

A photograph of three business professionals (two women and one man) sitting around a table in a meeting, engaged in conversation. The image is overlaid with a blue tint. The text 'Data for Breakfast' is written in a white, cursive font across the center of the image.


Data for Breakfast





Kalibri Labs

Kalibri Labs - Dev Team



Nick Donovan

- Lead Software Engineer
 - Boeing, NSA
 - Near real-time, high data volumes
 - Big data



Martin Farmer

- VP, Technology
 - GEC, AT&T, Lucent Technologies, GCHQ, NSA, Tensley Consulting
 - Near real-time, high data volumes
 - Big data



About Kalibri Labs



“Kalibri Labs evaluates and predicts revenue performance in the digital marketplace”


*We are the only company with the booking details from every day of every stay from the guest's folio. Our proprietary **Optimal Business Mix algorithm** uncovers the best a hotel can achieve and delivers an action plan to improve profit contribution.*

- 45 employees, 8 devs, 2 devops
- Over 33,000 hotels
- 8.2B rows of guest stay data (no PII)
- ~100M transactions per month
- Weekly updates from all clients




Our Journey

- Unique Selling Point (USP) of Snowflake over Redshift
- Proof of Concept (PoC)
- Transition to operations (TTO)
- Current operations



USP of Snowflake - Before and After!

ETL BEFORE



Previous bottlenecks


DB admin / periodic vacuuming

Downtime to scale up or down db resources

Larger queries would often fail

Multiple Databases

ETL AFTER



Snowflake value

No DBA required

Scale warehouse instantly

Time travel

Easy to clone database objects

USP Examples

- Scaling processing
 - from UI, JDBC or SQL

- Time Travel
 - absolute/relative time
 - statement id

- Clone Objects
 - database/tables

Context: ACCOUNTADMIN REGULAR_USE_ME... (M) OBM PUBLIC

Role: ACCOUNTADMIN Change

Warehouse: REGULAR_USE_MEDIUM (M) On

Database: REGULAR_USE_MEDIUM (M) On

BATCH_REMOVAL (M) Suspended

```
alter user mfarmer set DEFAULT_WAREHOUSE=regular_use_medium;
```

```
42
43 select * from kldw.public.f_stay at(offset => -60*5) limit 1;
```

Results Data Preview

✓ Query ID SQL 3.06s 1 rows

Filter result... [Download] [Copy]

Row	UID	KL_CO_CODE	CHECKOUT	CONFIRMATION	FOLIO	RESERVATION_I
1	3f9c5b285b...		36960734	36960734		1050

```
create sequence public.BATCH_ID_SEQUENCE clone devkldw.public.BATCH_ID_SEQUENCE;
drop table kldw.public.d.property;
create table kldw.public.d.property;
create table kldw.public.d.property;
create table kldw.public.d.property;
create table kldw.public.d.property;
```

Proof of Concept

Steps

Analysed representative queries

Wrote a simulator

Converted some java code to verify
JDBC access

With Snowflake

Medium warehouse comparable run times

Cost model similar

Able to stand up developer databases with clones

No DBA - Time travel very intuitive

PoC support excellent (migrated full database over)



Configurable
Warehouse



Production

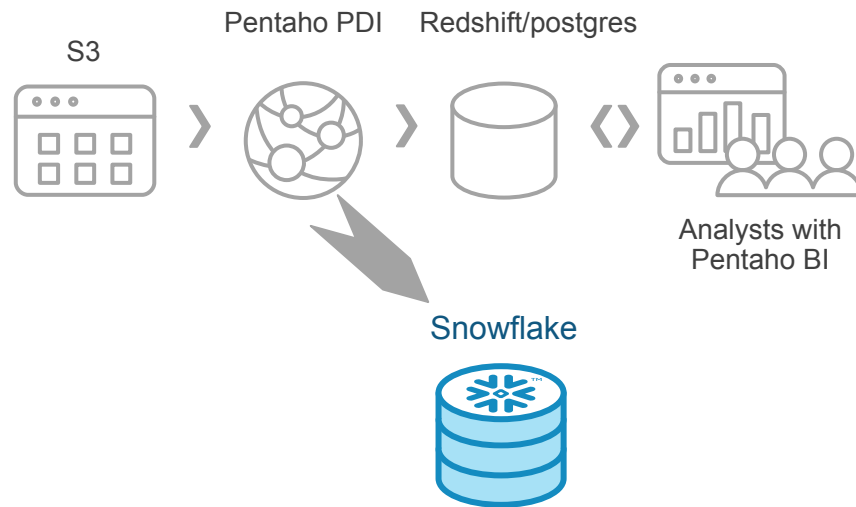


Developer
Clones




Transition to Operations

- Redshift and Snowflake in parallel
 - Snowpipe to sync data
- SnowSQL
 - Parameterized scripts
 - QA views
- External tables
 - Got on beta release
 - Used to mirror Redshift Spectrum tables



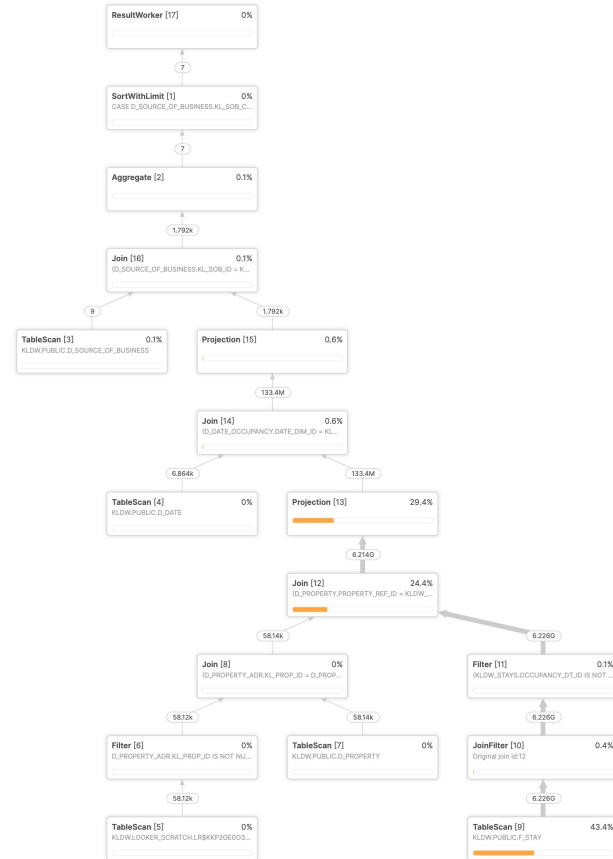
Additional Features

- Reader accounts
 - Creates a new URL for clients
 - Control shares from production
 - Control access and warehouse usage
 - Additional source of revenue
- Query profile
 - Identified a 12-hour query that we were able to optimize to 4 seconds




Monitoring

- Query profile
- Resource monitors
- SnowAlert
- Snowflake views



Upcoming Features

- Multi-cluster warehouses
 - Horizontal scaling
- Numeracy worksheets
 - Auto-complete
 - Sharing
 - Graphs
 - Reduce licenses for our BI tool

+  1 Shar

Updated 10m ag

Gross Merchandise Value ▾

Group by Month

```

5      sum(amount_returned) as total_returns,
6      sum(application_fee) as total_fees,
7      string_agg(orders.currency) as currency,
8      string_agg(orders.external_coupon_code) as coupons
9  from stripe.orders orders
10 join stripe.customers customers
11 on orders.|
12 where s
13 group b
14
15
16      amount
17      column integer
18
19      amount_returned
20      column integer
21
22      application_id
23      column string
24
25
26

```

Support



- Todd and team
- Community pages
 - Cases
 - Forums
 - New feature requests

A screenshot of the Snowflake Support page. At the top, a blue banner contains the text "From architecture concepts to expert tutorials - find the information you need to set up and use Snowflake" and a "GET STARTED" button. Below this is a white notification bar with a megaphone icon, the word "New", and the text "Snowflake Q&A forum is now on Stack Overflow" followed by a "GO TO STACK OVERFLOW »" link. The main content area consists of four white cards with rounded corners. The first card, "Q&A Forum", has a thumbs-up icon and text: "Join experts online to ask questions, discuss projects, and help one another". The second card, "User Groups", has a group of people icon and text: "Get together to talk about what we've built, check out new features, and learn from each other". The third card, "Documentation", has a document icon and text: "Access the latest product documentation to make the best out of Snowflake". The fourth card, "Support", has a hand holding a heart icon and text: "Need to contact Support? Log in to your Snowflake Community account".



Examples



Snowflake makes our business lean!

- Scalable warehouse
- Pushing work from Spark to Snowflake gave us performance gains



Snowflake saved our data!

- Time travel allowed us to recover production data without needing to go to backups
- Used for troubleshooting



We reduced a 6 hour Spark query to 22 minutes by pushing the query down to Snowflake natively.

Martin Farmer
VP, Technology

CLOSING THOUGHTS - Tech

- “It lives up to its promise”
 - No DBA
 - Time travel is awesome
 - Centralised our data
 - Rationalised 2 x Redshift and 2x postgres
 - Centralised our environment
 - Multiple dev and prod
- Flexibility
 - Reader accounts for client access
 - Scalable warehouses
 - Rich permissions model
 - SnowSQL

CLOSING THOUGHTS - Business

- Massively Streamlined Our Business
 - Multi-day jobs down to a few hours
- Used reader accounts to turn around new business opportunities
 - way quicker than before
 - can use production data that we choose to expose
- Higher Quality product delivered
 - nightly builds and QA
 - fresh clone of production database
 - run every client through dev code
 - twice a week
- Great Support
 - Sales and tech
 - Snow Lodge
- Bright Future
 - Premium to free
 - New features

THANK YOU

Data for Breakfast



ICONS



Snowflake DB



Data Warehouse



Data Lake



Datamart



BI Analytics



Server



Enterprise



Optimized Structured Data



Structured Data



S3 Data Protection



Alert



Process



User / Customer



Web App



Tools



Third Party



World / Global



Azure Blob



Check



GCS



Desktop



Mobile



Application



Document



IoT



Data Science



Cloud



Optimization



Metadata



Transactions



Refresh



Services



Metadata / Tag



Security



Results



Third Party Apps



Access Control



Easy Management



Dev



Communicate



Results



No



Edit



Optimize

