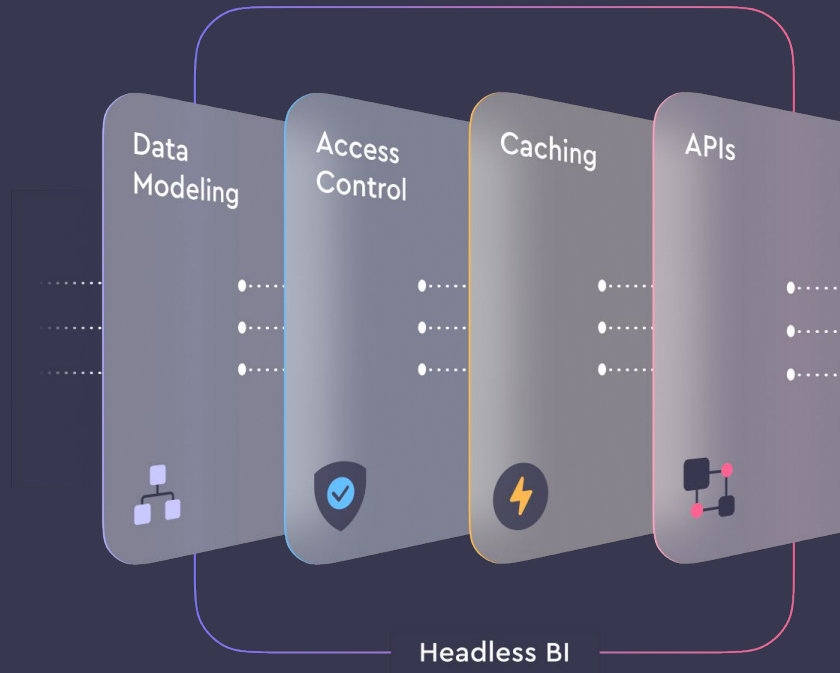


# Navigating the Modern Data Stack with Open-Source Headless BI



**Artyom Keydunov**

Co-founder & CEO at Cube



# What is headless BI?

# Basecase VC was the first to publicly talk about 'Headless BI' in early 2021

## Headless Business Intelligence

Ankur Goyal, Alana Anderson


January 7, 2021

*8 minute read*

It's 2021, and everyone is excited about the post-[Snowflake](#) world (worth just shy of \$80B at time of writing). At Base Case, we are bullish about the next generation of software companies built on top of cloud-based data warehouses like Snowflake. Specifically, we believe there's an open opportunity to solve a critical problem that all data-driven businesses face: calculating metrics consistently. Metrics like daily active users, funnel events, and churn signals are critical to scaling a product-led growth motion, but are still too difficult to shuttle into the core workflows that product and go-to-market teams engage in everyday.

The category of products that should be addressing this is business intelligence (BI) tools; however, by


# Proliferation of data applications



More data source

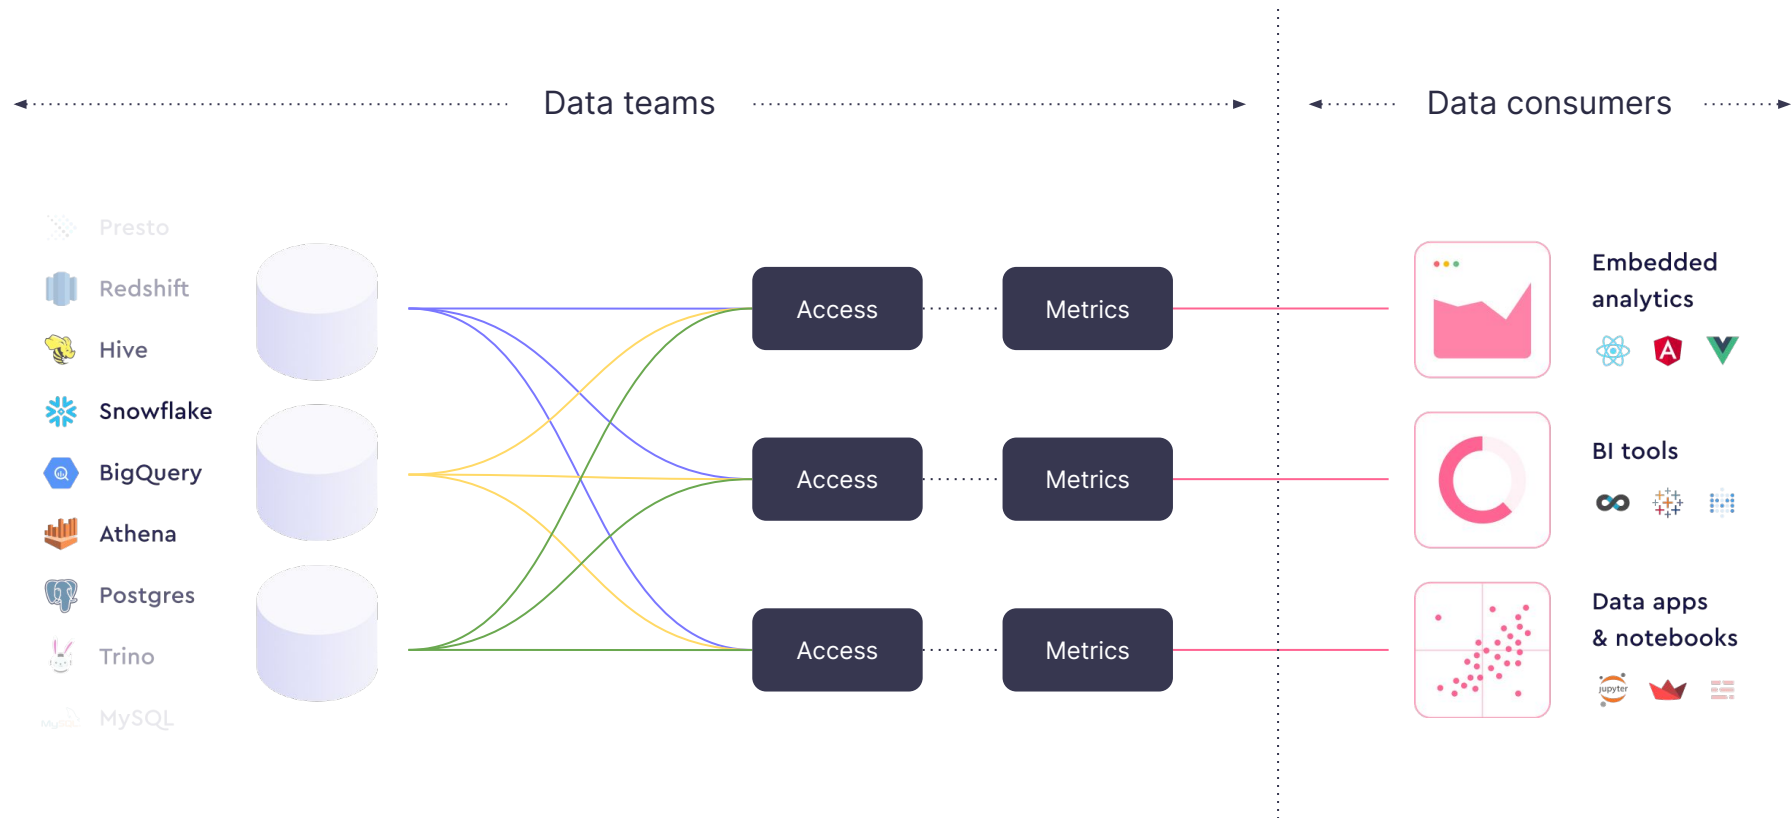


More data consumers

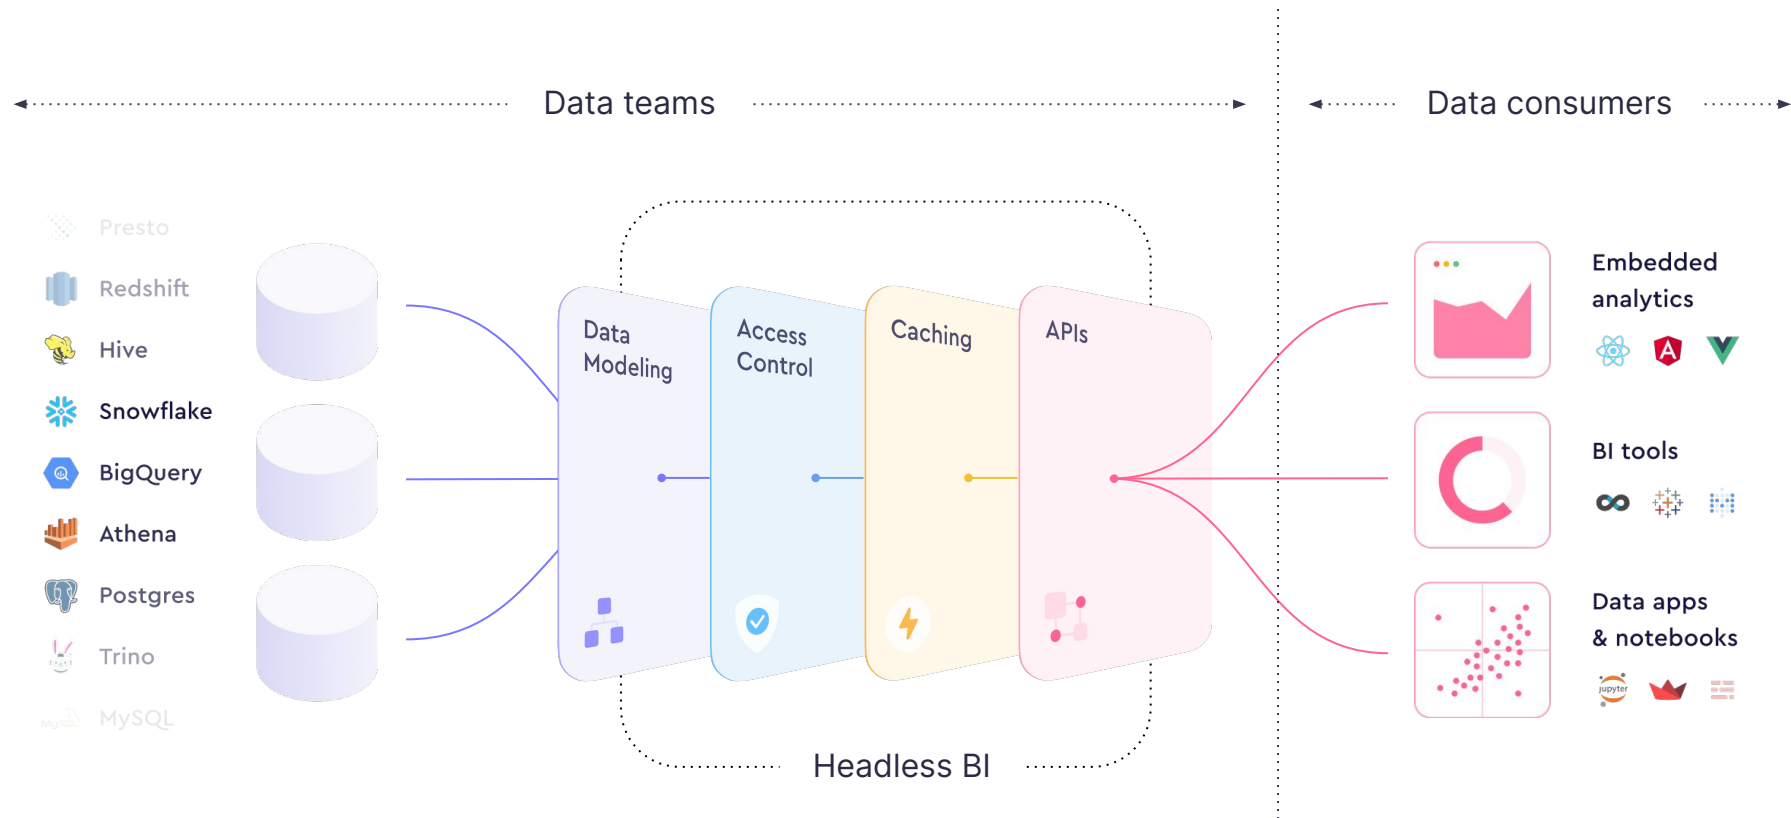


Proliferation of role-specific data applications

# Very *un*DRY

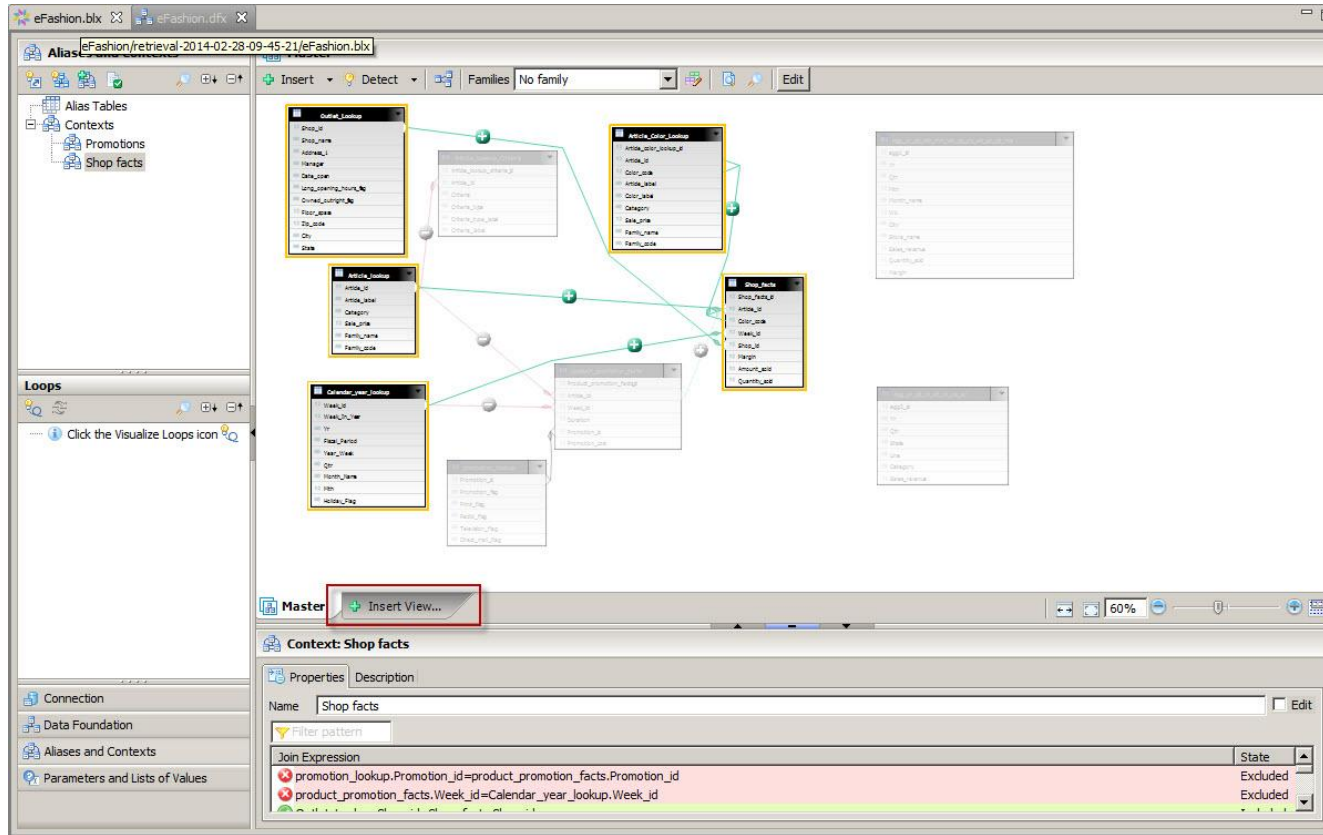


# Instead, extract common logic upstream

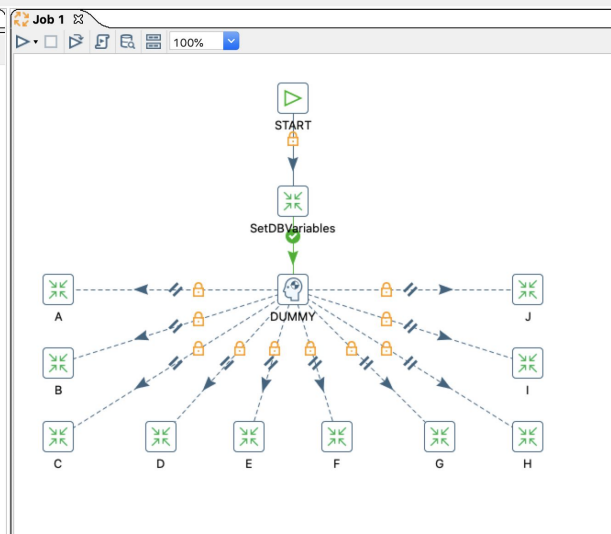

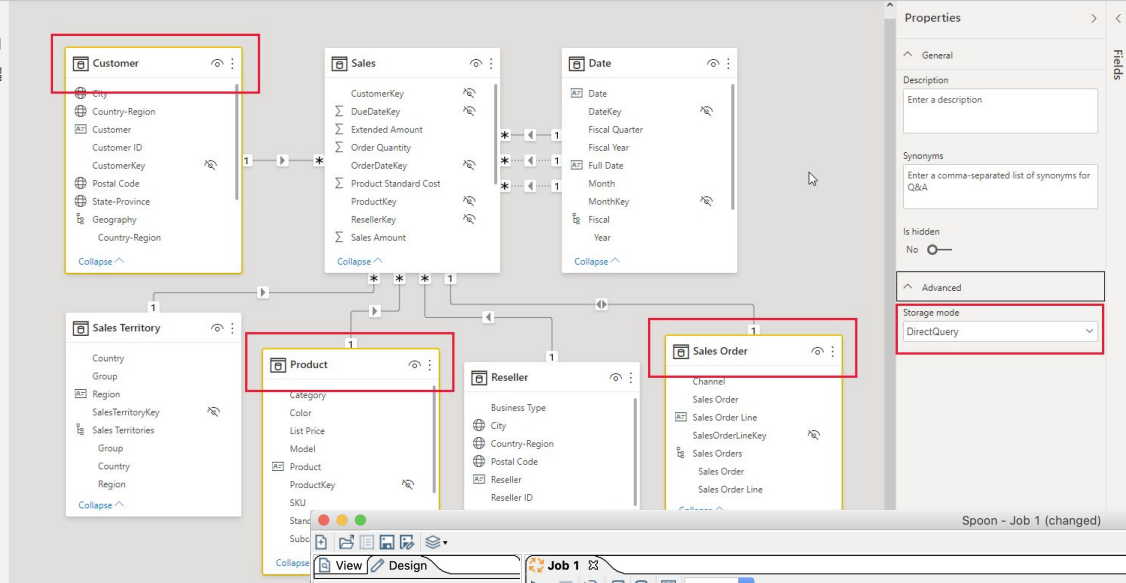
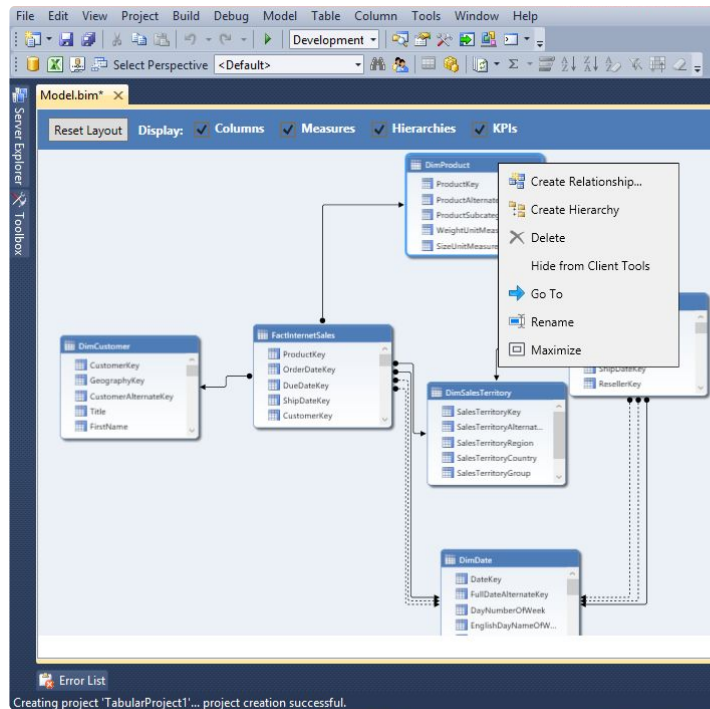


# The core of headless BI is the semantic layer—

—a concept that has existed for decades.







**The semantic layer was and still is  
a part of many, many BI tools.**

# Looker innovated the semantic layer by making it code-based

The screenshot displays the Looker web interface. At the top, a dark navigation bar contains the Looker logo, navigation links (Browse, Explore, Develop, Admin), and utility icons (search, help, user). Below the navigation bar, the page title "redshift" is visible. On the left sidebar, there's a section for "Your Personal Branch" (dev-nikolai-valiotti-gsdg) with a message about the production version and a "Pull from Production" button. Below this is a "Views" list with items like "adformat", "adformat\_container", "applications", "bundles", "campaign\_intents", "campaign\_portal\_details", "campaigns", "city", "country", "country\_details\_api", "creatives", "delivery\_jab\_categories", and "demand\_managers". The main content area shows the "redshift.model" editor. The code defines a connection to "redshift", includes views and dashboards, and defines a datagroup "redshift\_default\_datagroup". It also defines an explore "\_events\_agg" with a label "Events Hourly" and joins for "\_country", "\_state", and "\_city", each with a view label "Events Agg". The right sidebar contains a "Quick Help" section with a link to a help page and a "model:" section showing a JSON-like structure for a model definition.

```
1 connection: "redshift"
2
3 # include all the views
4 include: "/*.view"
5 include: "/*.dashboard"
6
7 datagroup: redshift_default_datagroup {
8   # sql_trigger: SELECT MAX(id) FROM etl_log;;
9   # SELECT DATE(CURRENT_TIMESTAMP)
10  max_cache_age: "1 hour"
11 }
12
13
14 persist_with: redshift_default_datagroup
15
16 # -----
17 explore: _events_agg{
18   label: "Events Hourly"
19
20   group_label: ""
21
22   join: _country{
23     view_label: " Events Agg"
24     type:left_outer
25     sql_on: ${_events_agg.country_id}=${_country.country_id};;
26     relationship:many_to_one
27   }
28   join: _state{
29     view_label: " Events Agg"
30     type:left_outer
31     sql_on: ${_events_agg.state_id}=${_state.state_id};;
32     relationship:many_to_one
33   }
34   join: _city{
35     view_label: " Events Agg"
36     type:left_outer
37     sql_on: ${_events_agg.city_id}=${_city.city_id};;
38     relationship:many_to_one
```

Quick Help →

A `model` references a combination of related explores. Unlike other LookML elements, a model is not declared explicitly with the `model` keyword.

```
model: {
  access_grant: identifier
  case_sensitive: yes or no
  connection: "string"
  datagroup: identifier
  explore: identifier
  fiscal_month_offset: number
  include: "string"
  label:
    possibly-localized-string
  map_layer: identifier
  named_value_format: identifier
  persist_for: "string"
  persist_with: datagroup-ref
  test: identifier
  view: identifier
  week_start_day: monday or ...
}
```

# LookML was a major factor in Looker's success

TECH DRIVERS

## Google cloud boss Thomas Kurian makes his first big move — buys Looker for \$2.6 billion

PUBLISHED THU, JUN 6 2019-9:26 AM EDT | UPDATED THU, JUN 6 2019-12:55 PM EDT



Lauren Feiner  
@LAUREN\_FEINER



Jordan Novet  
@JORDANNOVET

SHARE    

### KEY POINTS

- Google says it plans to acquire Looker for \$2.6 billion in cash.
- The acquisition of Looker is the first major acquisition under former Oracle executive Thomas Kurian.
- Kurian joined Google Cloud as CEO in November.

Ad covered  
content

# What if a semantic layer could be accessed outside of a BI tool?

Enterprise

## Google's Looker partners with Tableau


Frederic Lardinois @fredericl / 5:00 AM PDT • October 12, 2021

 Comment

In 2019, Google [acquired](#) the business intelligence service [Looker](#) for \$2.6 billion and Salesforce [picked up Tableau](#) for \$15.7 billion. But today, the two competitors are growing closer together thanks to a new integration between the two. Specifically, Tableau users will soon be able to access Looker's semantic layer, while Google's Looker users will soon be able to use Tableau's visualization layer on top of the Looker platform.

This may seem like an unlikely partnership at first, but in a way, this allows both services to play to their strengths. Tableau's advanced visualization capabilities have always been a draw for its users, but unlike Looker, it was born on the desktop and despite all of the company's efforts, that

# Headless BI: keep the semantic layer, caching, and access control—but decouple visualizations from them.



# 2020–2022: The rise of Headless BI tools in the modern data stack

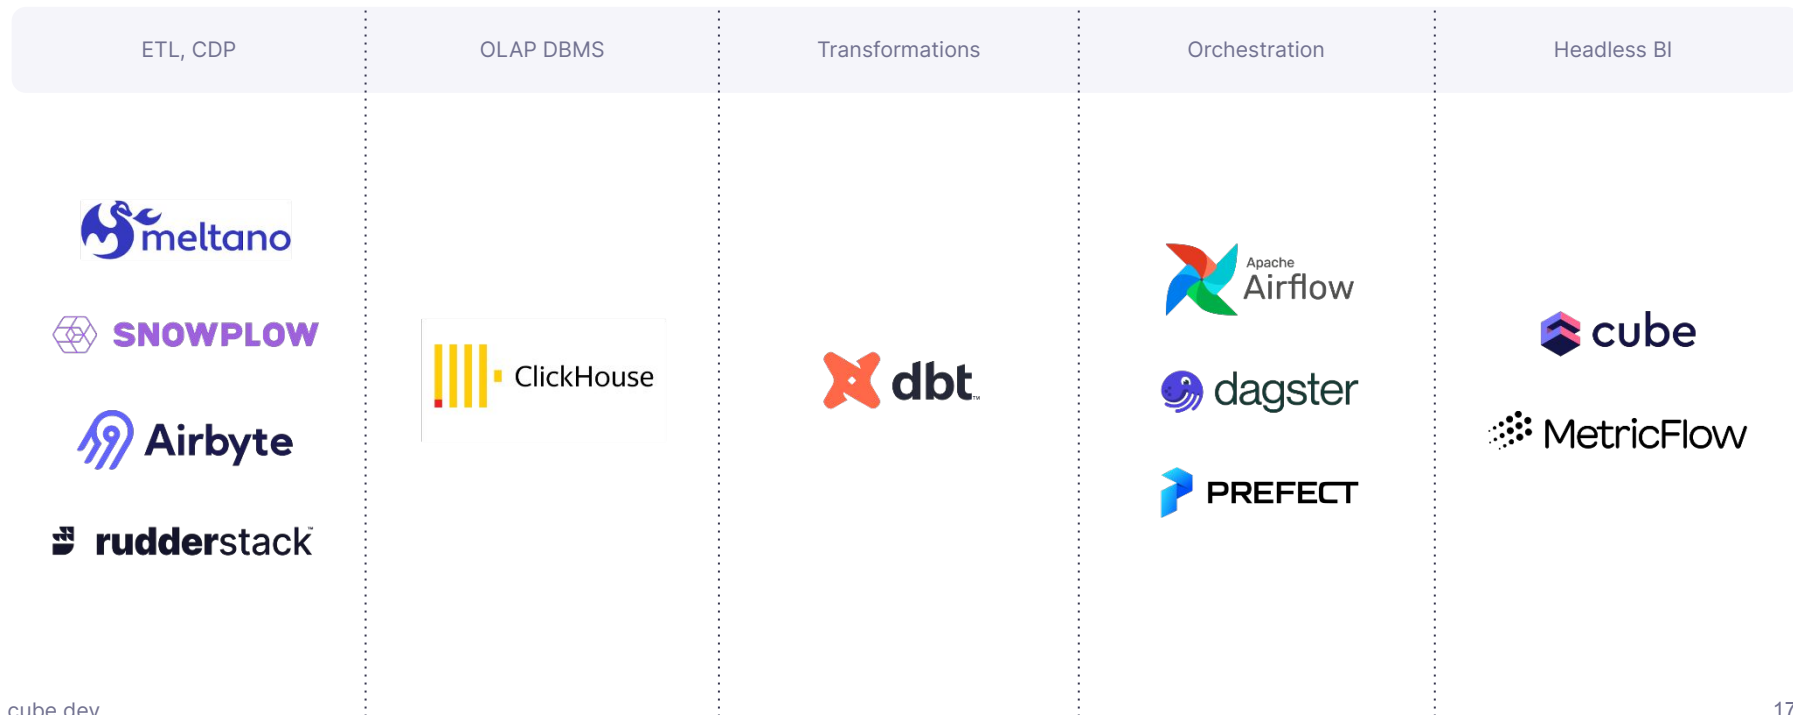


# Why open-source?

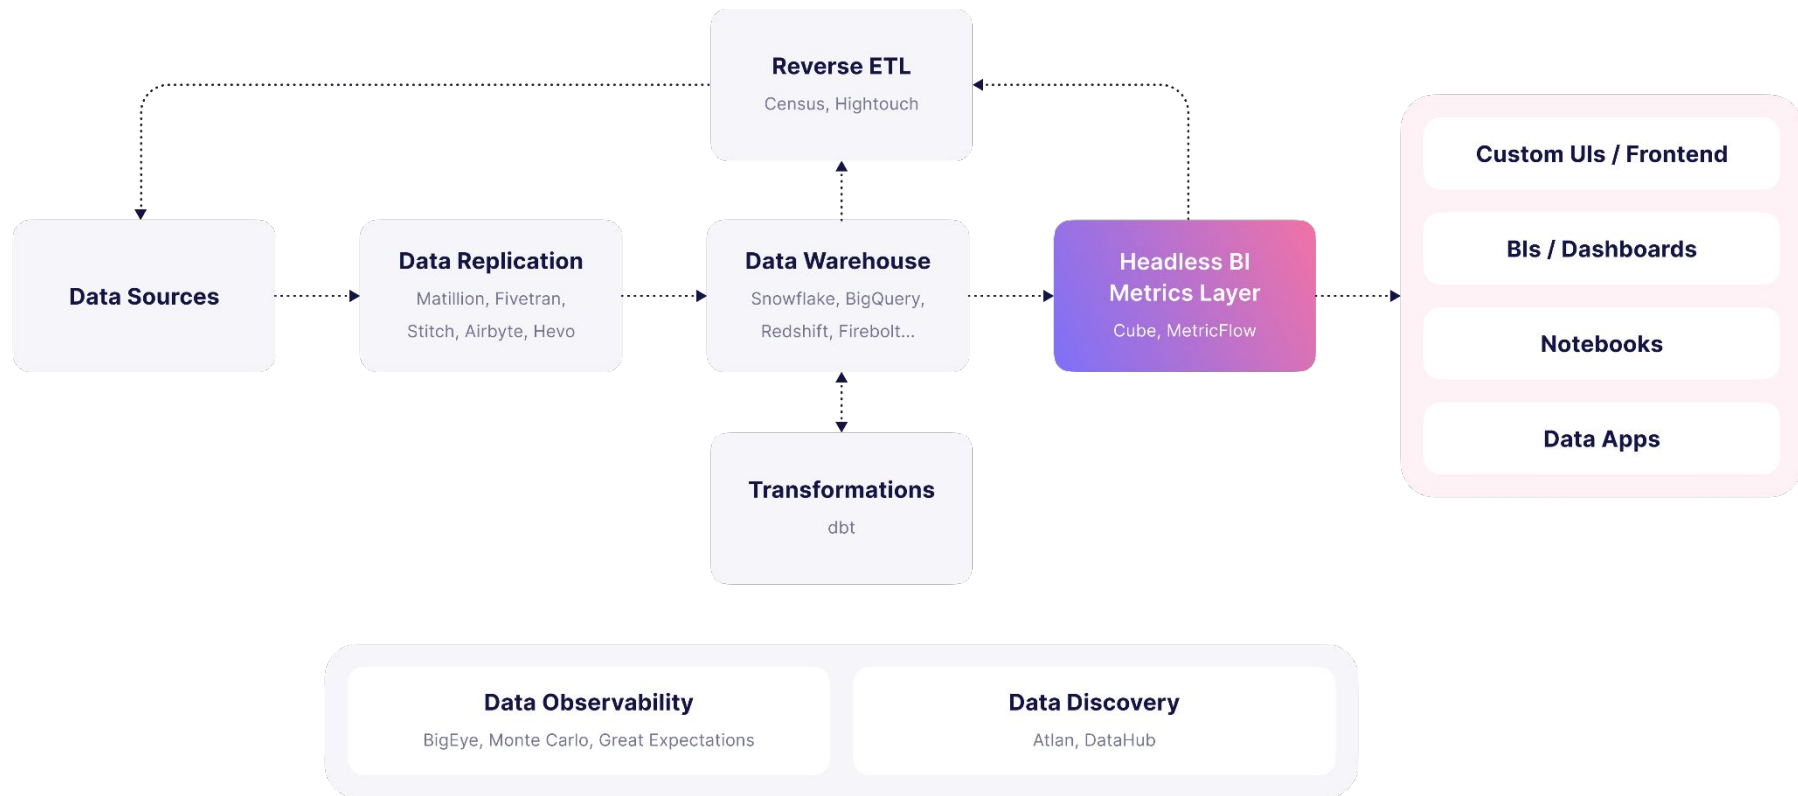


# The modern data stack is inherently more open

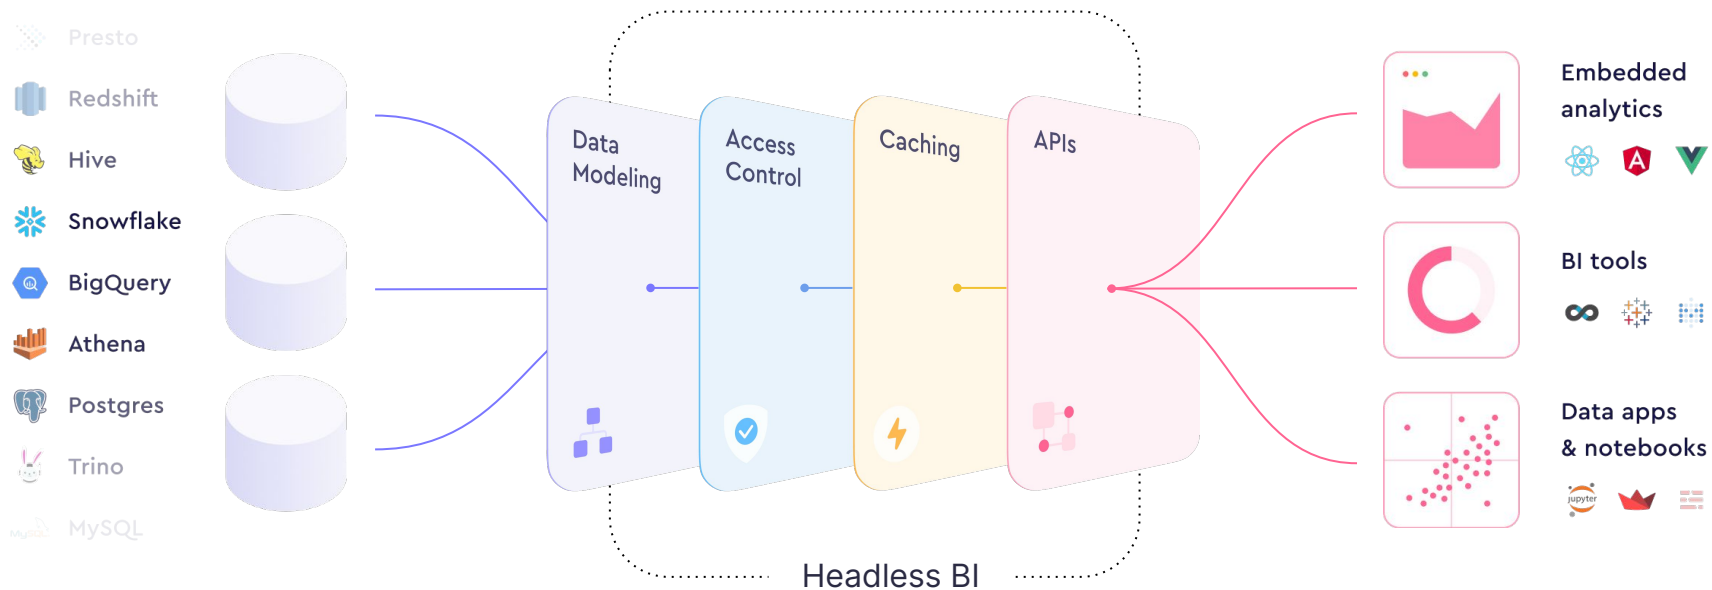
Many core components of the modern data stack are open-sourced



# Headless BI in the modern data stack




# Architecture of Headless BI



# Push down compute to CDW

Build on top of the innovation of previous generation BI.  
BI tools generate and execute SQL queries in CDWs.



- Power BI direct query
- Tableau live connections
- Thoughtspot
- Looker

# Data modeling (semantic) layer in code

```
cubes:  
  name: active_users  
  sql: SELECT * from events  
  
  measures:  
    - name: weekly_active  
      type: count_distinct  
      rolling_window:  
        trailing: 7 day  
        offset: start  
  
  dimensions:  
    - name: time  
      type: time  
      sql: timestamp
```




```
data_source:  
  name: transactions  
  sql_table: schema.transactions  
  
  identifiers:  
    - name: transaction_id  
      type: primary  
  
  measures:  
    - name: revenue  
      description:  
        expr: price * quantity  
        agg: sum
```



# Multi-tenancy and security control

Headless BI generates SQL from the data model within the specific security context that was supplied via incoming query, e.g. with JWT token in HTTP request.



# Caching

Caching via aggregations inside CDW or purpose-built aggregation storage

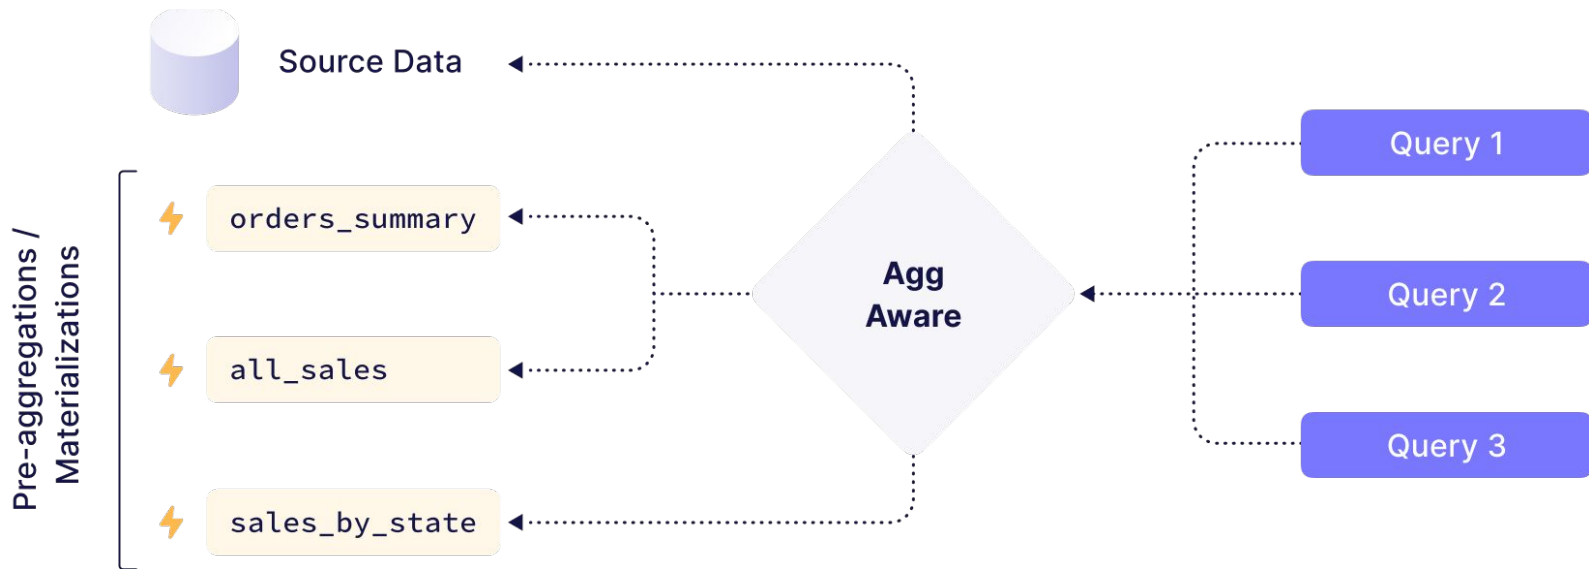
```
cubes:  
  pre_aggregations:  
    - name: orders_summary  
      measures:  
        - count  
        - total_value  
      dimensions:  
        - status  
      time_dimension: created_at  
      granularity: week
```



```
materialization:  
  name: user_bookings_summary  
  metrics:  
    - bookings  
    - booking_value  
    - customer_service_tickets  
    - guest_host_messages  
  dimensions:  
    - metric_time  
    - guest_signup_channel
```



# Querying with caching: aggregate awareness

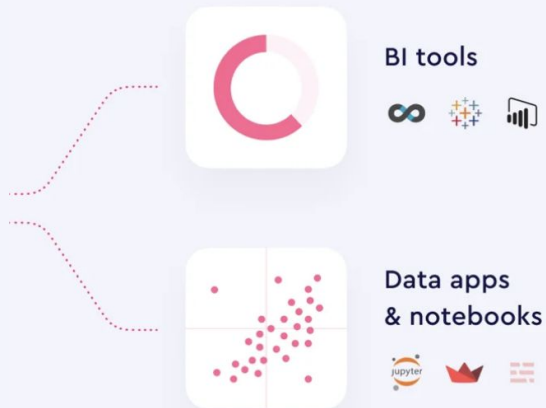




# Consumption level — APIs: SQL

Importance of SQL API; the SQL is lingua franca of data world

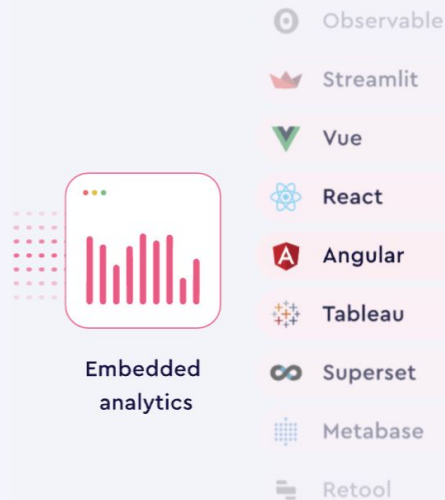
```
SELECT
  status,
  MEASURE(revenue)
FROM orders
WHERE STATUS != "completed"
GROUP BY 1
```



# Consumption level — APIs: REST/GraphQL

Rest/GraphQL for web developers

```
query {  
  cube {  
    orders(where: {  
      status: { notEquals: "completed" }  
    }) {  
      count  
      status  
      createdAt {  
        month  
      }  
    }  
  }  
}
```



**Thank you!**

