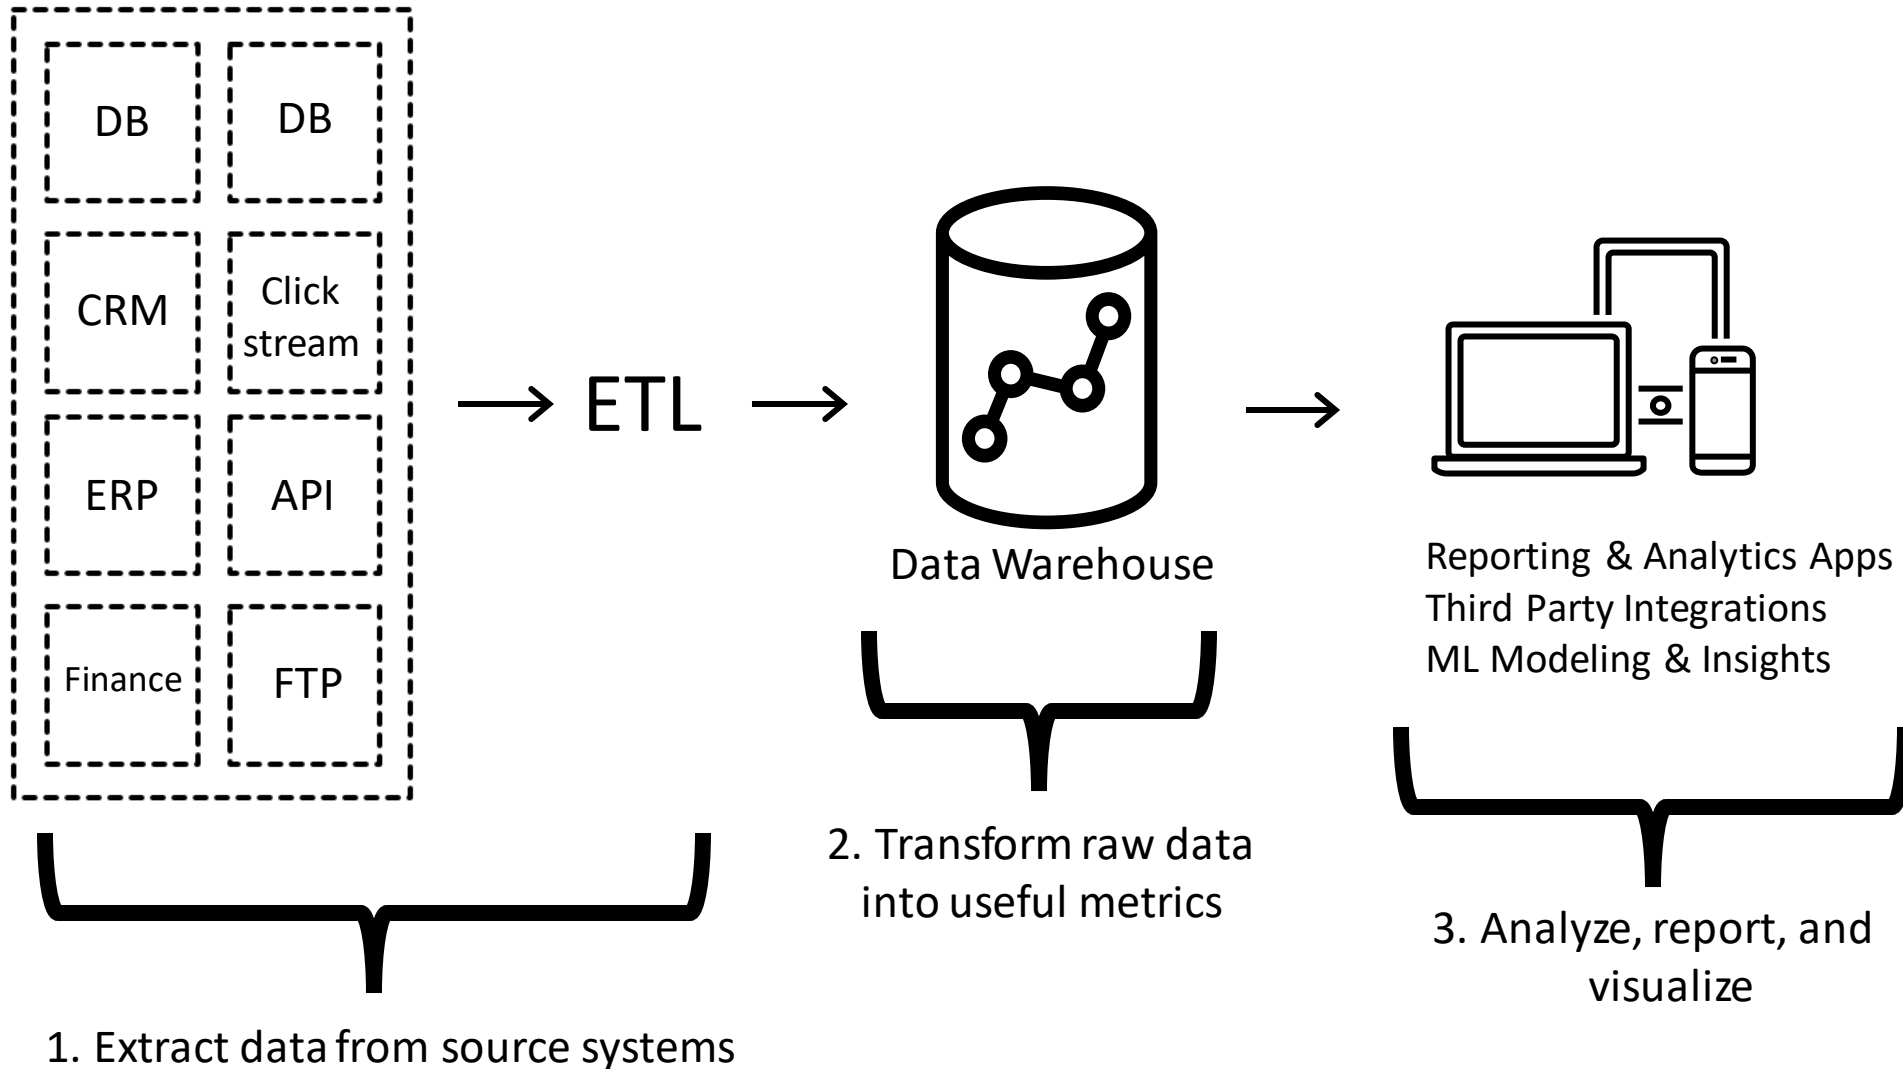
SOULCYCLE

FURTHERMORE
FROM EQUINOX

INTRODUCING
**EQUINOX
HOTELS**

we operate more than 200 locations within every major city across the country
in addition to London and Canada

Analytics Overview



Why A Data Warehouse?

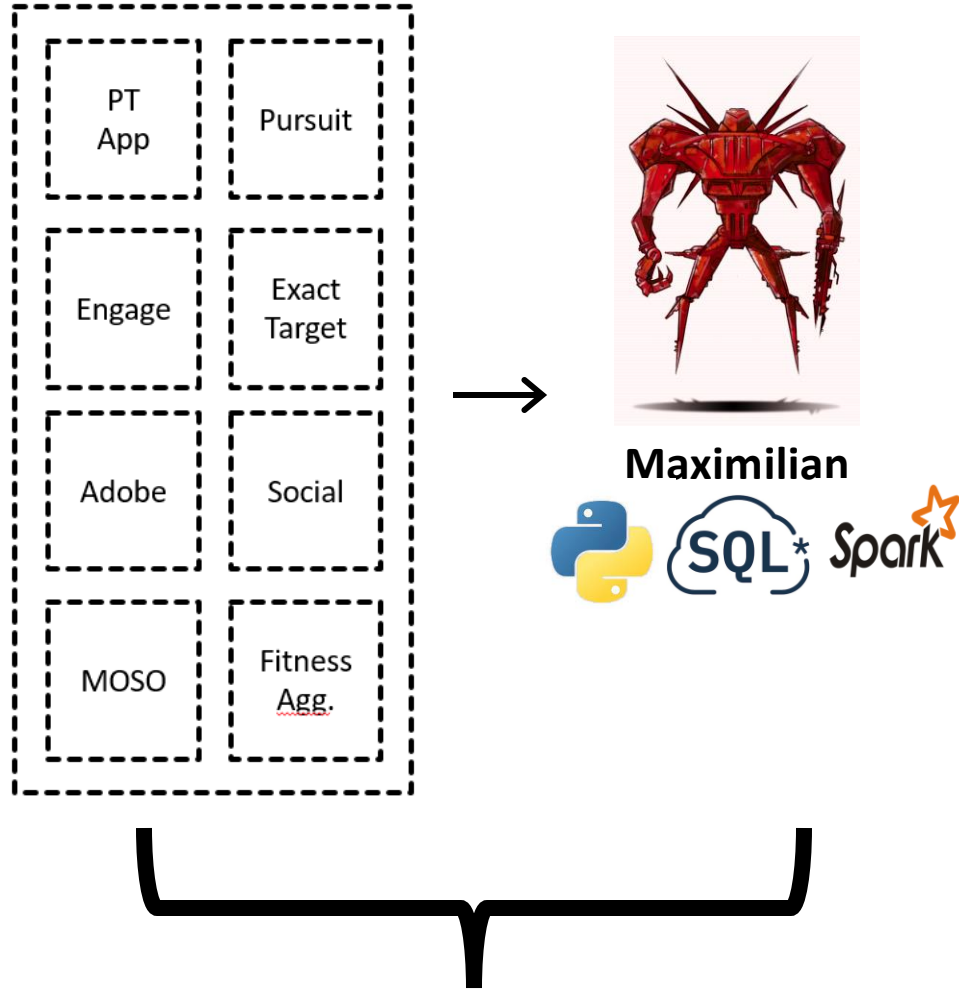
- Prevent data discrepancies
- Lower employee learning curve
- Avoid duplicating logic in multiple systems
- Isolate production DBs from analytic workloads

Why A Data Warehouse?

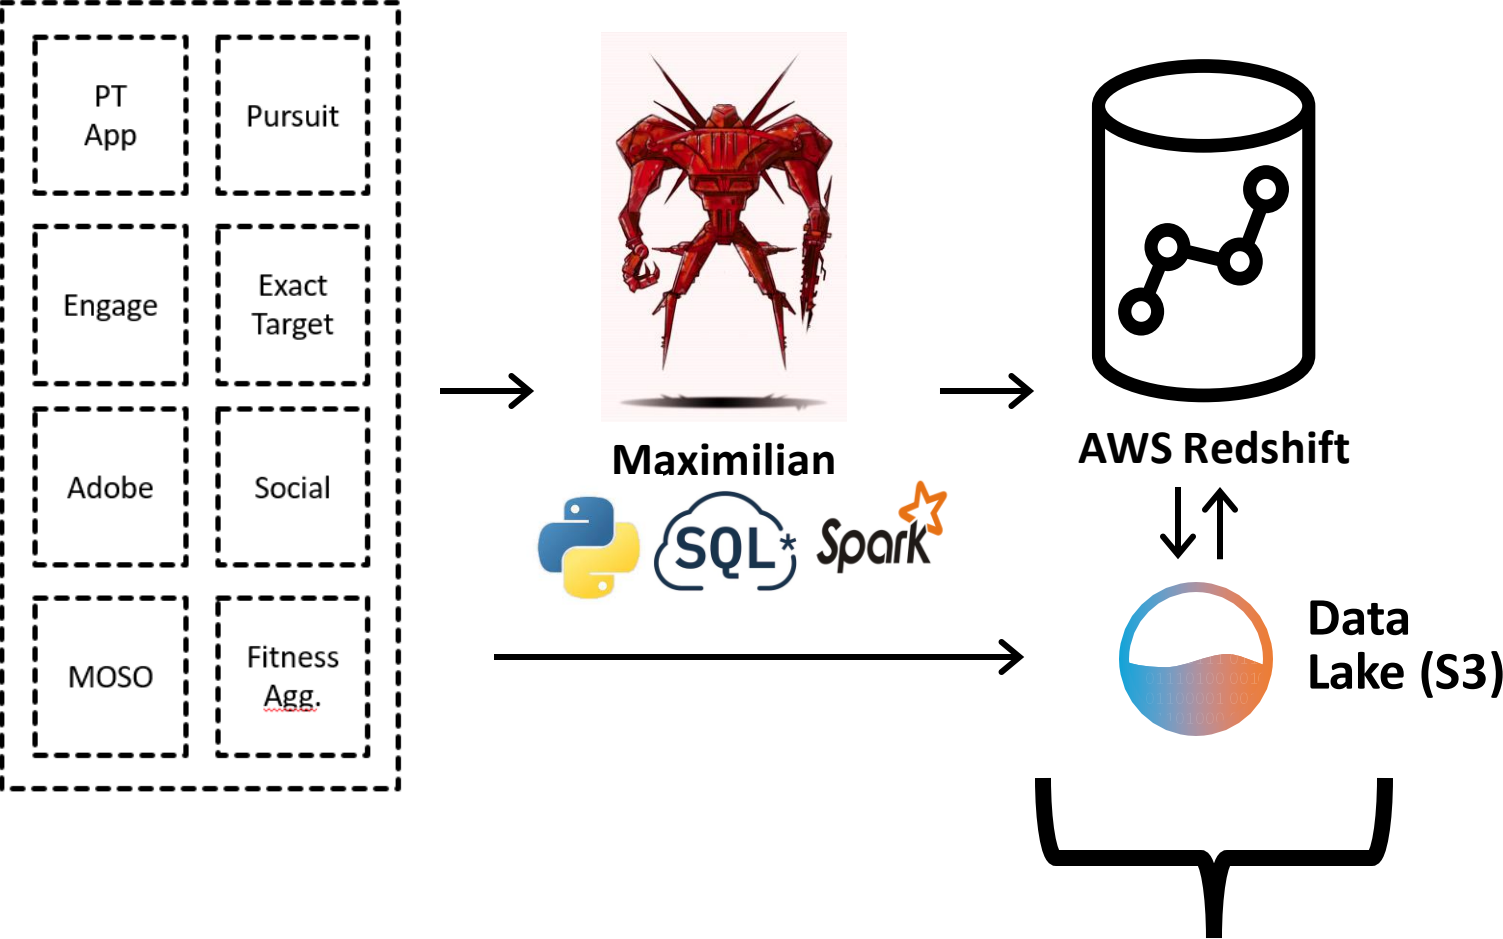
- Prevent data discrepancies
- Lower employee learning curve
- Avoid duplicating logic in multiple systems
- Isolate production DBs from analytic workloads

Presents a **single point of failure** so we created a **data replication failover procedure**

Data Ecosystem @ Equinox



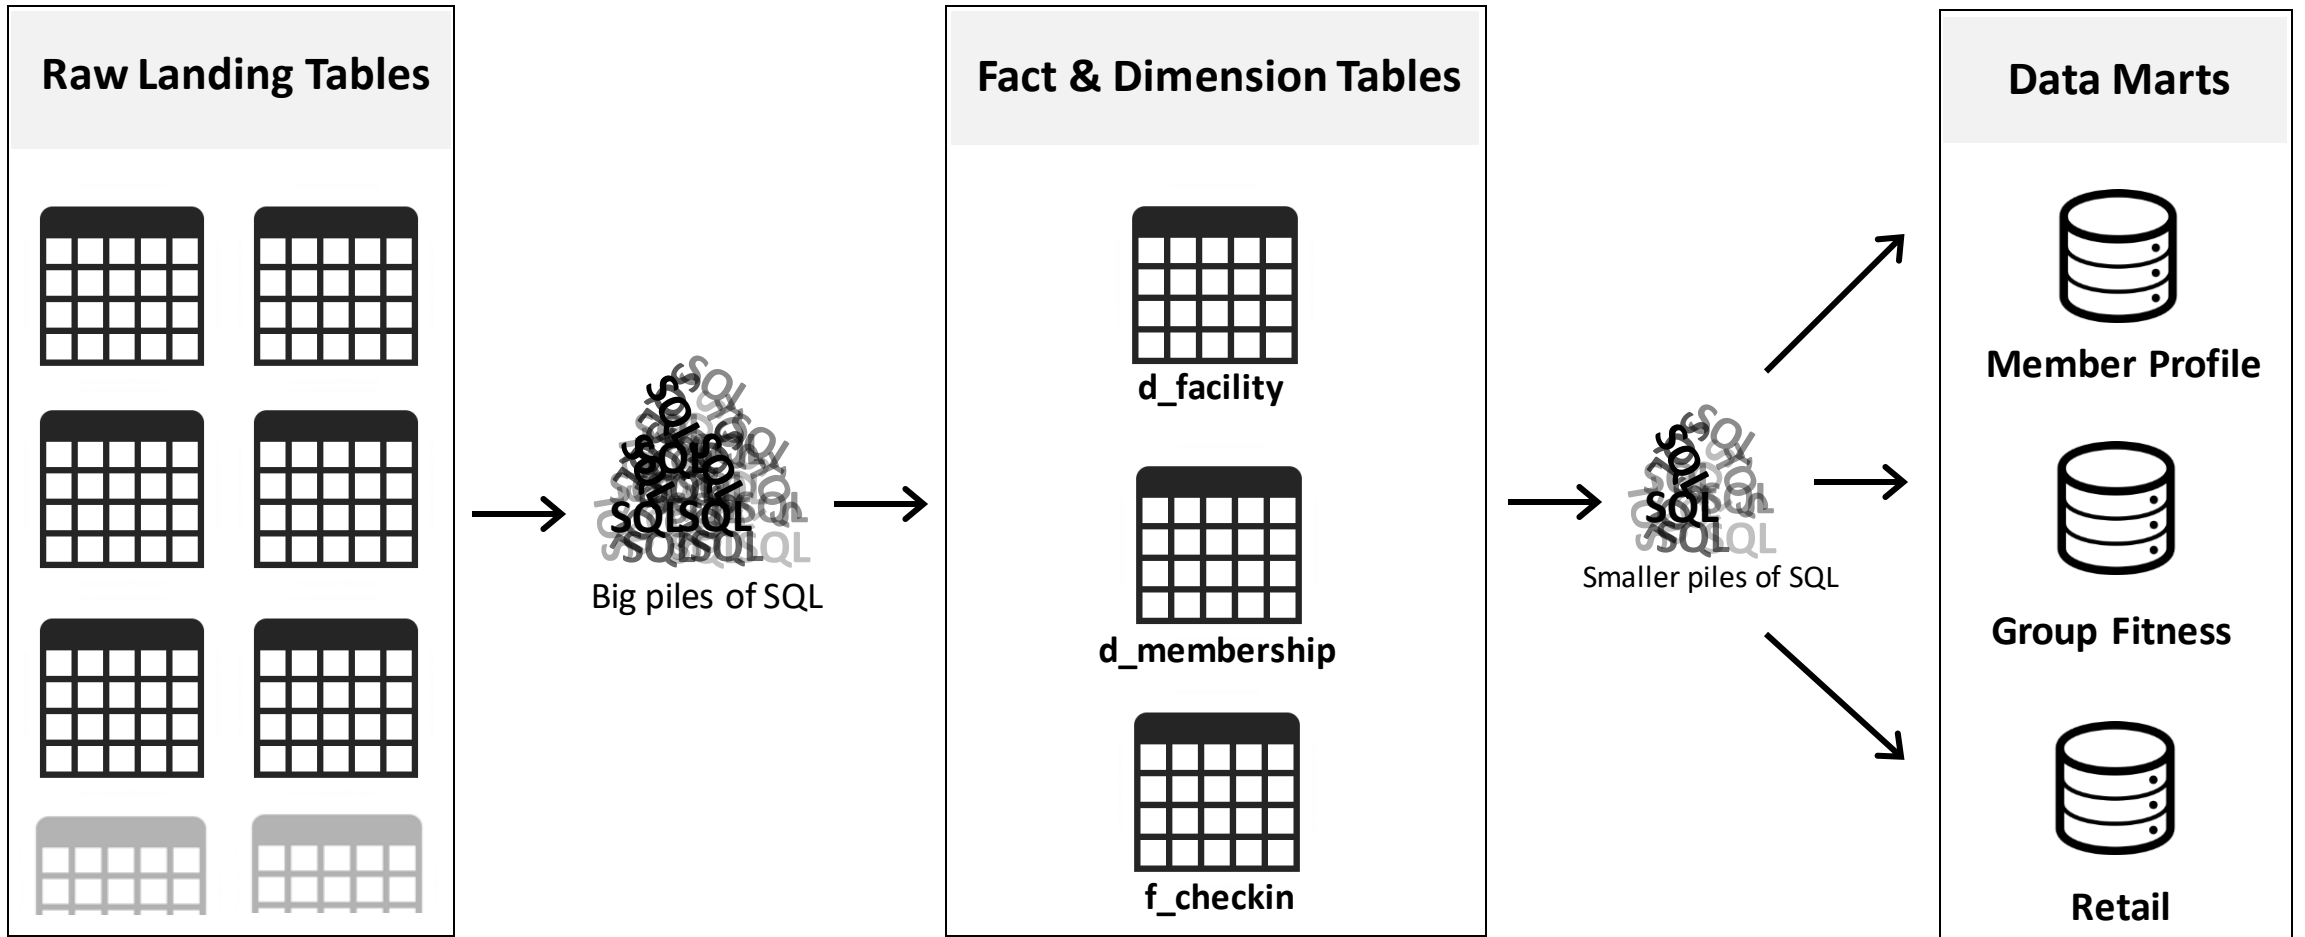
Data Ecosystem @ Equinox



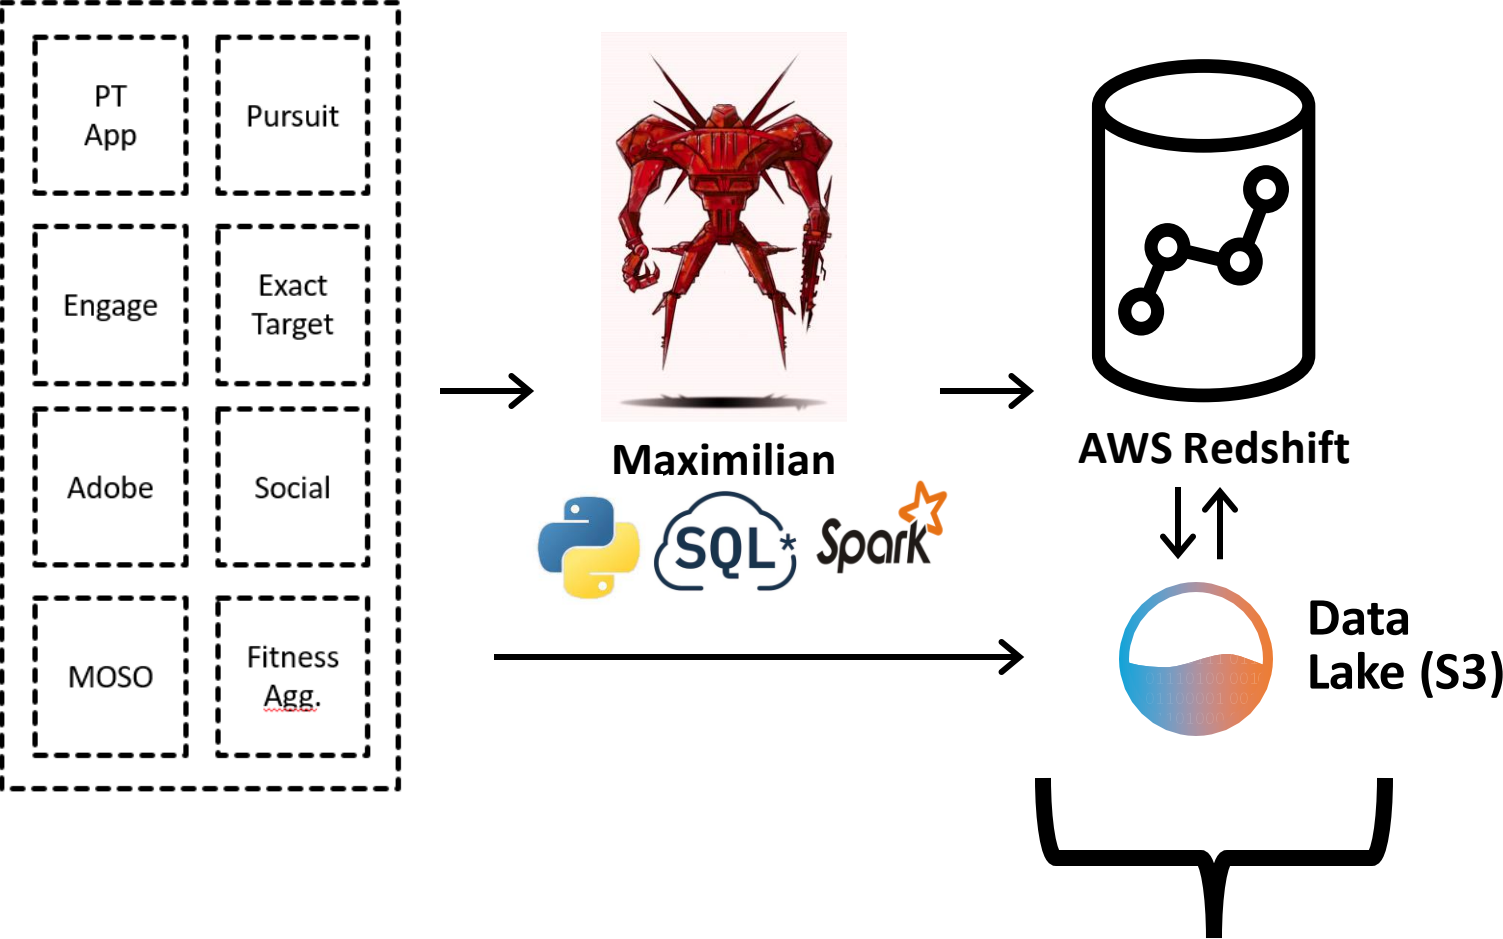
2. Transform raw data into useful metrics



Redshift Warehouse Structure (J.A.R.V.I.S.)



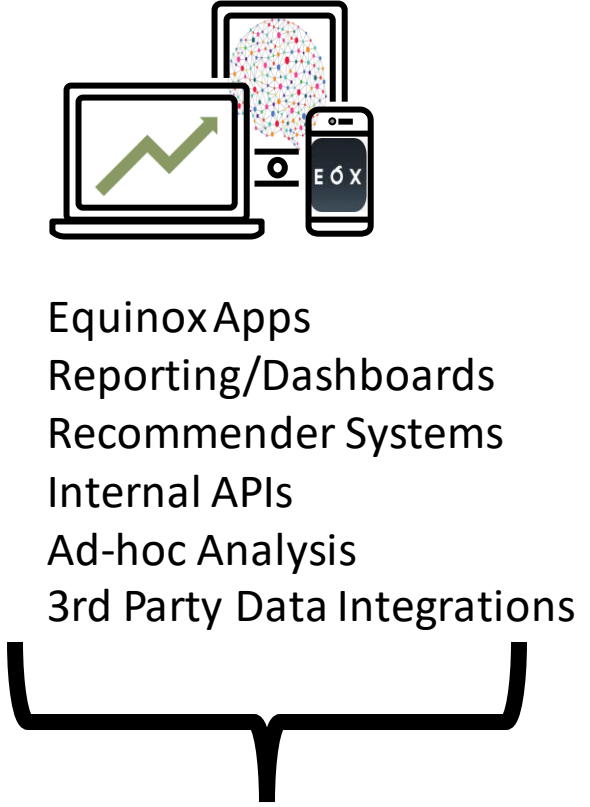
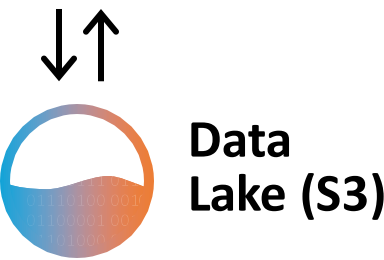
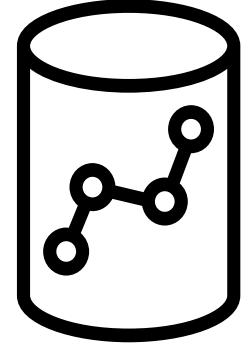
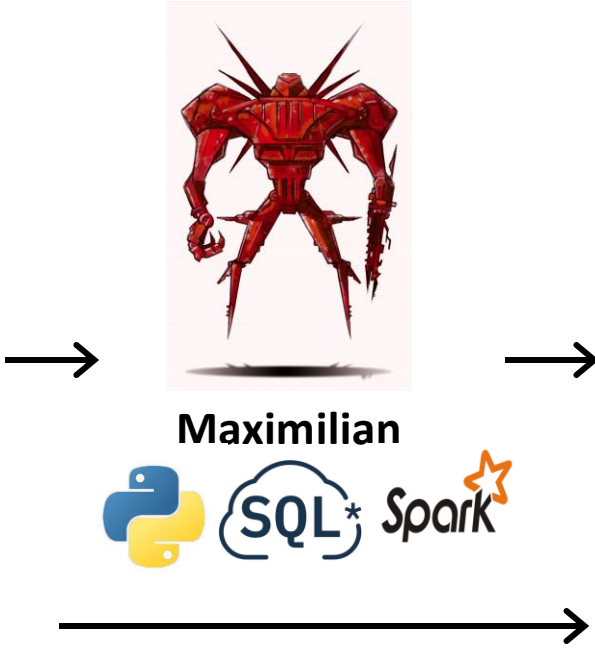
Data Ecosystem @ Equinox



2. Transform raw data into useful metrics

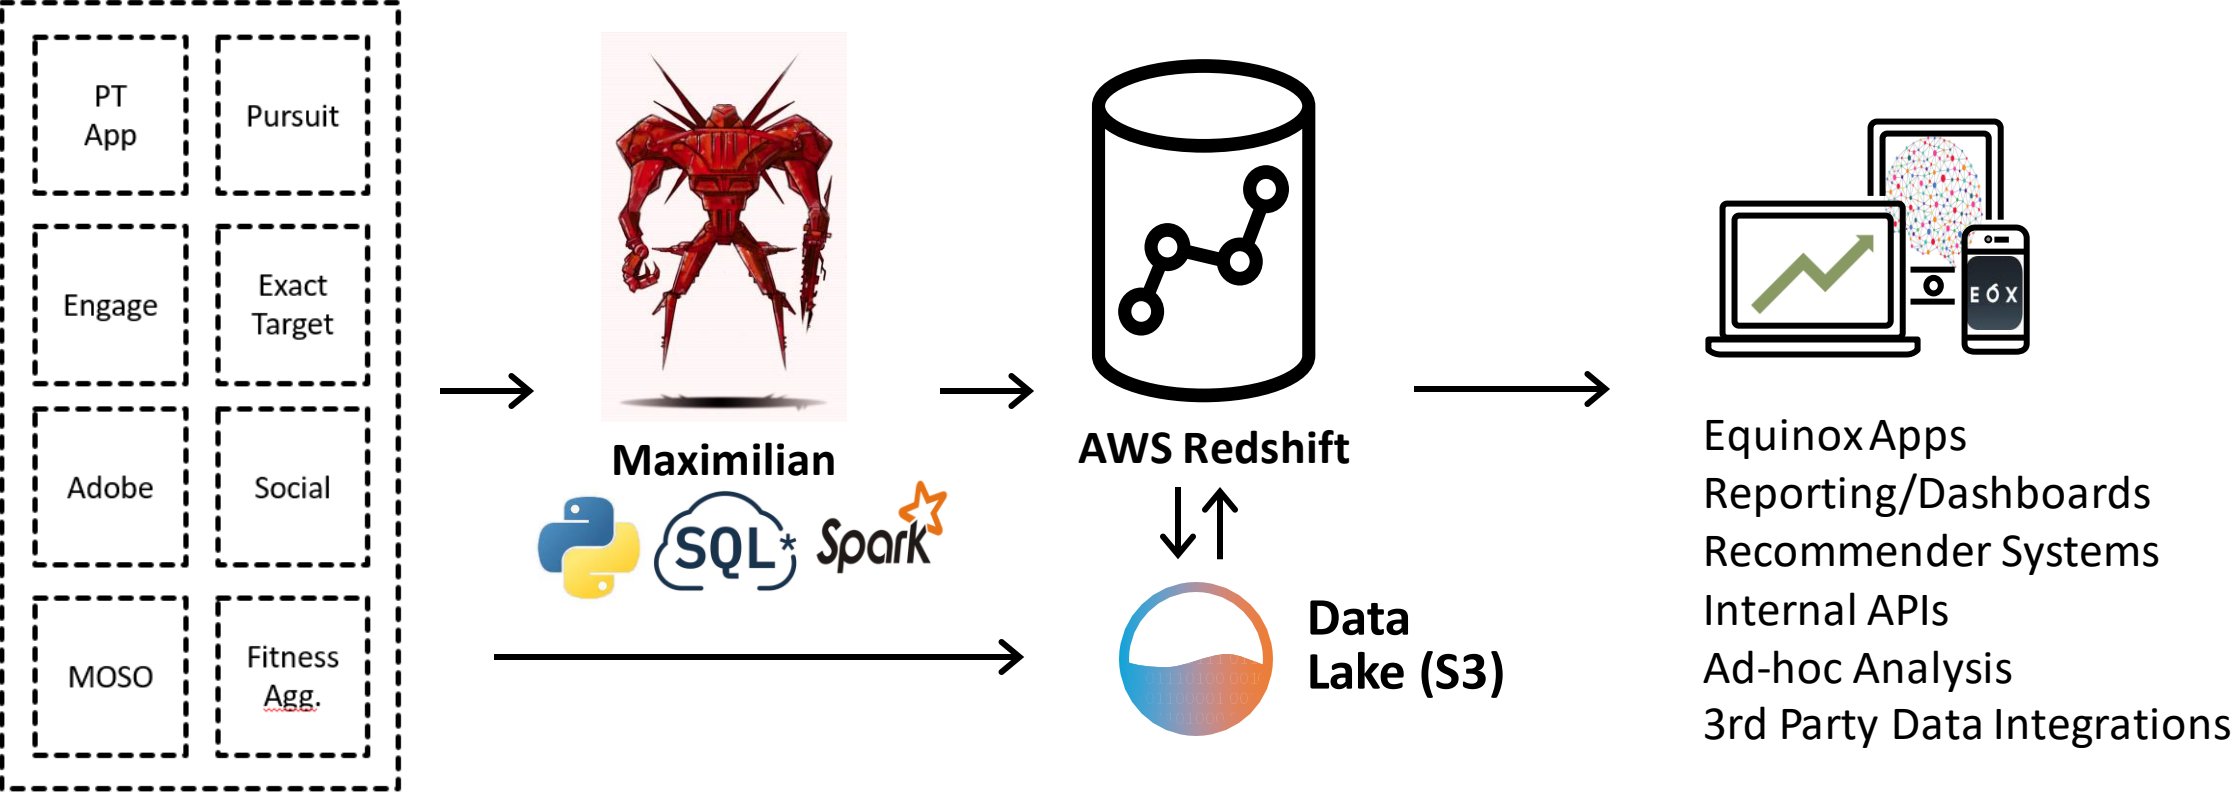
Data Ecosystem @ Equinox

PT App	Pursuit
Engage	Exact Target
Adobe	Social
MOSO	Fitness Agg.



3. Analyze, report, and visualize

Data Ecosystem @ Equinox

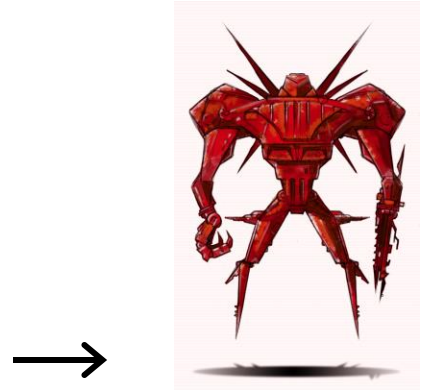


4. Data Monitoring & Quality

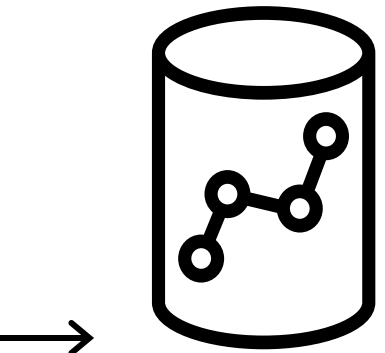
Data Ecosystem @ Equinox



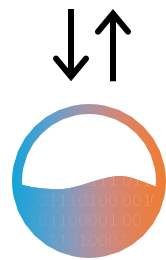
PT App	Pursuit
Engage	Exact Target
Adobe	Social
MOSO	Fitness Agg.



Maximilian
Python SQL* Spark



AWS Redshift

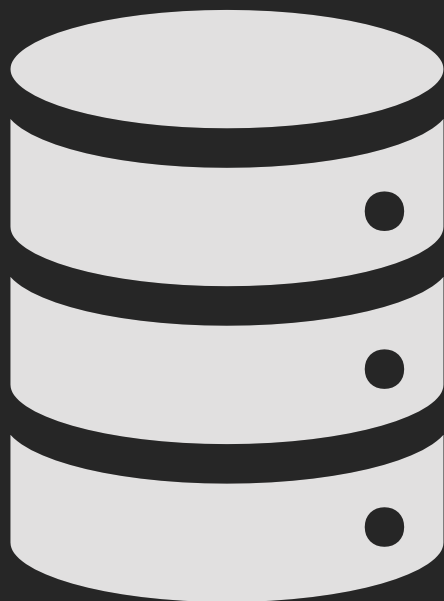


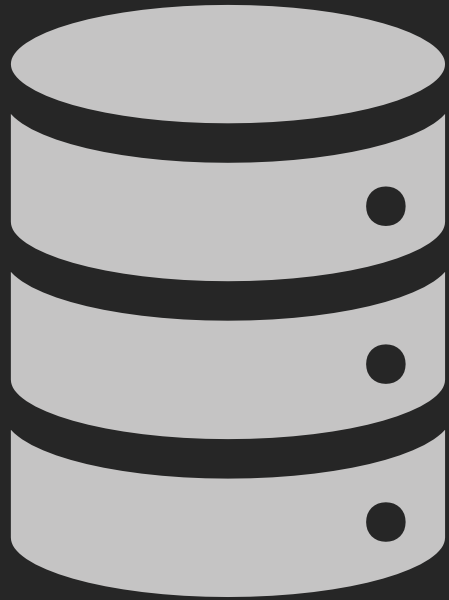
Data Lake (S3)



Equinox Apps
Reporting/Dashboards
Recommender Systems
Internal APIs
Ad-hoc analysis
3rd party data integrations

4. Data Monitoring & Quality





On-line Transactional Processing
(OLTP)



On-line Analytical Processing
(OLAP)

Key Features of Redshift

Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- Immutable 1MB block storage
- Massively-parallel processing compute engine
- Native multi-node (leader + workers) architecture
- Distribution key & sort key table settings
- Workload Management queue settings

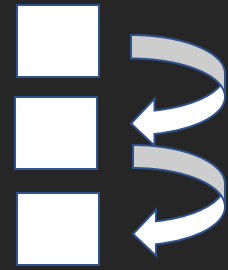
Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- **Column-oriented storage**
- Immutable 1MB block storage
- Massively-parallel processing compute engine
- Native multi-node (leader + workers) architecture
- Distribution key & sort key table settings
- Workload Management queue settings



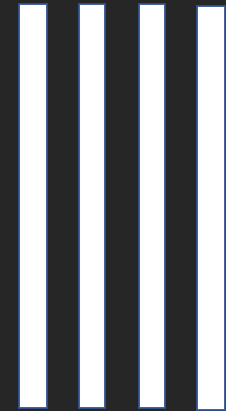
Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- **Immutable 1MB block storage**
- Massively-parallel processing compute engine
- Native multi-node (leader + workers) architecture
- Distribution key & sort key table settings
- Workload Management queue settings



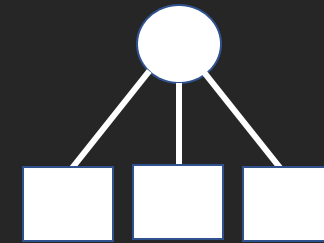
Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- Immutable 1MB block storage
- **Massively-parallel processing compute engine**
- Native multi-node (leader + workers) architecture
- Distribution key & sort key table settings
- Workload Management queue settings



Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- Immutable 1MB block storage
- Massively-parallel processing compute engine
- **Native multi-node (leader + workers) architecture**
- Distribution key & sort key table settings
- Workload Management queue settings



Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- Immutable 1MB block storage
- Massively-parallel processing compute engine
- Native multi-node (leader + workers) architecture
- **Distribution key & sort key table settings**
- Workload Management queue settings



Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- Immutable 1MB block storage
- Massively-parallel processing compute engine
- Native multi-node (leader + workers) architecture
- Distribution key & sort key table settings
- Workload Management queue settings



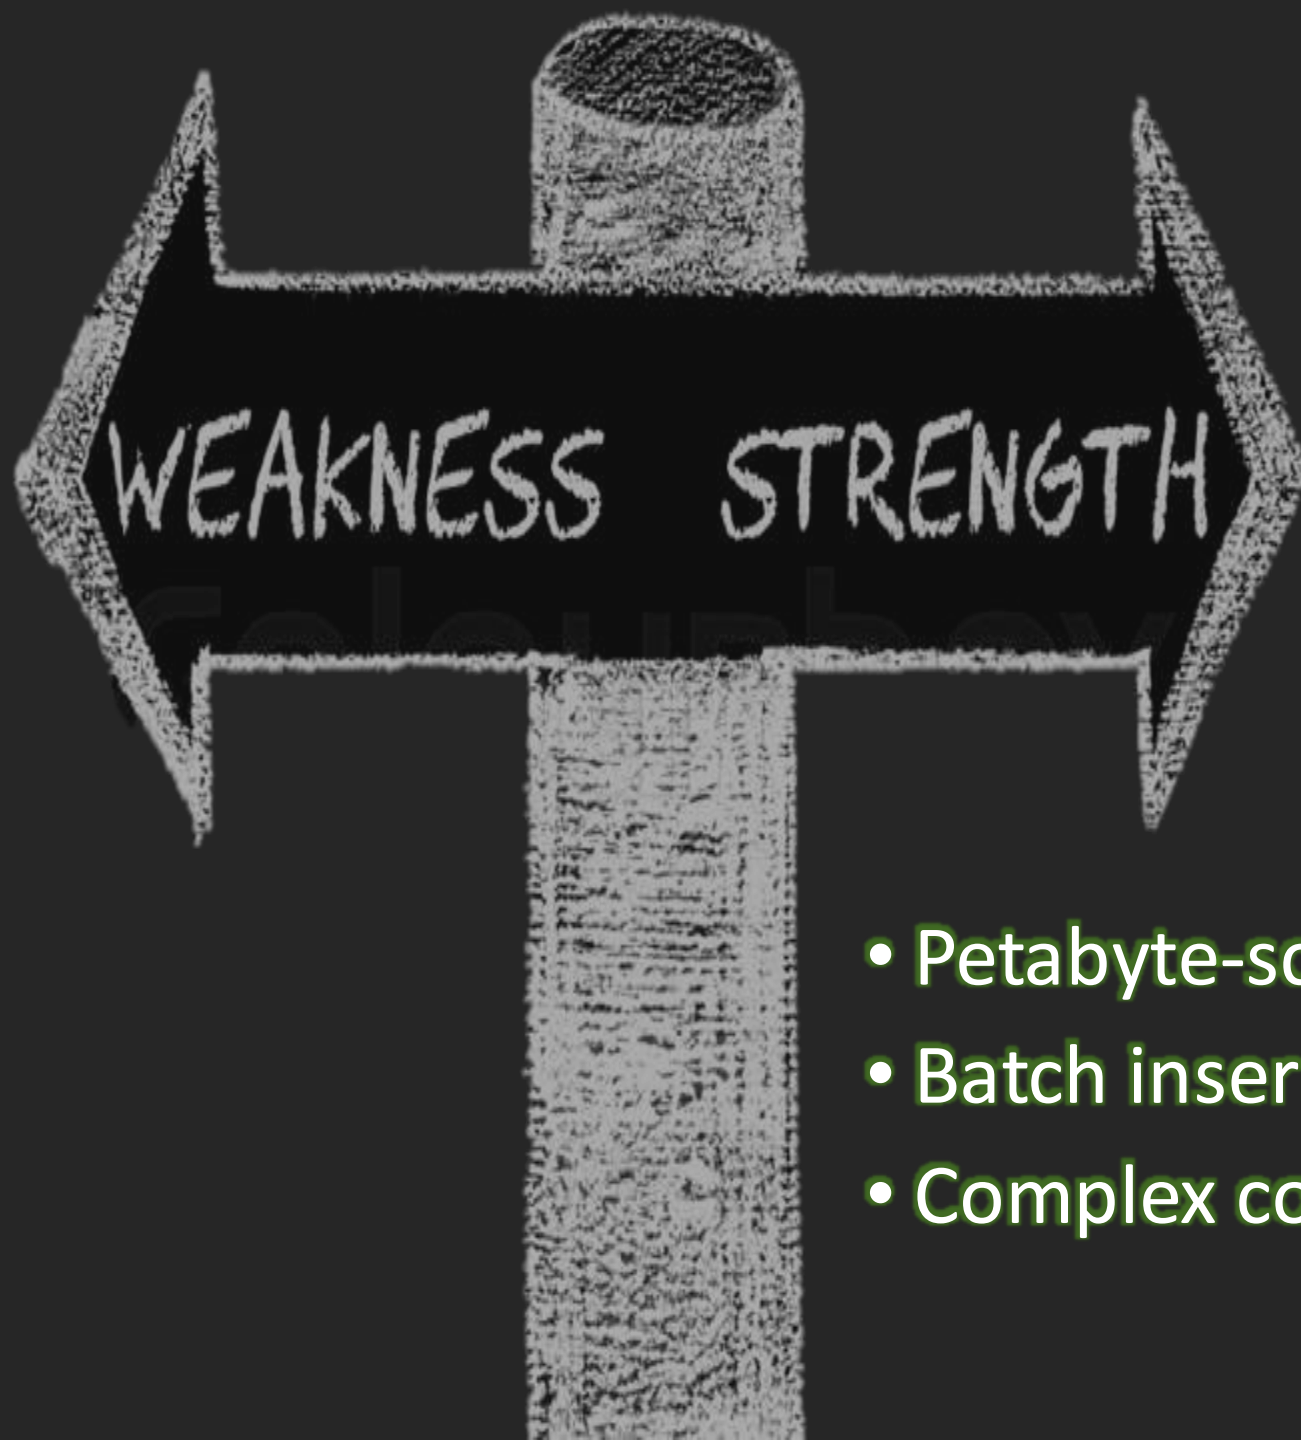
Key Features of Redshift

- Released by AWS in early 2013, based on Postgres v8.0.2
- Column-oriented storage
- Immutable 1MB block storage
- Massively-parallel processing compute engine
- Native multi-node (leader + workers) architecture
- Distribution key & sort key table settings
- Workload Management queue settings

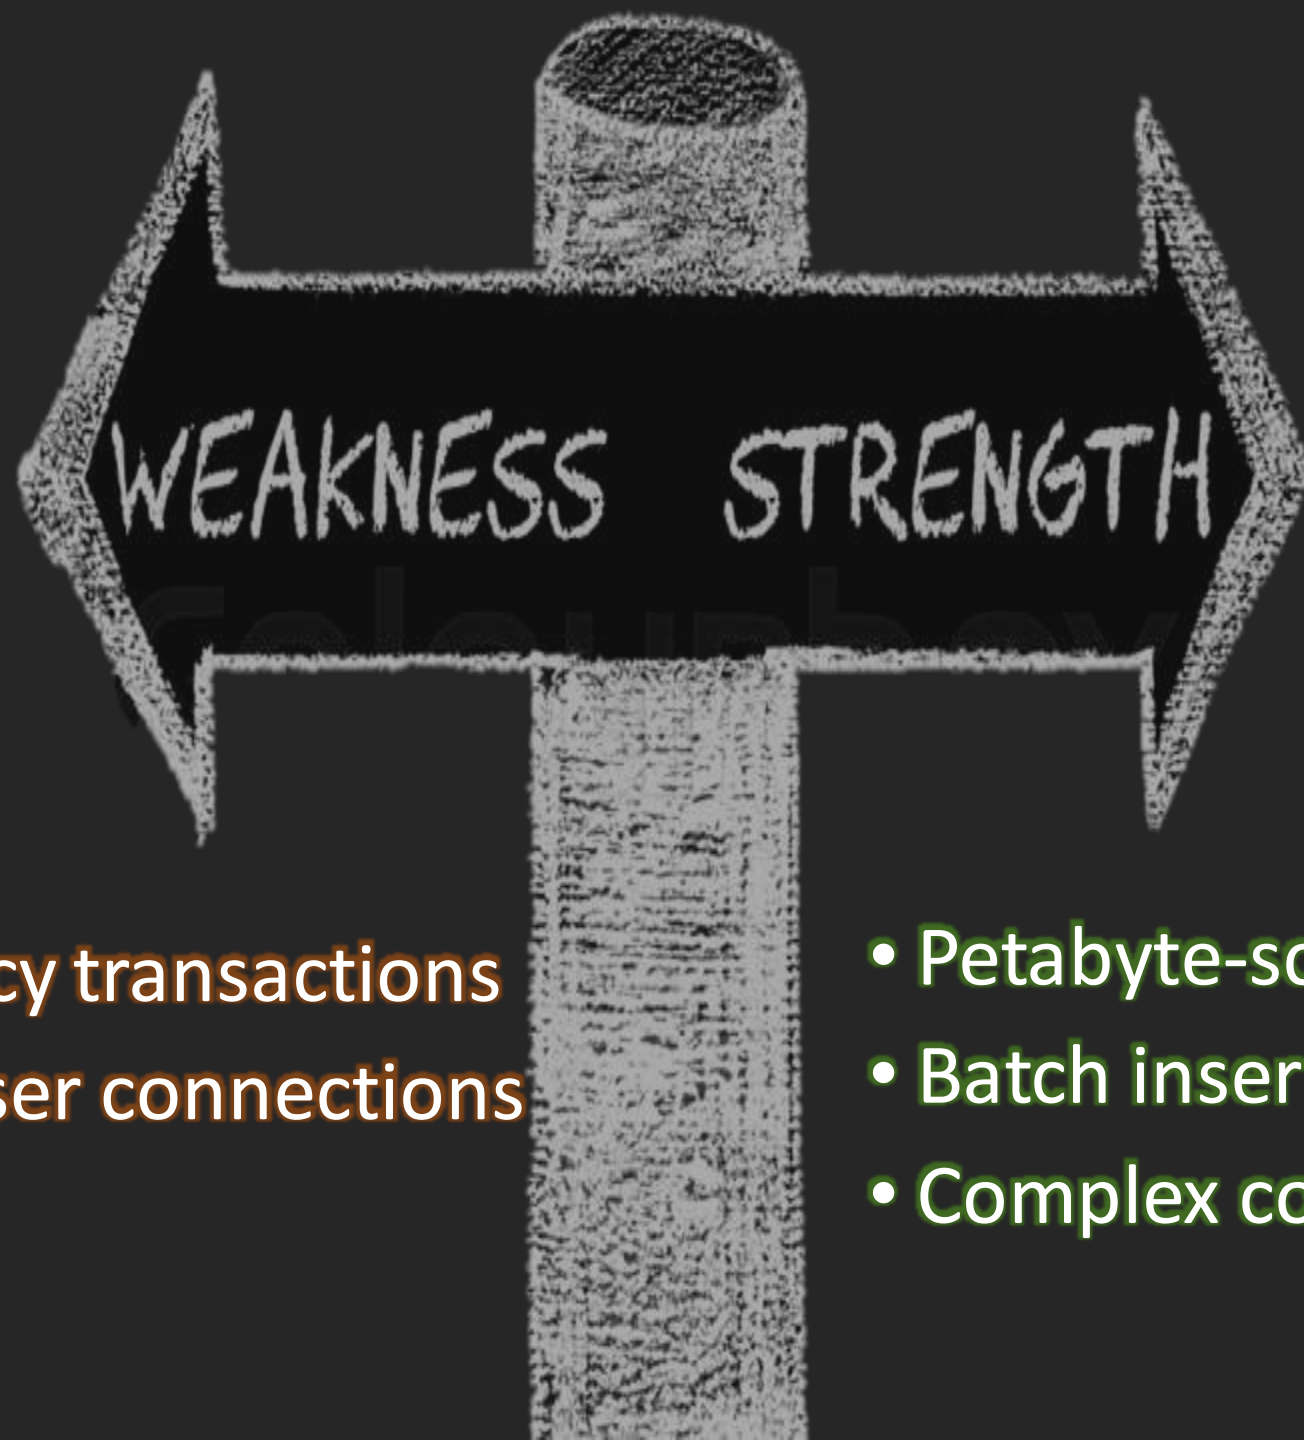


WEAKNESS

STRENGTH



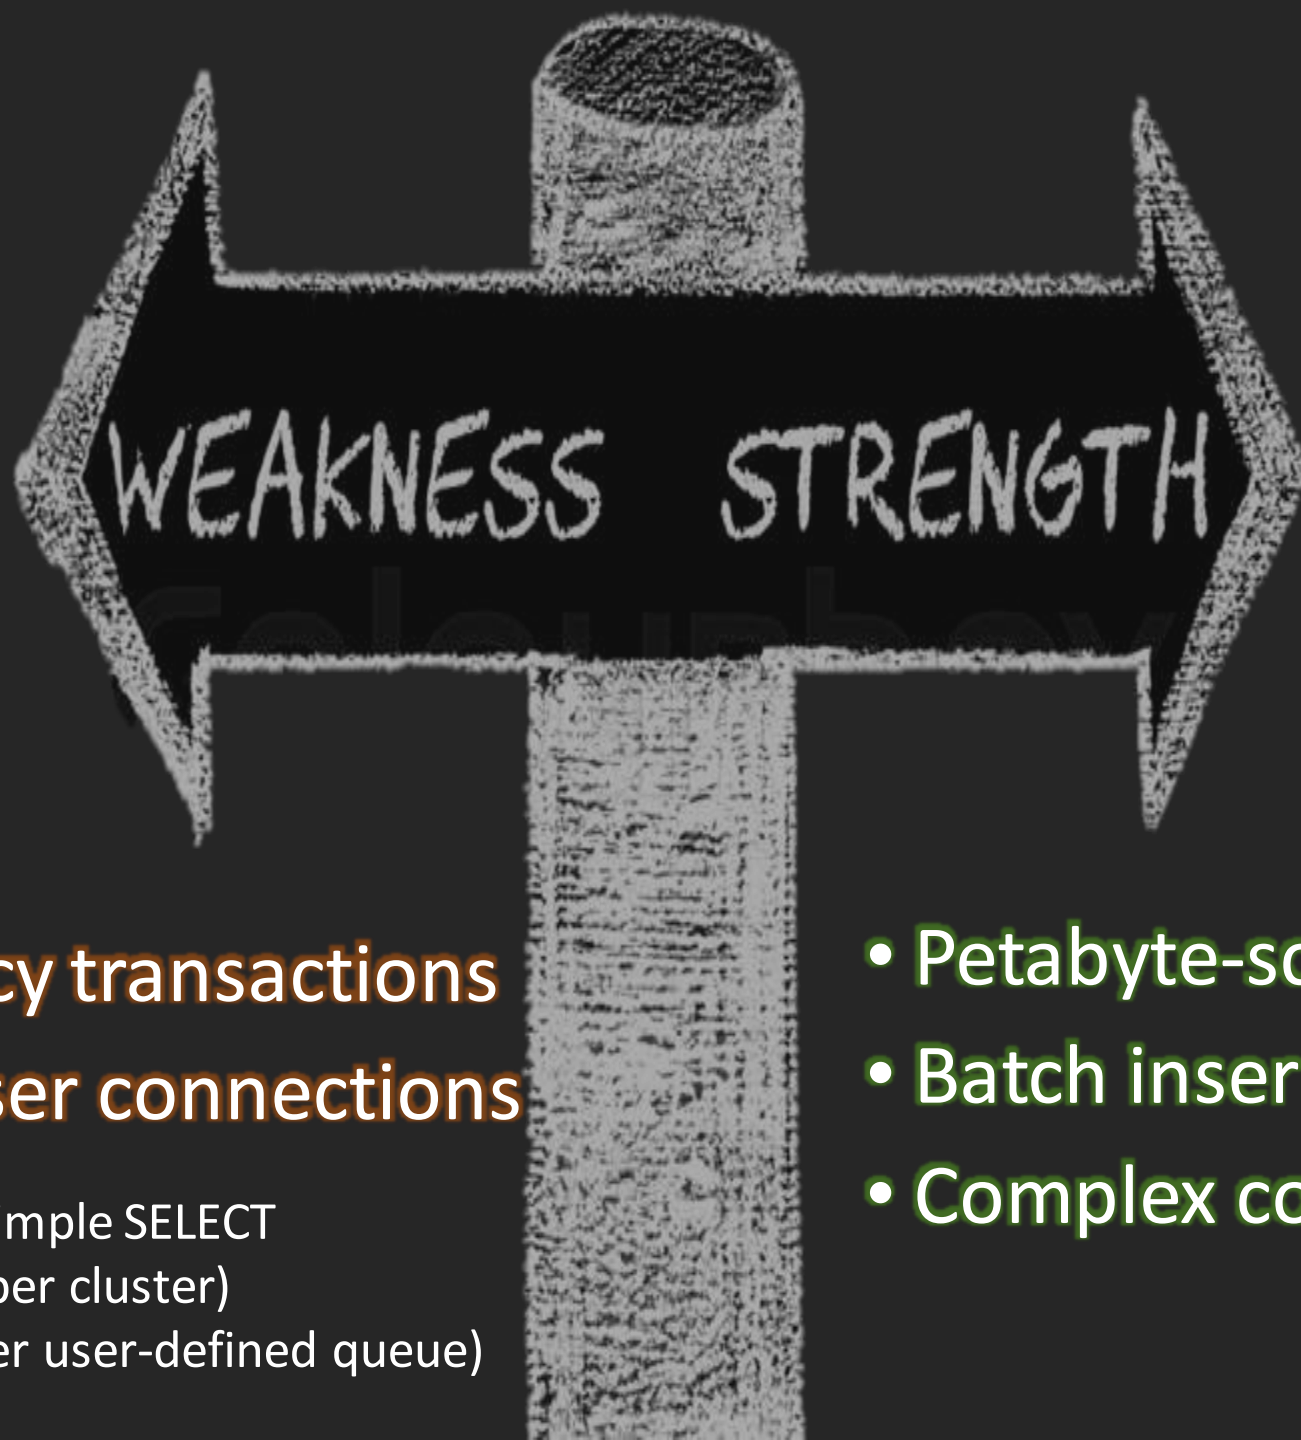
- Petabyte-scale disk storage
- Batch insertions & retrieval
- Complex computations



WEAKNESS STRENGTH

- High-frequency transactions
- Concurrent user connections

- Petabyte-scale disk storage
- Batch insertions & retrieval
- Complex computations

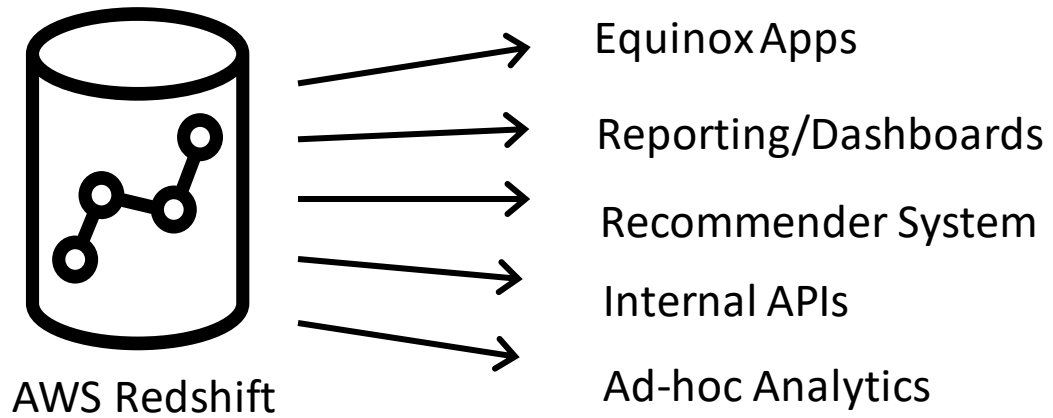


- High-frequency transactions
- Concurrent user connections

10x - 100x slower on simple SELECT
500 connection limit (per cluster)
50 connection limit (per user-defined queue)

- Petabyte-scale disk storage
- Batch insertions & retrieval
- Complex computations

Warehouse Consumers



Warehouse Consumers



Sales Review Details

From Date: 3/1/2019 To Date: 3/17/2019

Account Cancellations

3DayCXL

Acct Code	Club	Member Id	Member Name	Member Status	Date	Payment Type	Class	Commission Status	Count
115	Park Avenue	1003931724	Mercino, Pedro	Cancelled	3/6/19	BM	Select Monthly	NO COMM	-1
115	Park Avenue	1003921211	Mohr, Anthony	Cancelled	3/12/19	PIF	All Access Yearly	INC SIGNATURES, INCOMPLETE CONTACT INFO	-1
Total Count:									-2

3DayReactivate

Acct Code	Club	Member Id	Member Name	Member Status	Date	Payment Type	Class	Commission Status	Count
115	Park Avenue	1003931724	Mercino, Pedro	Active	3/6/19	BM	All Access Monthly	DUE 1	1
Total Count:									1

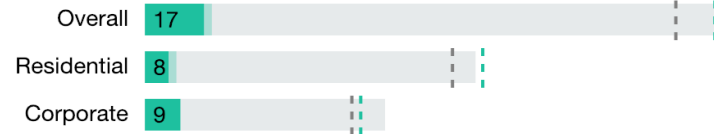
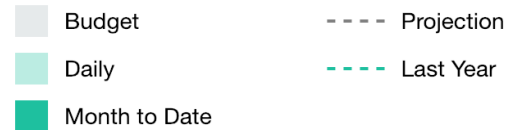
New

Acct Code	Club	Member Id	Member Name	Member Status	Date	Payment Type	Class	Commission Status	Count
115	Park Avenue	1001742000	Mercino, Ryan	Active	3/4/19	BM	All Access Monthly	DUE 3	1
115	Park Avenue	1003929636	Morales, Mayra	Active	3/4/19	BM	Select Monthly	DUE 1	1
115	Park Avenue	1003900694	Khan, Asif	Active	3/5/19	BM	All Access Monthly	INC CORPORATE	1
115	Park Avenue	1003932836	Mitten, Kristina	Active	3/5/19	BM	Select Monthly	INC SIGNATURES	1
115	Park Avenue	1003932984	Boer, Samantha	Active	3/5/19	BM	Select Monthly	DUE 1	1
115	Park Avenue	1001796524	Reuter, Paul	Active	3/6/19	BM	Select Monthly	DUE 1	1
115	Park Avenue	1003931724	Mercino, Pedro	Active	3/6/19	BM	Select Monthly	NO COMM	1
115	Park Avenue	1003198113	Simmons, Natalie	Pending Start	3/7/19	BM	All Access Monthly	DUE 1	1
115	Park Avenue	1003630343	Morales, Kelli	Pending Start	3/7/19	BM	Select Monthly	DUE 3	1
115	Park Avenue	1003932836	Mitten, Kristina	Active	3/5/19	BM	Select Monthly	INC SIGNATURES	1

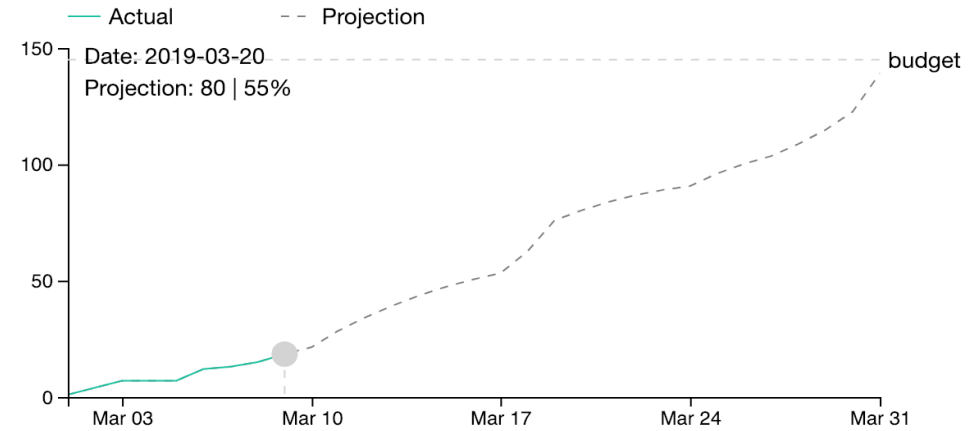


SALES

The progress bar charts show the daily actual, month to date and budget relationship of sales. The gray bars are the end of month budget as 100 percent, the green ones are the month to date actual in percentage of budget, with the exact value showing as label. The daily new actual value is marked out in lighter green color. Green dot lines and gray dot lines are the projection and last year value as percentage of each category.



Net Sales Projection Line



3 Days Cancel

1

3 Days Reactivate

1

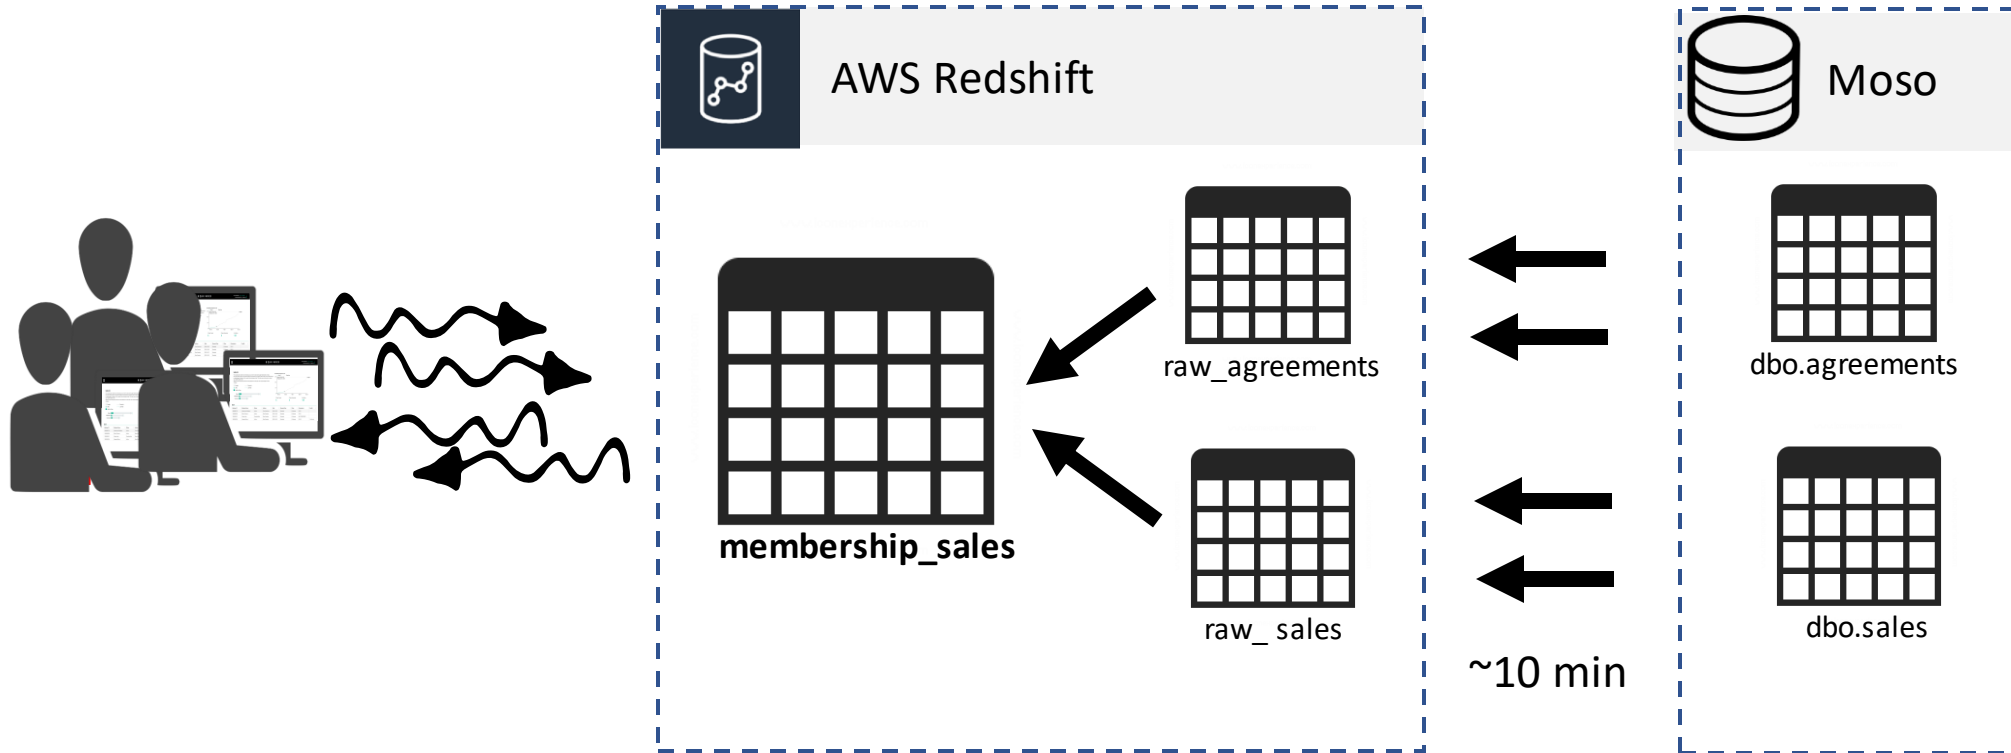
Net Sales

17

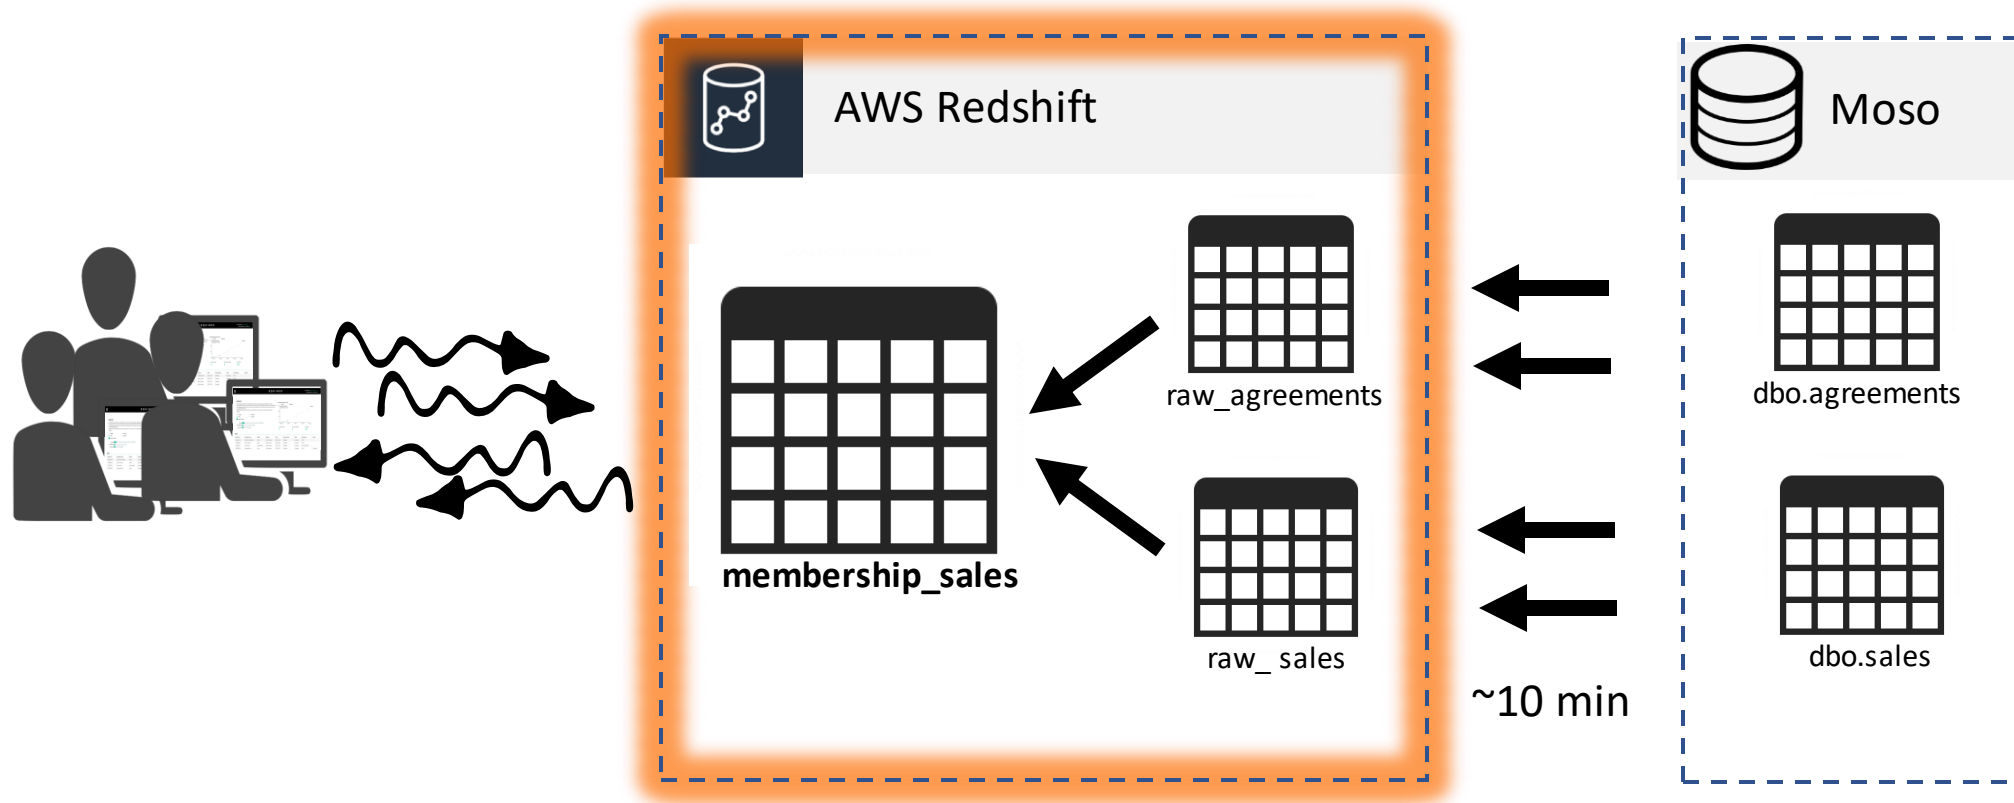
NEW

Member ID	Member Name	Status	Advisor	Date	Payment Type	Class	Commission	Equifit
1003858732	[Redacted]	Active	Tony Marchena	2019-03-02	Perpetual	All Access	DUE 1	
1003936006	[Redacted]	Active	Ryan Sanders	2019-03-06	Perpetual	All Access	DUE 1	
1003932537	[Redacted]	Pending Start	Ryan Sanders	2019-03-07	Perpetual	All Access	INC CORPORATE	
1003202293	[Redacted]	Active	Ryan Sanders	2019-03-03	Perpetual	All Access	DUE 2	Confirmed

Sales Reporting Architecture



Sales Reporting Architecture **Problem**



10x - 100x slower on simple SELECT
500 connection limit (per cluster)
50 connection limit (per user-defined queue)

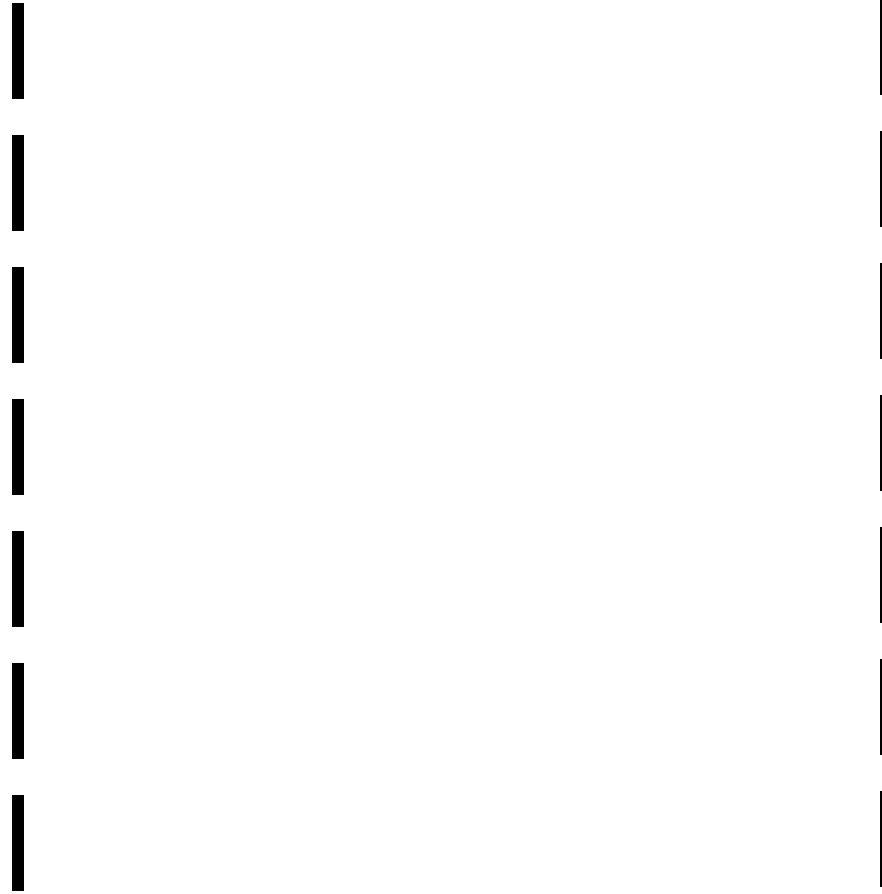
Potential Solutions

|
|
|
|
|
|
|

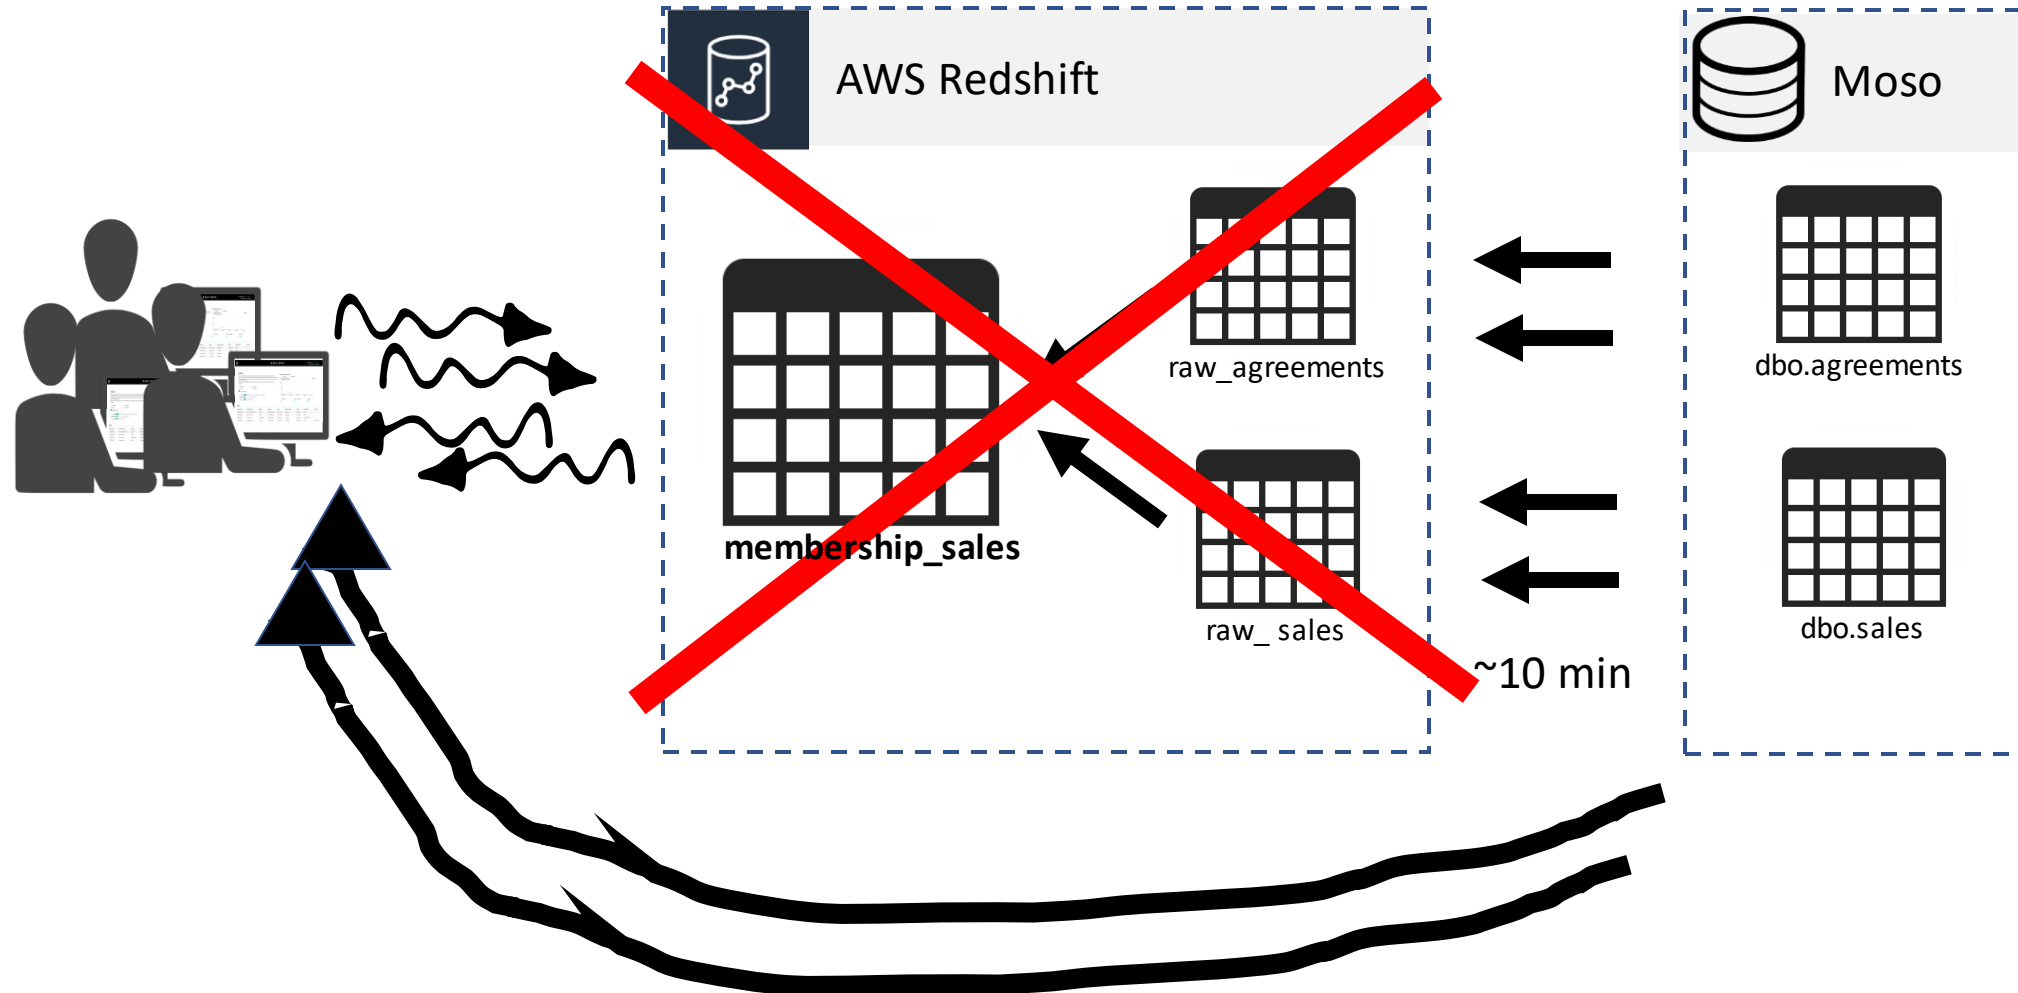
|
|
|
|
|
|
|

Potential Solutions

1. Pull from source DB



Sales Reporting Architecture



Potential Solutions

1. Pull from source DB

Pros

- Source DB is OLTP

Cons

- Sales logic is complex!
- Burdens prod DB

Potential Solutions

1. Pull from source DB

Pros

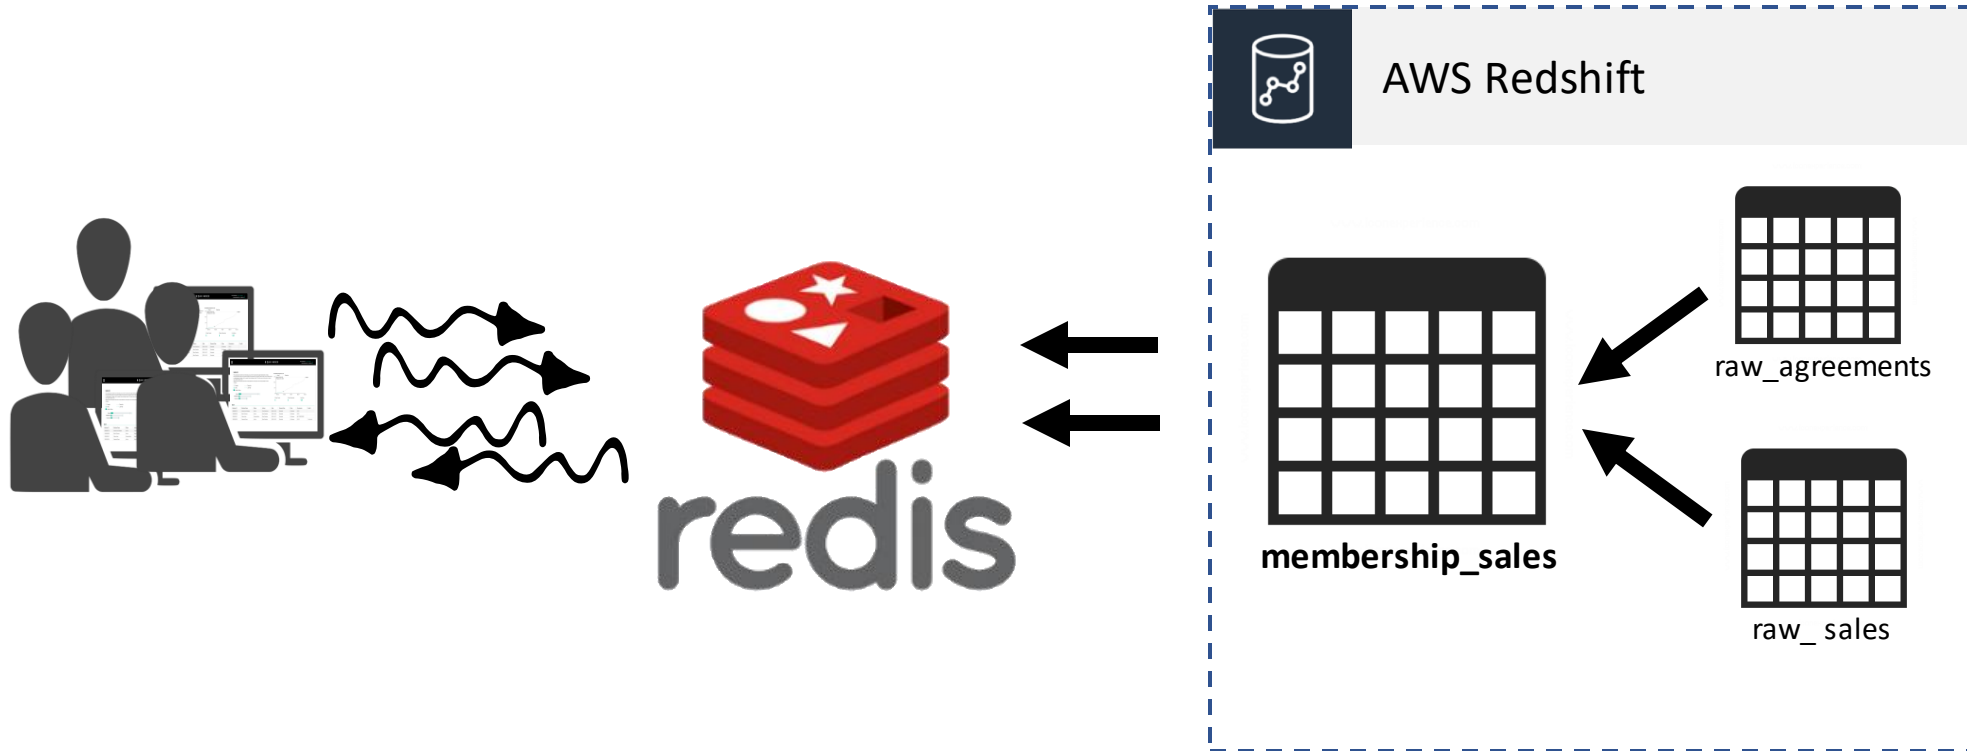
- Source DB is OLTP

Cons

- Sales logic is complex!
- Burdens prod DB

2. Cache in Redis

Sales Reporting Architecture



Potential Solutions

1. Pull from source DB

Pros

- Source DB is OLTP

Cons

- Sales logic is complex!
- Burdens prod DB

2. Cache in Redis

Pros

- Very fast performance

Cons

- Non-relational data structure
- Keys must be created for every data view

Potential Solutions

1. Pull from source DB

Pros

- Source DB is OLTP

Cons

- Sales logic is complex!
- Burdens prod DB

2. Cache in Redis

Pros

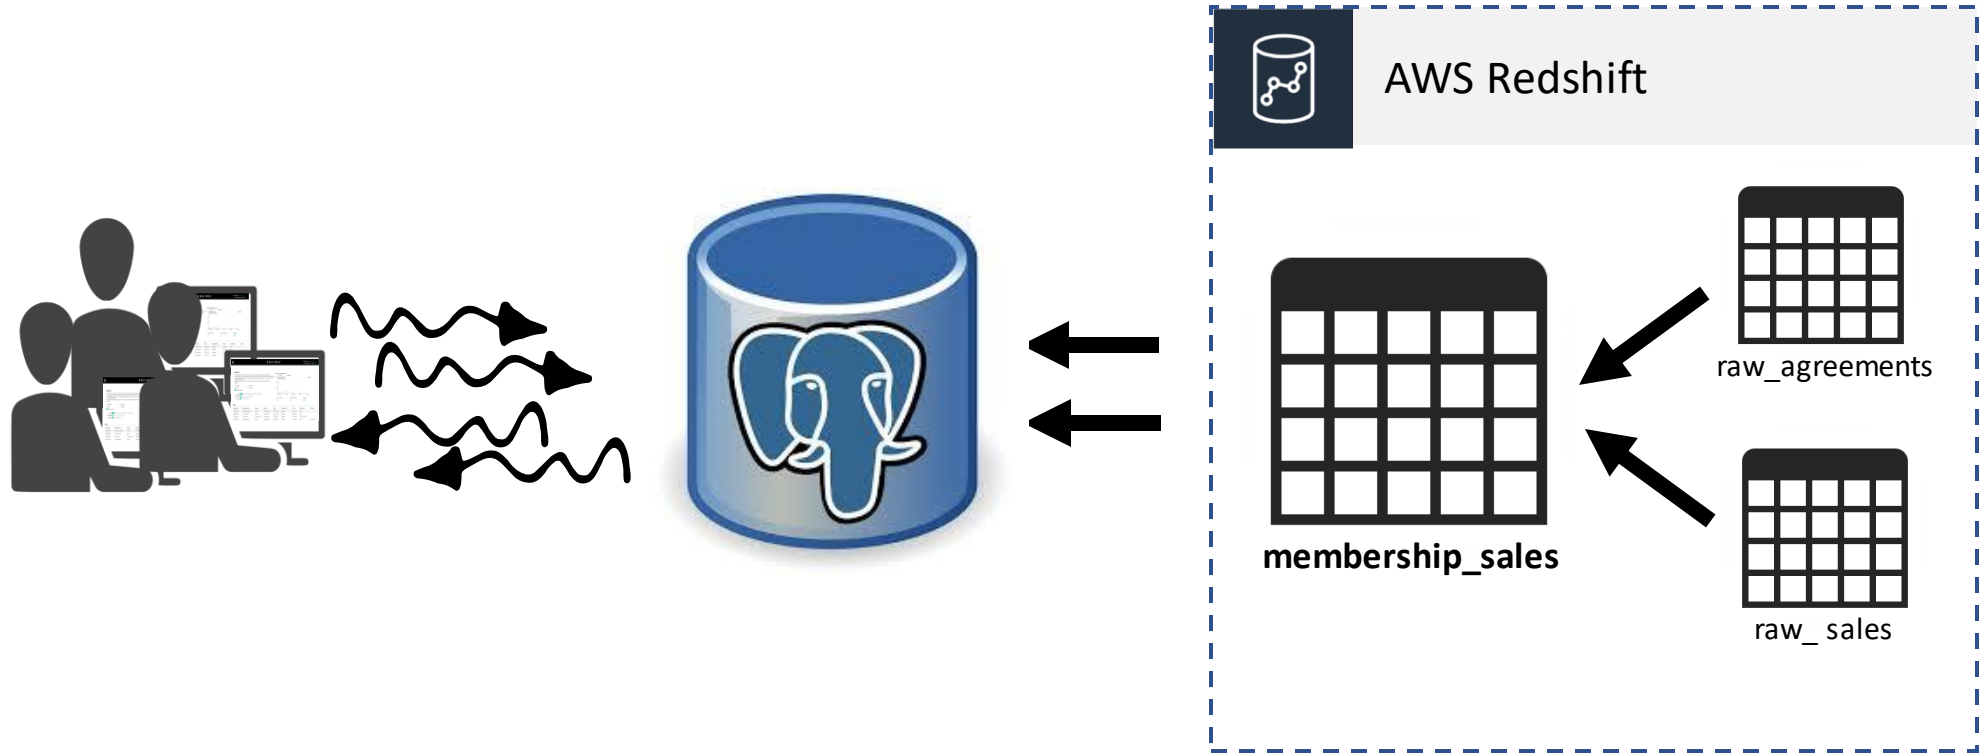
- Very fast performance

Cons

- Non-relational data structure
- Keys must be created for every data view

3. Copy to another DB

Sales Reporting Architecture



Potential Solutions

1. Pull from source DB

Pros

- Source DB is OLTP

Cons

- Sales logic is complex!
- Burdens prod DB

2. Cache in Redis

Pros

- Very fast performance

Cons

- Non-relational data structure
- Keys must be created for every data view

3. Copy to another DB

Pros

- Maintain relational format

Cons

- Additional ETL step

Potential Solutions

1. Pull from source DB

Pros

- Source DB is OLTP

Cons

- Sales logic is complex!
- Burdens prod DB

2. Cache in Redis

Pros

- Very fast performance

Cons

- Non-relational data structure
- Keys must be created for every data view

3. Copy to another DB

Pros

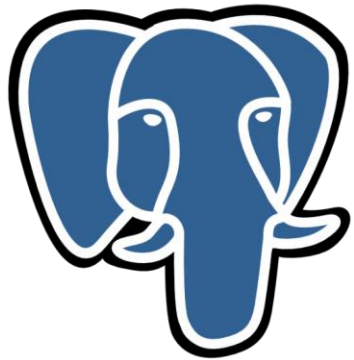
- Maintain relational format

Cons

- Additional ETL step



PostgreSQL Foreign Data Wrapper

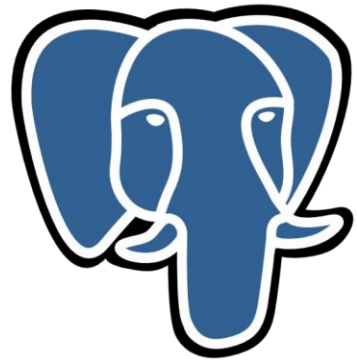


PostgreSQL



amazon
REDSHIFT

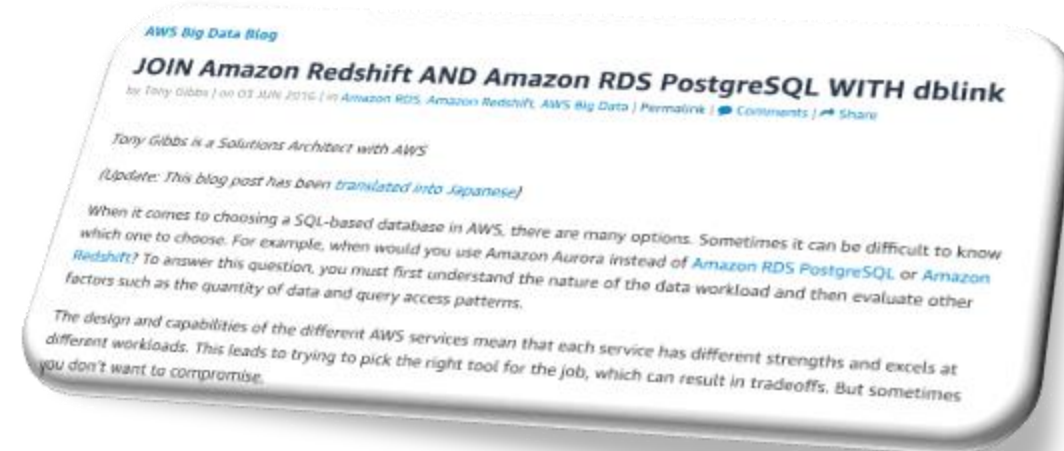
PostgreSQL Foreign Data Wrapper



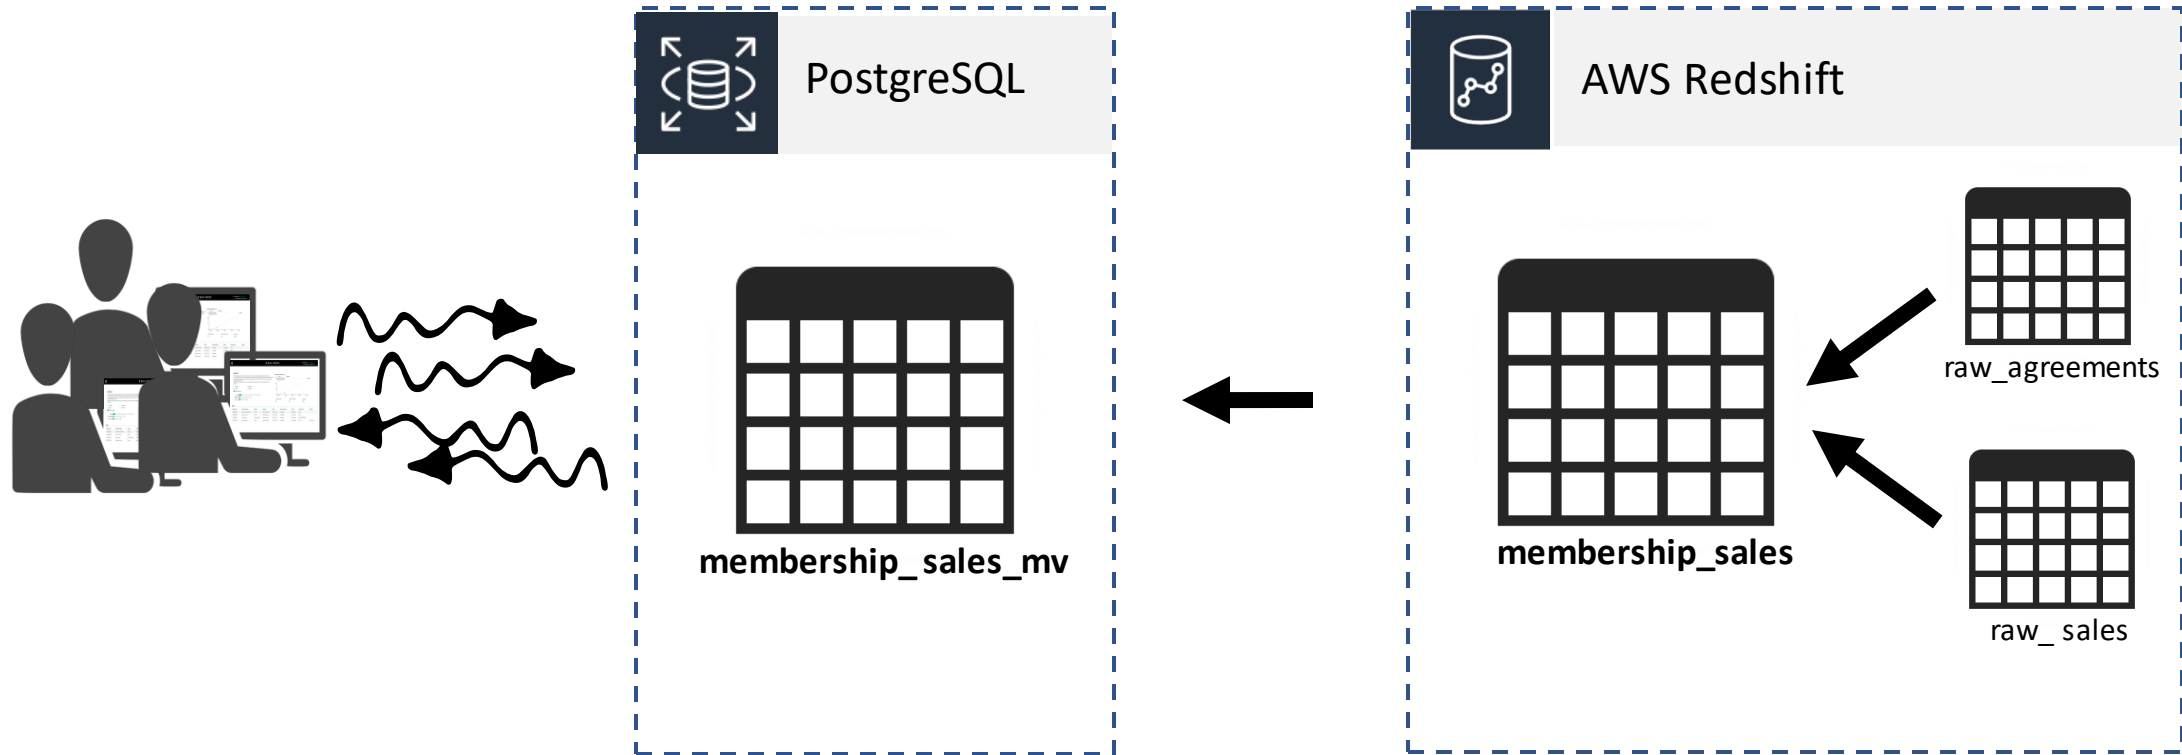
PostgreSQL



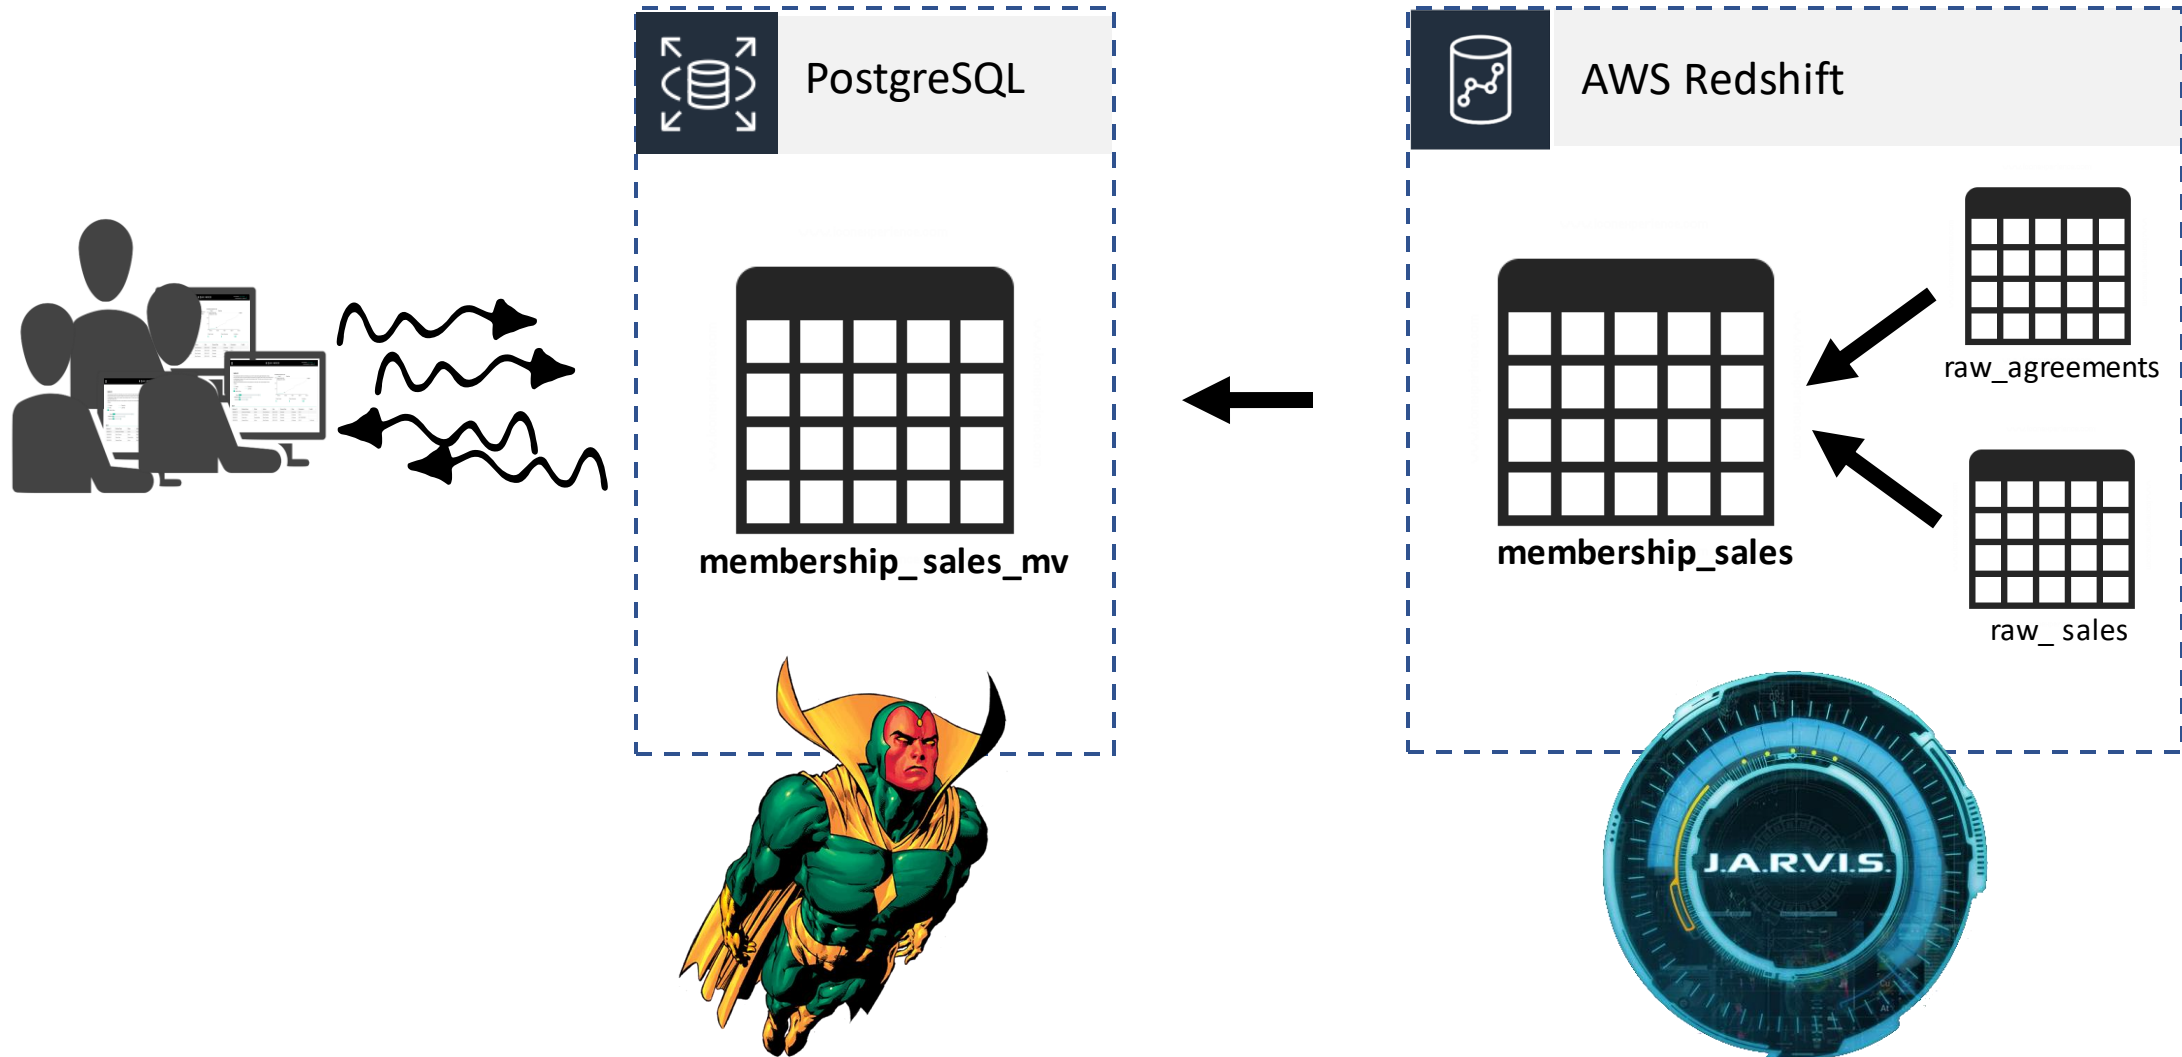
amazon
REDSHIFT



Proposed Sales Reporting Architecture



Proposed Sales Reporting Architecture



Environment Set-up

- ✓ Create Redshift cluster
- ✓ Create PostgreSQL server (9.5+)
 - RDS recommended
 - For self-managed, install Postgres contrib package:
 - `sudo yum install postgresql10-contrib.x86_64`
- ✓ Networking (AWS)
 - Co-locate in same Availability Zone
 - Configure Security Group

Creating the Link

--1 enable the required extensions

```
CREATE EXTENSION postgres_fdw;  
CREATE EXTENSION dblink;
```

--2 create the external server

```
CREATE SERVER jarvis  
    FOREIGN DATA WRAPPER postgres_fdw  
    OPTIONS (host 'REDSHIFT_ENDPOINT', port '5439',  
            dbname 'REDSHIFT_DB_NAME', sslmode 'require');
```

--3 save redshift login to this external server

```
CREATE USER MAPPING FOR PG_USERNAME  
    SERVER Jarvis OPTIONS  
    (user 'RS_USERNAME', password 'RS_PASSWORD');
```

Running queries on PostgreSQL

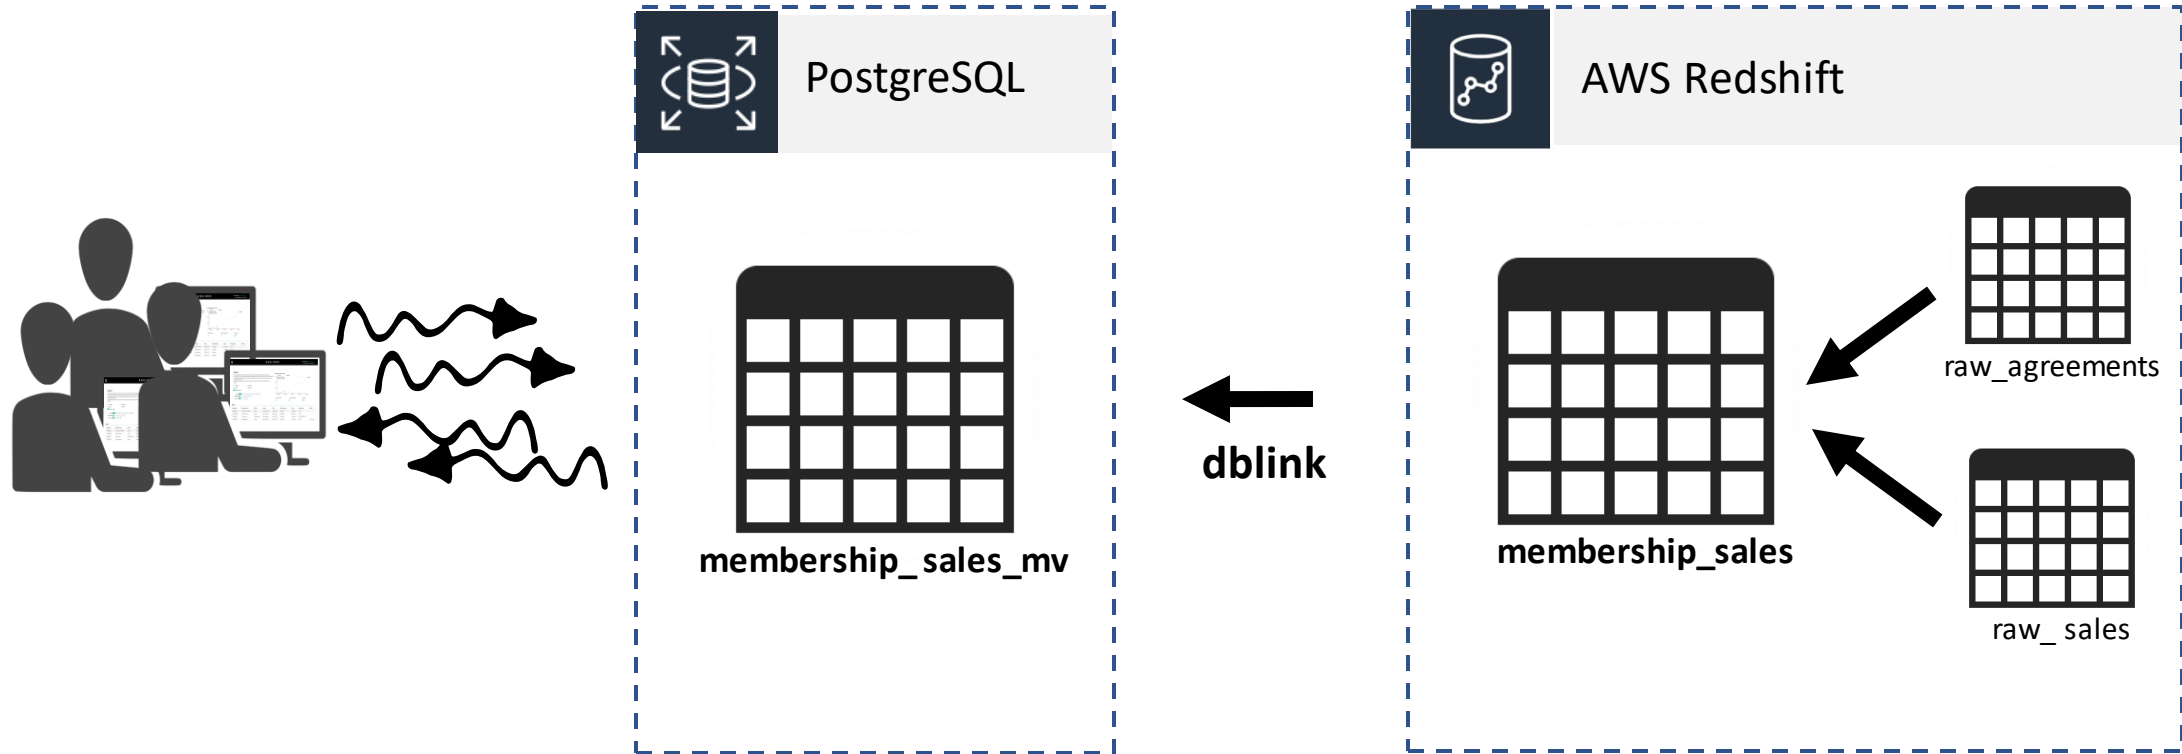
```
SELECT *  
FROM dblink('jarvis', $REDSHIFT$  
    SELECT  
        member_sales_id,  
        member_id,  
        sales_action,  
        sales_action_date  
    FROM  
        rs_landing.raw_sales $REDSHIFT$)  
AS sales_actions (  
    member_sales_id varchar(50),  
    member_id varchar(50),  
    sales_action varchar(50),  
    sales_action_date date  
);
```

Leveraging a Materialized View

```
CREATE MATERIALIZED VIEW pg.membership_sales_copy AS
(
  SELECT *
  FROM dblink('jarvis', $REDSHIFT$
  SELECT
    member_sales_id,
    member_id,
    sales_action,
    sales_action_date
  FROM
    rs.membership_sales $REDSHIFT$
) AS membership_sales_copy (
  member_sales_id varchar(50),
  member_id varchar(50),
  sales_action varchar(50),
  sales_action_date date
);

REFRESH materialized VIEW pg.membership_sales_copy;
```

Sales Reporting Architecture



Materialized View Roadblock



PostgreSQL

membership_sales_mv

contract_id	batch_id	timestamp
1111	101	2019-03-10
2222	101	2019-03-10
3333	101	2019-03-10
4444	101	2019-03-10
...
5555	102	2019-03-21
6666	102	2019-03-21



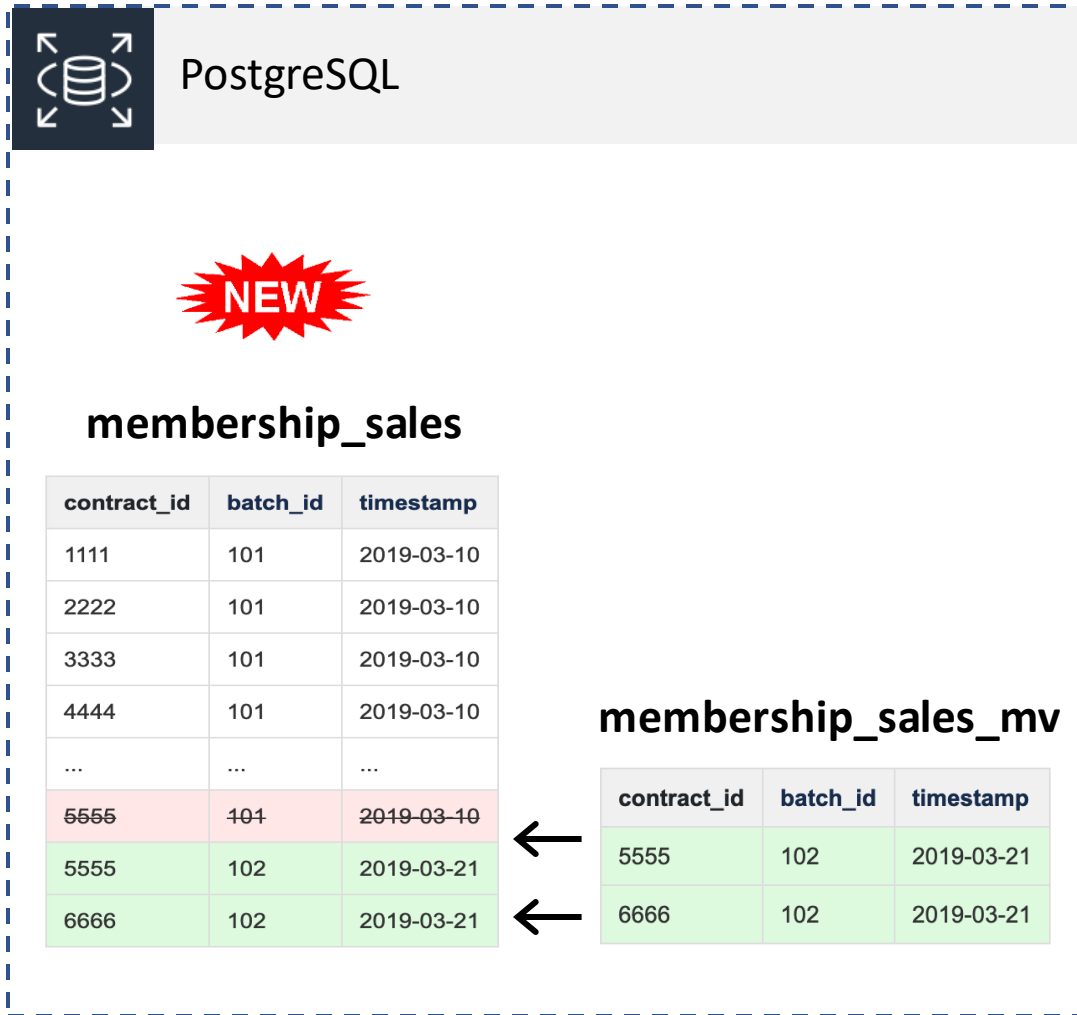
AWS Redshift

membership_sales

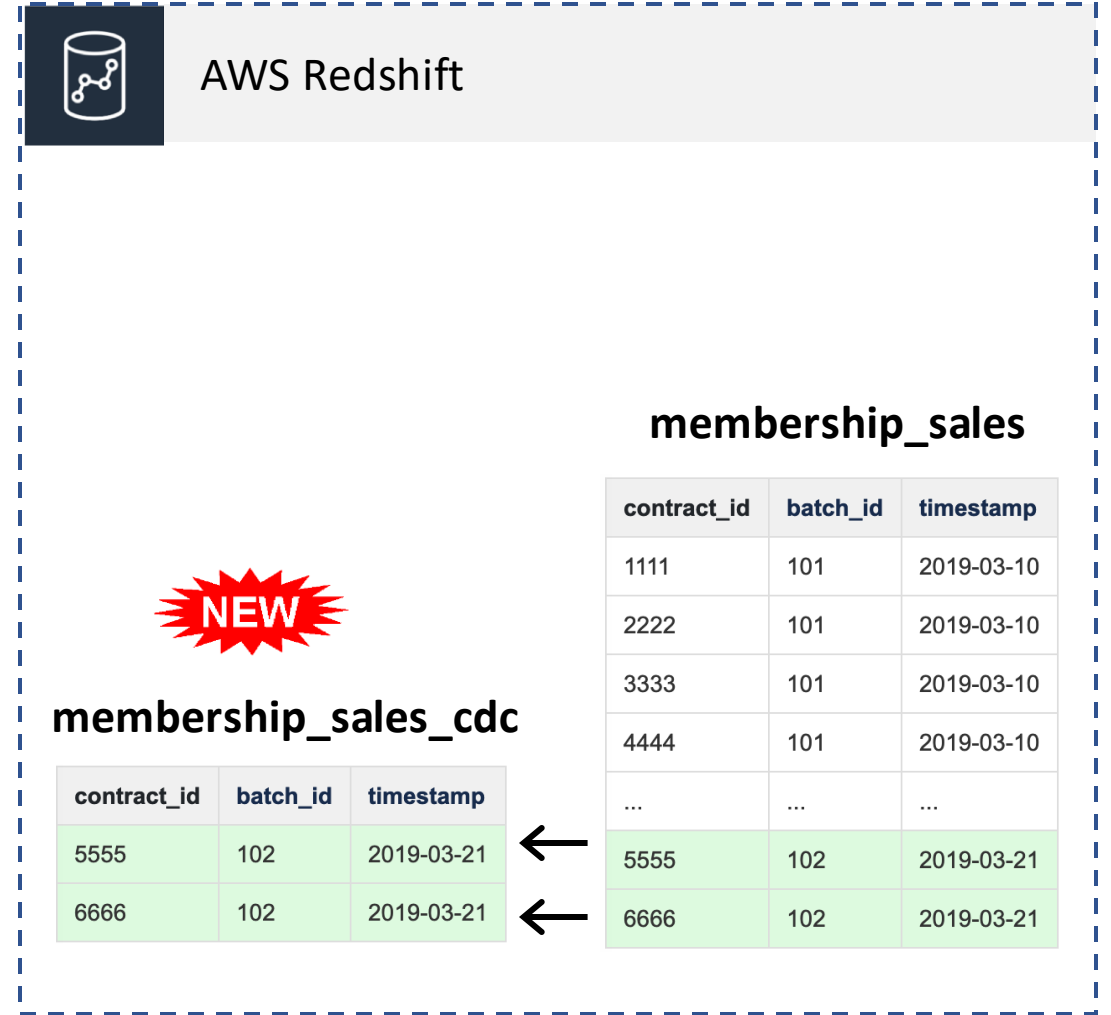
contract_id	batch_id	timestamp
1111	101	2019-03-10
2222	101	2019-03-10
3333	101	2019-03-10
4444	101	2019-03-10
...
5555	102	2019-03-21
6666	102	2019-03-21

Millions
Of records
←
~10 mins

Change Data Capture for Large Tables



←
dblink



--Step 1: Create staging table in Redshift with last few hours of sales actions

```
--CREATE TABLE rs_landing.stage_sales_action
```

```
DELETE FROM rs.membership_sales_cdc
```

```
INSERT INTO rs.membership_sales_cdc
```

```
SELECT member_sales_id, member_id, sales_action, sales_action_date
```

```
FROM rs.membership_sales
```

```
WHERE date >= ' ${?from_date}';
```

--Step 2: Refresh materialized view in Postgres

```
REFRESH materialized VIEW pg.membership_sales_mv;
```

--Step 3: Upsert logic to populate final table in Postgres from materialized view

--temp table to hold last batch

```
DROP TABLE IF EXISTS cdc_sales;
```

```
CREATE TEMP TABLE cdc_sales AS
```

```
SELECT * FROM pg.membership_sales_mv;
```

--update changed records, member_sales_id as the key to identify a unique record

```
UPDATE pg.membership_sales ms
```

```
SET sa.member_id = s.member_id,
```

```
    ms.sales_action = s.sales_action,
```

```
    ms.sales_action_date = s.sales_action_date
```

```
FROM cdc_sales s
```

```
WHERE s.member_sales_id = ms.member_sales_id;
```

--delete the records we just updated from temp table

```
DELETE FROM cdc_sales s USING pg.membership_sales ms
```

```
WHERE s.member_sales_id = ms.member_sales_id;
```

--insert new records not found in membership_sales

```
INSERT INTO pg.membership_sales
```

```
SELECT * FROM cdc_sales;
```

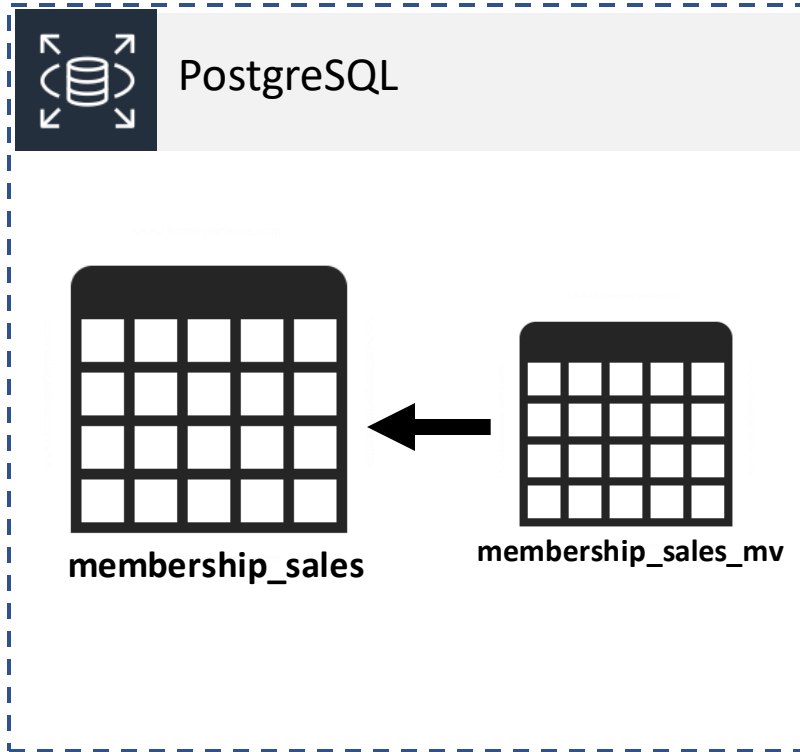
--drop temp table

```
DROP TABLE cdc_sales;
```

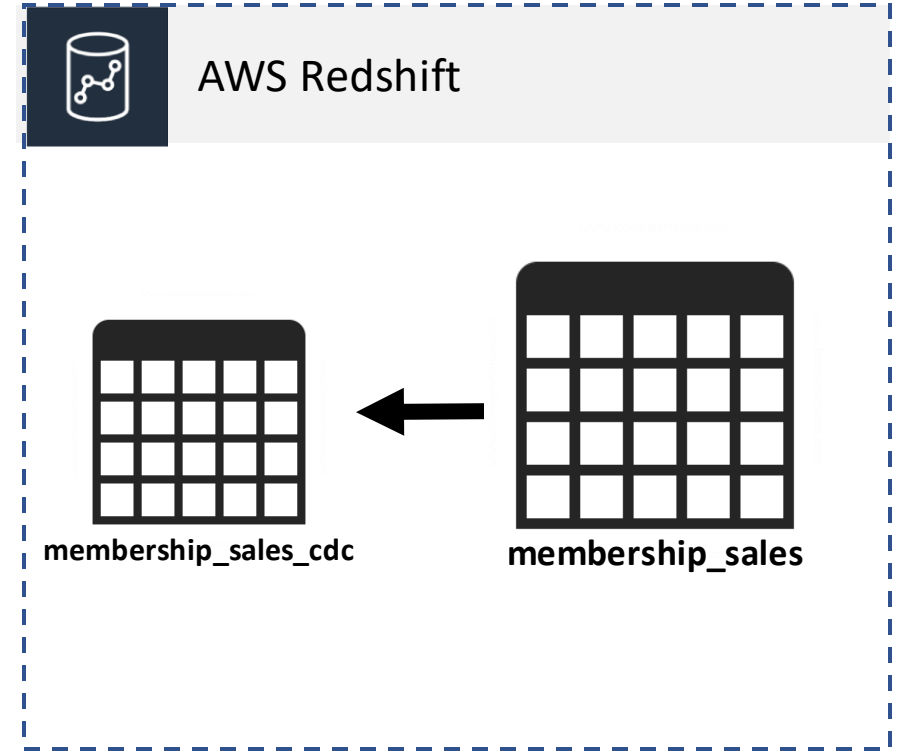
Redshift

PostgreSQL

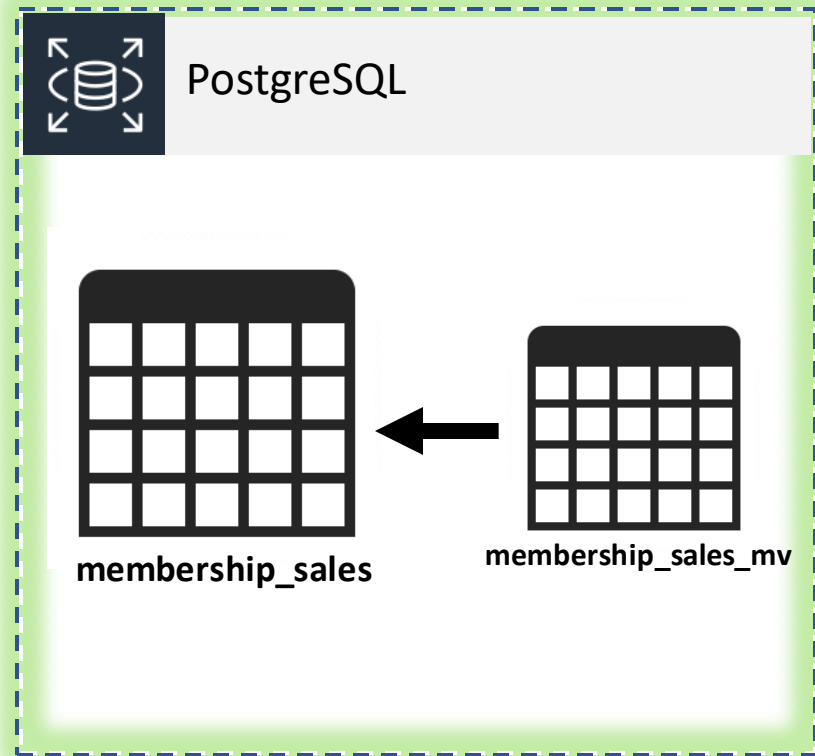
Sales Reporting Architecture



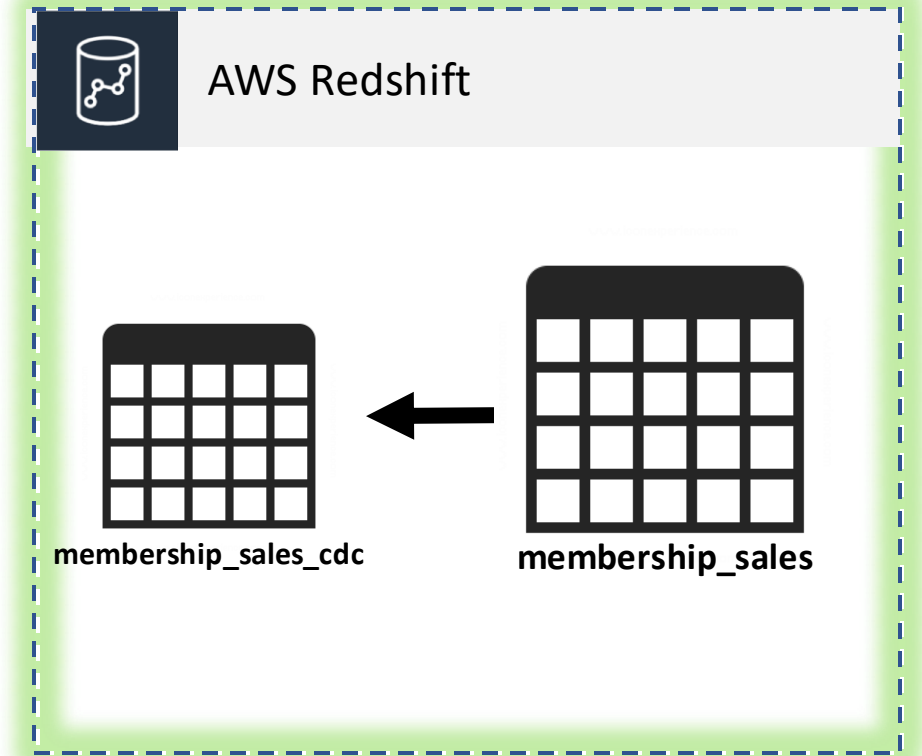
←
dblink



Sales Reporting Architecture **Solution**



← dblink



Learnings & Notes

- Minimal maintenance on Postgres instance
- Won't reflect source deletions
- Limited to a few tables
- Flexible for schema evolution



Looking to the Future...

- Make use of foreign table
- Front-end scaling with read-replicas
- Extensible to other datastores
- Event-based streaming architecture

More to Explore

Appendix F. Additional Supplied Modules

Table of Contents

- F.1. adminpack
- F.2. amcheck
 - F.2.1. Functions
 - F.2.2. Using amcheck effectively
 - F.2.3. Repairing corruption
- F.3. auth_delay
 - F.3.1. Configuration Parameters
 - F.3.2. Author
- F.4. auto_explain
 - F.4.1. Configuration Parameters
 - F.4.2. Example
 - F.4.3. Author
- F.5. bloom
 - F.5.1. Parameters
 - F.5.2. Examples
 - F.5.3. Operator Class Interface
 - F.5.4. Limitations
 - F.5.5. Authors
- F.6. btree_gin
 - F.6.1. Example Usage
 - F.6.2. Authors
- F.7. btree_gist
 - F.7.1. Example Usage
 - F.7.2. Authors
- F.8. chkpass
 - F.8.1. Author
- F.9. citext

Generic SQL Database Wrappers

Data Source	Type	License	Code	Install	Doc	
ODBC	Native		github			CartoD
JDBC	Native		github			
JDBC2	Native		github			
SQL_Alchemy	Multicorn	PostgreSQL	GitHub	PGXN	documentation	Can be used to acc
VirtDB	Native	GPL	GitHub			A generic

Specific SQL Database Wrappers

Data Source	Type	License	Code	Install	Doc	
PostgreSQL	Native	PostgreSQL	git.postgresql.org		documentation	
Oracle	Native	PostgreSQL	github	PGXN	website	
MySQL	Native		github	PGXN	example	
Informix	Native	PostgreSQL	github			
Firebird	Native		github	PGXN		
SQLite	Native		github			
SQLite	Native	PostgreSQL	github	PGXN	README	An FDV
Sybase / MS SQL Server	Native		github	PGXN		
MonetDB	Native		github			

NoSQL Database Wrappers

Data Source	Type	License	Code	Install	Doc
BigTable or HBase	Native Rust Binding (RPGFFI)	MIT	Github		
Cassandra	Multicorn	MIT	Github	Rankactive	
Cassandra2	Native	MIT	Github		
Cassandra	Multicorn	PostgreSQL	Github		
ClickHouse	Multicorn	BSD	Github		READM

Q&A

We Are Hiring

Email ITCareers@Equinox.com

- Head of Engineering
- React Native Engineer
- Sr. React Native Engineer
- API Engineer
- SDET Java - Architect
- SDET Javascript - Architect
- SDET Java
- SDET Javascript
- Sr. UX Researcher

Equinox Tech Blog <http://tech.equinox.com/>