

Postgres as BI platform

Andy Fefelov
mastery.pro

Agenda

- Problem statement
- Open source solution
- ROLAP
- Our architecture review
- Postgres features suitable for BI
 - ETL vs ELT (stage-nds-ddm)
 - Column data storage
 - Configuration
 - Special features
- Pros and cons of our solution

Problem statement

- Our customer is one of the largest pharmacy supply chain group in Ireland
- 4 types of dispensary software
- 250 pharmacies
- To be analyzed:
 - Orders
 - Scripts (prescription, recipe)
 - Claims
- Goals to be achieved:
 - Purchasing policy optimization
 - Marketing killing feature

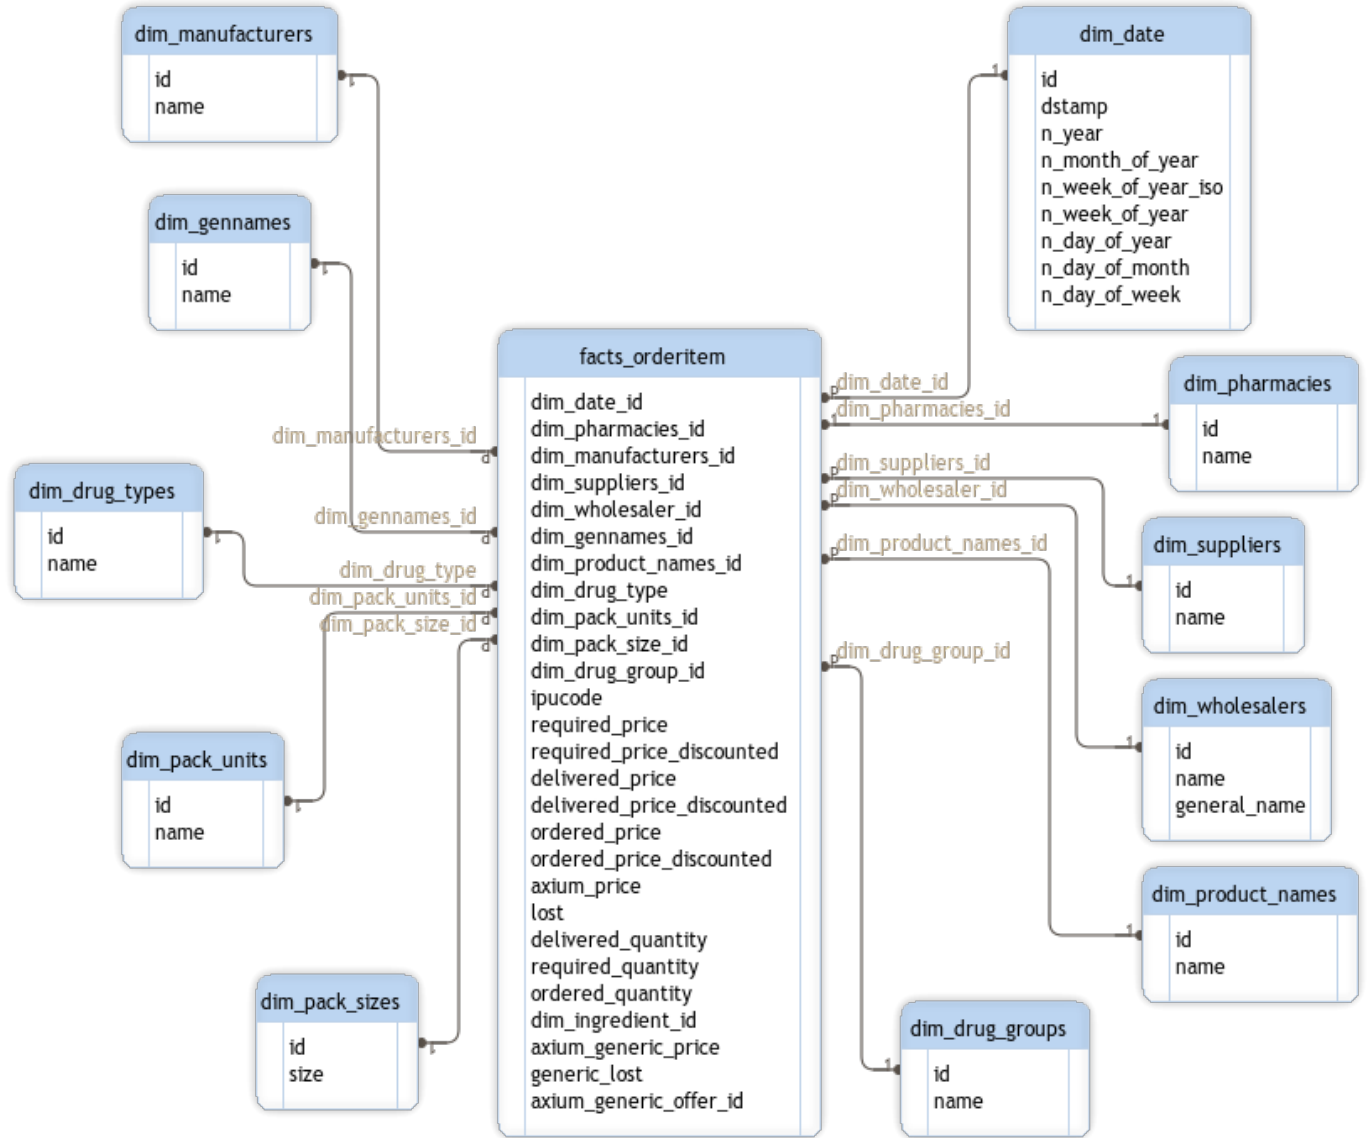
Open source

- SpagoBI
- Pentaho
- Mondrian
- Saiku
- Cubes (databrewery)

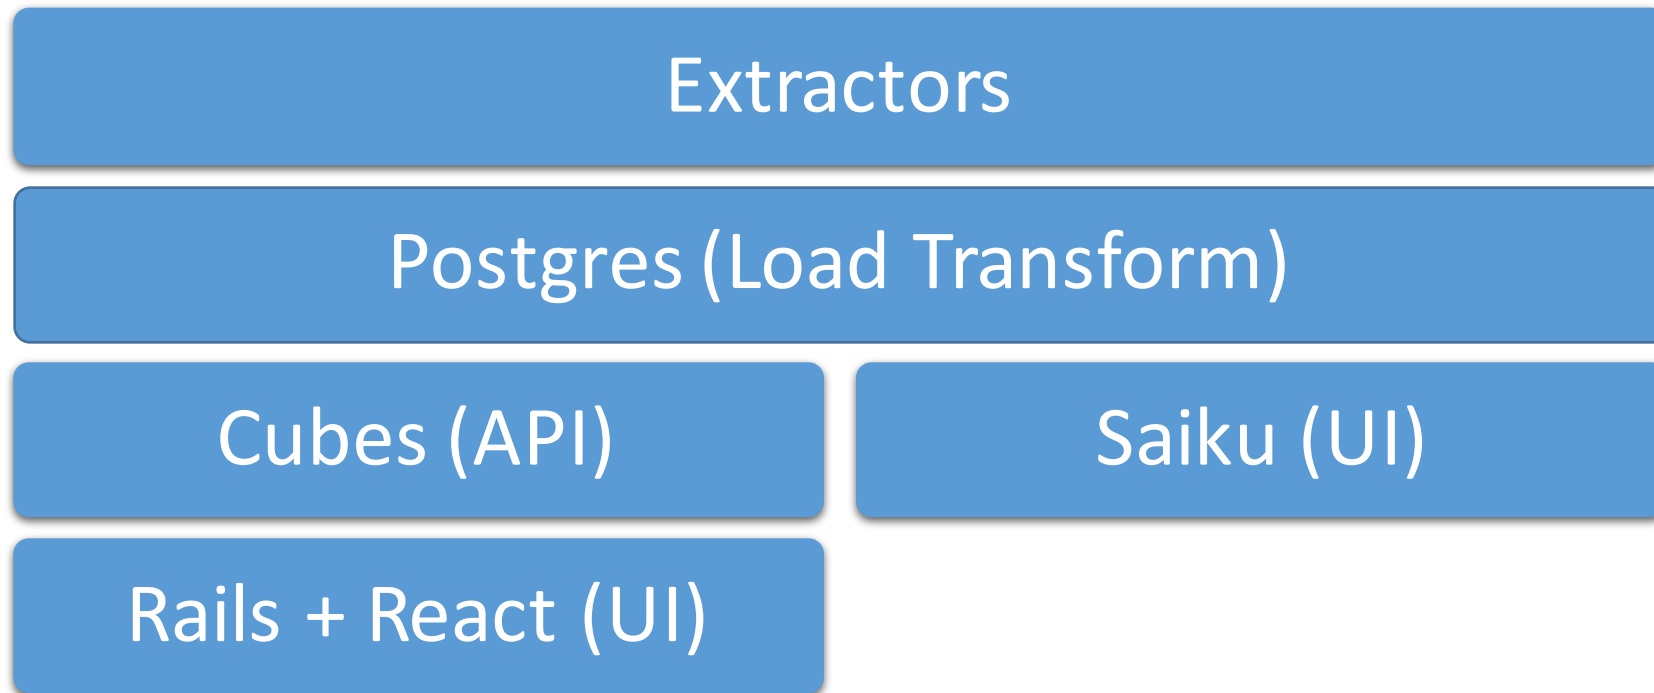
ROLAP (R-ROLAP)

- Star scheme
 - Facts
 - Dimensions
 - Measures
- No pre-calculated aggregates
- SSD
- Column storage
- ???
- Profit!

ROLAP




Our architecture




Architecture - extractors

- Cyclone_client
 - Mssql (2008-2012)
 - Golang
 - CSV + rsync over ssh
- Kachok
 - Web scrapper
- Skytools replication
 - From existing products



Architecture – API + UI

- Cubes - cubes.databrewery.org
 - Easy drilling-down
 - Slicing and dicing
 - Serves aggregates, dimension details, facts
 - Provides all necessary metadata for a reporting application
- Rails, React
 - Authorization
 - d3, dc, crossfilter
- Saiku
 - Only for back office



Architecture – Postgres (load, transform)

stage

- raw data
- load_something_to_nds(_pharmacy_id **integer**)

nds

- normalized data store
- load_something_to_ddm(_pharmacy_id **integer**)

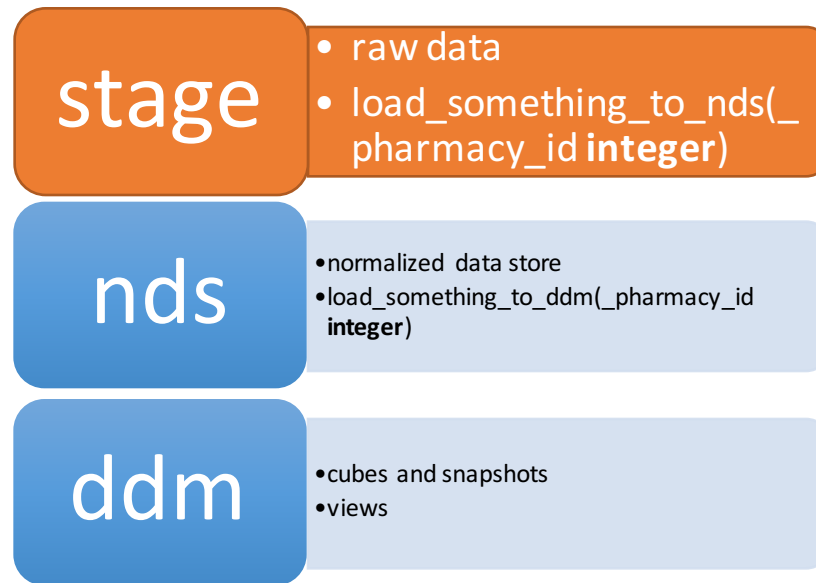
ddm

- cubes and snapshots
- views

Architecture– Postgres (load, transform)

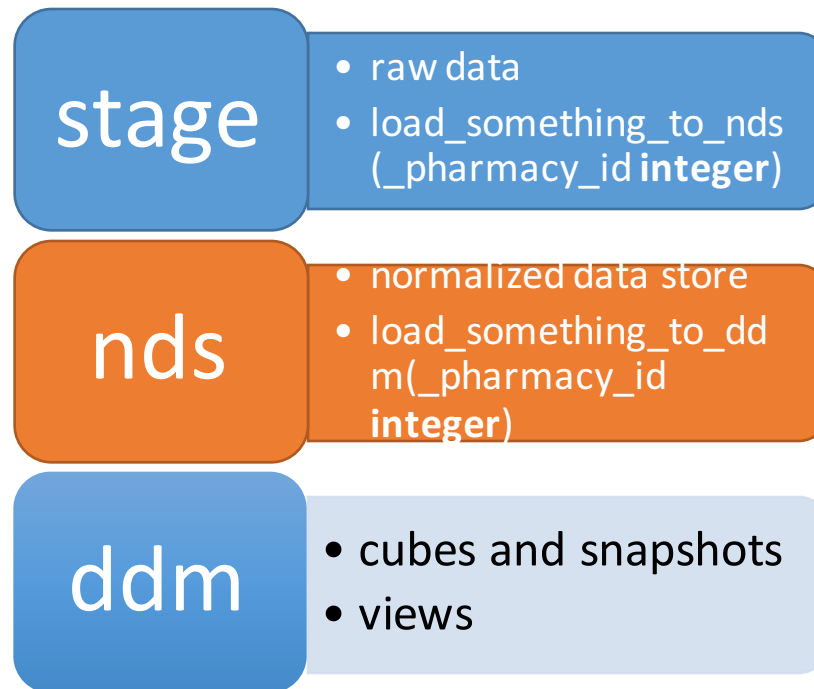
Stage

- «Raw» data
- Cleaned up completely in every ELT cycle
- Is as data source for NDS



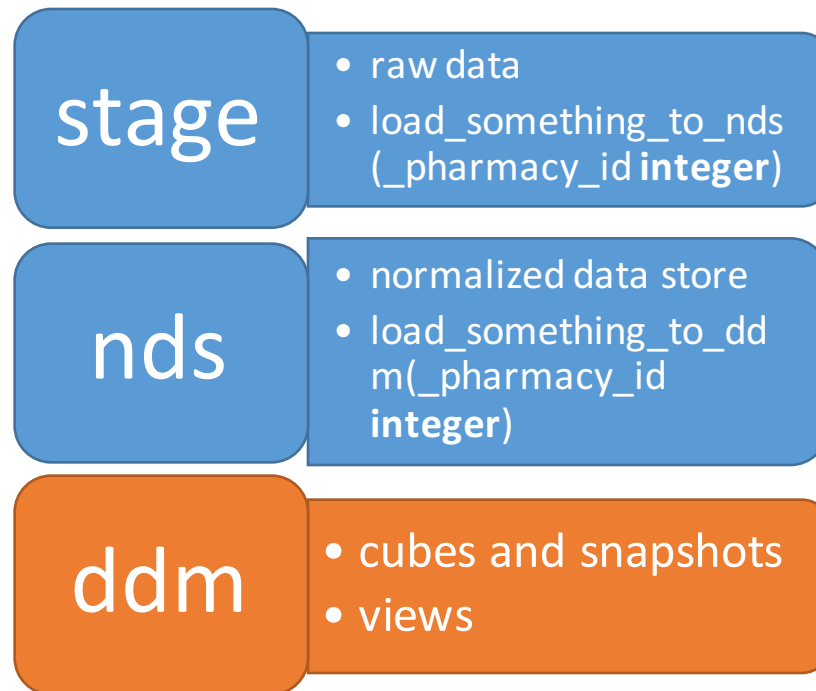
Architecture– Postgres (load, transform)

- Normalized Data Store
 - Here data is normalizing and validating
 - Is a source for ddm
 - Measures for ddm is calculated there
 - delta calculating for loading into ddm based on last_updated field

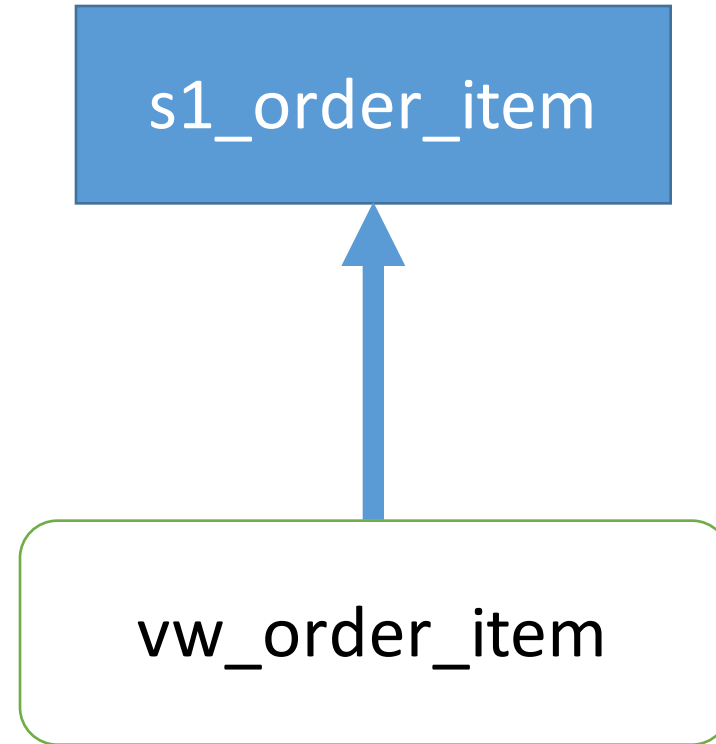



Architecture – Postgres (load, transform)


- Dimensional data model
 - Cubes
 - Snapshots
 - Deploy calmly
 - Analyze before-after release states
 - View is entry point for application



Architecture – Postgres (snapshots)



Architecture – Postgres (snapshots)



Column storage

- Suitable for:
 - aggregations
 - showing fixed numbers of columns
- `cstore_fdw` -> https://github.com/citusdata/cstore_fdw
 - Compression: Reduces in-memory and on-disk data size by 2-4x. Can be extended to support different codecs.
 - Column projections: Only reads column data relevant to the query. Improves performance for I/O bound queries.
 - Skip indexes: Stores min/max statistics for row groups, and uses them to skip over unrelated rows.

Column storage

- Our experience:
 - Is not faster than vanilla postgres (say hello to cubes)
 - Volume reduced up to 12 times. Wow.
 - No way to backup traditional way (no need?)
 - No support for delete/update (snapshots)

Configuration

- Load profile:
 - Big volume RW I/O
 - Most of I/O is sitting in stage, nds
 - ddm is not high loaded
- shared_buffers = ½ RAM
- work_mem = 2GB
- maintenance_work_mem = 3GB
- temp_buffers = 2GB
- effective_cache_size = ½ RAM
- max_wal_size = 32GB

Features

- DDM could be placed in dedicated server (londiste, pg_logical)
- Use COPY/BULK INSERTS, don't use UPDATE (ke ke ke)
- You should think about horizontal and vertical partitioning, please find proper keys for that
- You should think about parallelism from very beginning
- Use TABLESPACES/PARTIAL INDEXES (and more and more disks)
- You should use data store policy
- Statistics should be collected in tempfs volume

Features vol 2

- Use migrations – sqitch by theory
- You'd better test ELT - sqitch by theory
- Use pg_stat_statements (add this into monitoring)
- Use profiling – PLPROFILER3
- Sometimes, you have (not) to use cstore_fdw
- Sometimes, you have (not) to use unlogged tables

Pros and cons

- Cons

- No easy way to scale horizontally
- Reasonable difficult deploy

- Pros

- Local data (no big network transfers)
- Effectively parallelized (thanks to pharmacy_id)
- **PL/pgSQL**

Thank you

andy@mastery.pro

Speed limit

- cubes is not fast (due to serialization)
 - json (12 sec)
 - ujson (4 sec)
 - postgres json output (1.5 sec) db self time 0.3-0.7 sec