

**OSA CON 23**



MAYBE THE REAL MODERN DATA  
STACK WAS THE OSS WE GOT  
ALONG THE WAY

---

**Pedram Navid**

Head of Data Engineering, Dagster Labs  
[@pdrmnvd](https://twitter.com/pdrmnvd) / [linkedin.com/in/pedramnavid](https://www.linkedin.com/in/pedramnavid)

December 12-14 2023

# Modern Data Stack: Origins



# MDS: Origins



**Craig Van Spall**  
@CraigVanSpall



Not many understand the nuts and bolts of the modern data driven business. Why is a modern data stack relevant f... [bit.ly/1qOz4Np](https://bit.ly/1qOz4Np)

10:49 AM · Sep 29, 2014

# MDS: Origins



Product Announcements

## What you need to know about the sunset of consumer Google+ on April 2

January 30, 2019

**David Conway**  
Product Manager

# MDS: Origins



**Gonçalo Gaiolas**

@ggaiolas



Come build a modern [#data](#) stack to completely transform how a company utilizes to serve its customers. [#datascience](#) [lnkd.in/dB5\\_Kx8](https://lnkd.in/dB5_Kx8)

7:56 AM · Oct 3, 2016

# MDS: Origins



The screenshot shows the top navigation bar of the OutSystems website with the following elements: the OutSystems logo, a search icon, and menu items for Platform, Solutions, Developers, Why OutSystems, and Pricing. On the right side of the navigation bar are two buttons: 'Schedule demo' and 'Start Free'. The main content area has a black background with the text '404' in large white font, followed by 'Sorry, this page was not found.' Below this text are two buttons: 'Back to the Home Page' (a red button) and 'Discover the Platform' (a white button with a black border).

# MDS: Origins



**Treasure Data** @TreasureData · May 1, 2017

"Any organization can immediately have a **modern data stack** w/o investing years and massive engineering resources." [buff.ly/2oNB6Vr](https://buff.ly/2oNB6Vr)

The diagram illustrates the Customer Data Platform (CDP) architecture, divided into three main stages: **Connect**, **Store, Unify and Enrich**, and **Activate**.

- Connect (1st Party (customer data)):** Includes Ad placement logs, Web logs, Mobile app logs, Social data, POS data, CRM data, Customer Survey, E-Commerce, and User demographics.
- Store, Unify and Enrich (Customer Data Platform):** A central cloud icon represents the CDP. Below it, an **Enrichment** section shows data from **2nd & 3rd Party DMP** being added. This includes Statistics, Weather, Income, Age/Demographics, Data Exchange, Interest, POS data, and Occupation.
- Activate:** Data is used for **Smarter Advertising** (Google, Facebook, Twitter), **Personalized Communication** (Social, Messaging apps, Email, Website personalization), and **Accessible Data Science & BI** (Python, R, Tableau).

**HOT STARTUP IS BECOMING THE TREASURE TROVE FOR LIVE DATA MANAGEMENT**

The explosion of mobile apps, new business applications, scalable data storage, and extensible web services has created a flood of data and is driving reinvention of business models and operations. Treasure Data empowers decision makers through making Big Data simple, actionable and accessible.

GAIL BARNES April 25, 2017 · 6 min read

**SUBMIT YOUR STARTUP** SNAPMUNK | STARTUP DIRECTORY

What to do if you are Japanese, don't speak much English, but you have some good friends and a talent as big as the Big Data that your startup explains and makes sense of? If you're Kazuki Ohta, you decide to move to Silicon Valley and build your startup there, and, if you're as lucky as you are smart, you score a meeting the 'Kite VC', aka Bill Tai, who commits to funding you during your first meeting. The kicker is that you don't realize that he has committed because your English isn't that good yet.

Kaz Ohta

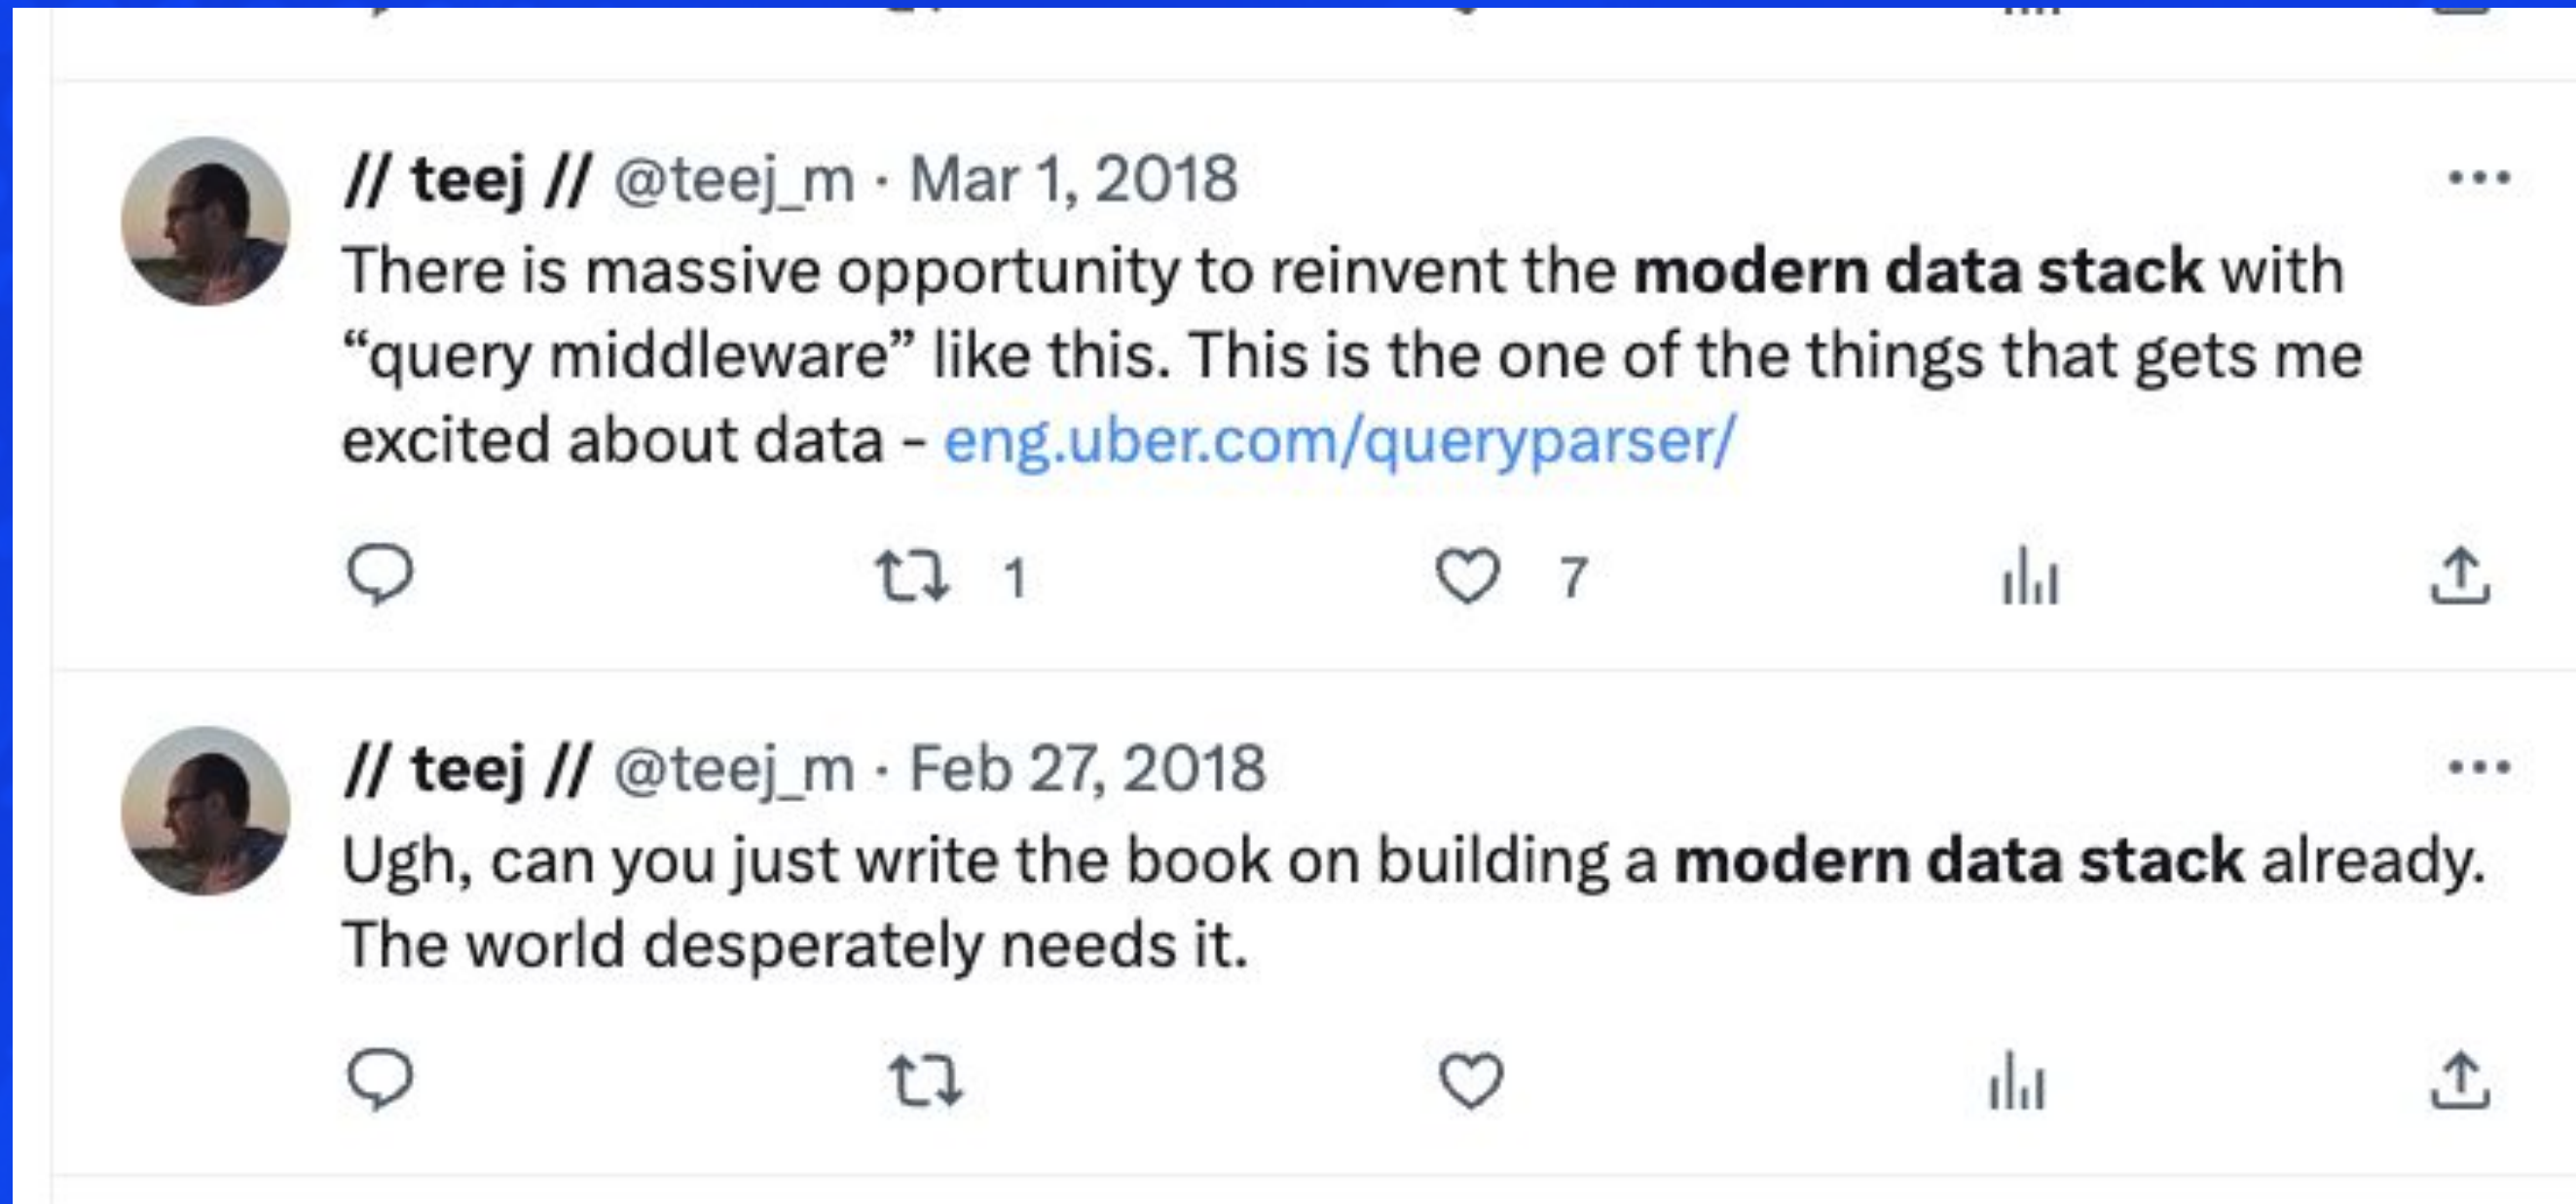
"The BEST fun 'founder' story is this one," Bill Tai told SnapMunk. "Fifteen minutes into our first meeting, I decided I would write a personal check to fund Kaz Ohta and Hiro Yoshikawa, to start a 'Big Data' company – Treasure Data – with and around them. I spent the rest of the meeting brainstorming about what to build and who to pull in to enhance the chance of success. Years later, I learn that Kaz didn't know that I had COMMITTED in the meeting, because he was not conversationally fluent in English back then and likely not expecting instant backing. Treasure Data, Inc. is now ingesting 4 trillion rows of data per month for hundreds of enterprise customers 'round the world..."

**Save Money While Going Green With Arcadia Power**

**3 Compelling Reasons Why Your Tech Startup Needs A .TECH Domain**

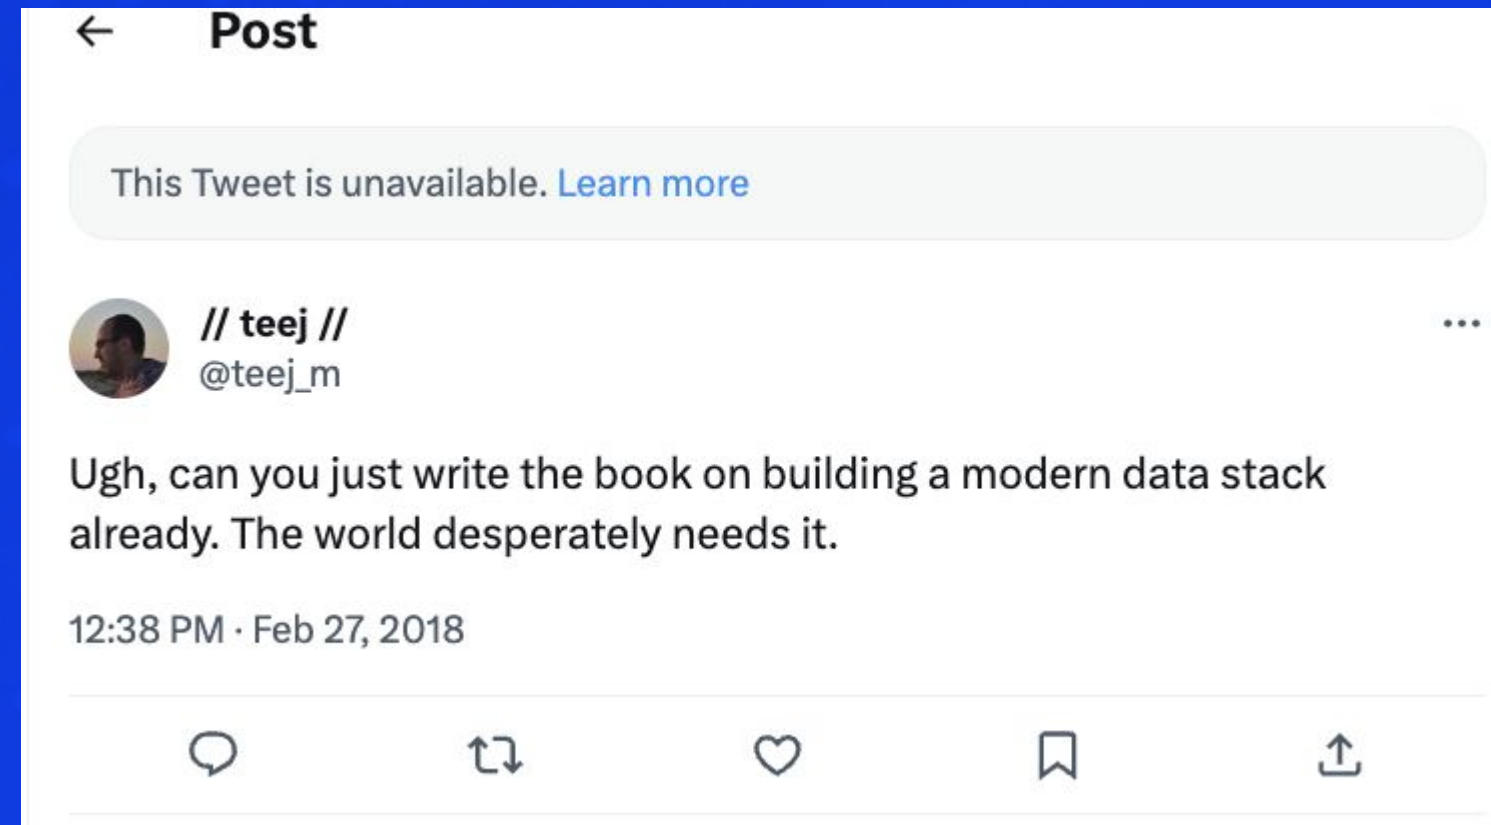
**Why the Last, Shittiest Ever Black Mirror Episode Is Definitely Its Best Ever Episode**

# MDS: Origins





# MDS: Origins





**Mode Analytics** @ModeAnalytics · Apr 17, 2019

Mc  
be  
da



**Snowflake** @SnowflakeDB · Aug 9, 2019

With a modern #data stack, fitness tracking app @Strava gains insight into the customer journey and makes more informed marketing decisions.

#SnowflakeDB ❄️

Learn how Snowflake + @fivetran + @tableau teamed to make this a reality:

fivetran.com  
Strava Centralizes Data, Runs With the Results | C...  
With a modern data stack, fitness tracking app  
Strava gains insight into the customer journey and ...



**Jake Klein (Taylor's Version)** @JakeKlein94 · Mar 13, 2019

Wanna do big data but not really sure where to get started? Check out this

he

SPet

modern

ng, led  
the



**OpsMat**  
The latest  
**Data Sta**  
Accelerate  
[ow.ly/xQ](http://ow.ly/xQ)



**Software**  
How to E



**Aditya A**  
From Da  
**Data Sta**



**Daryl Ki**  
From Da  
**Data Sta**

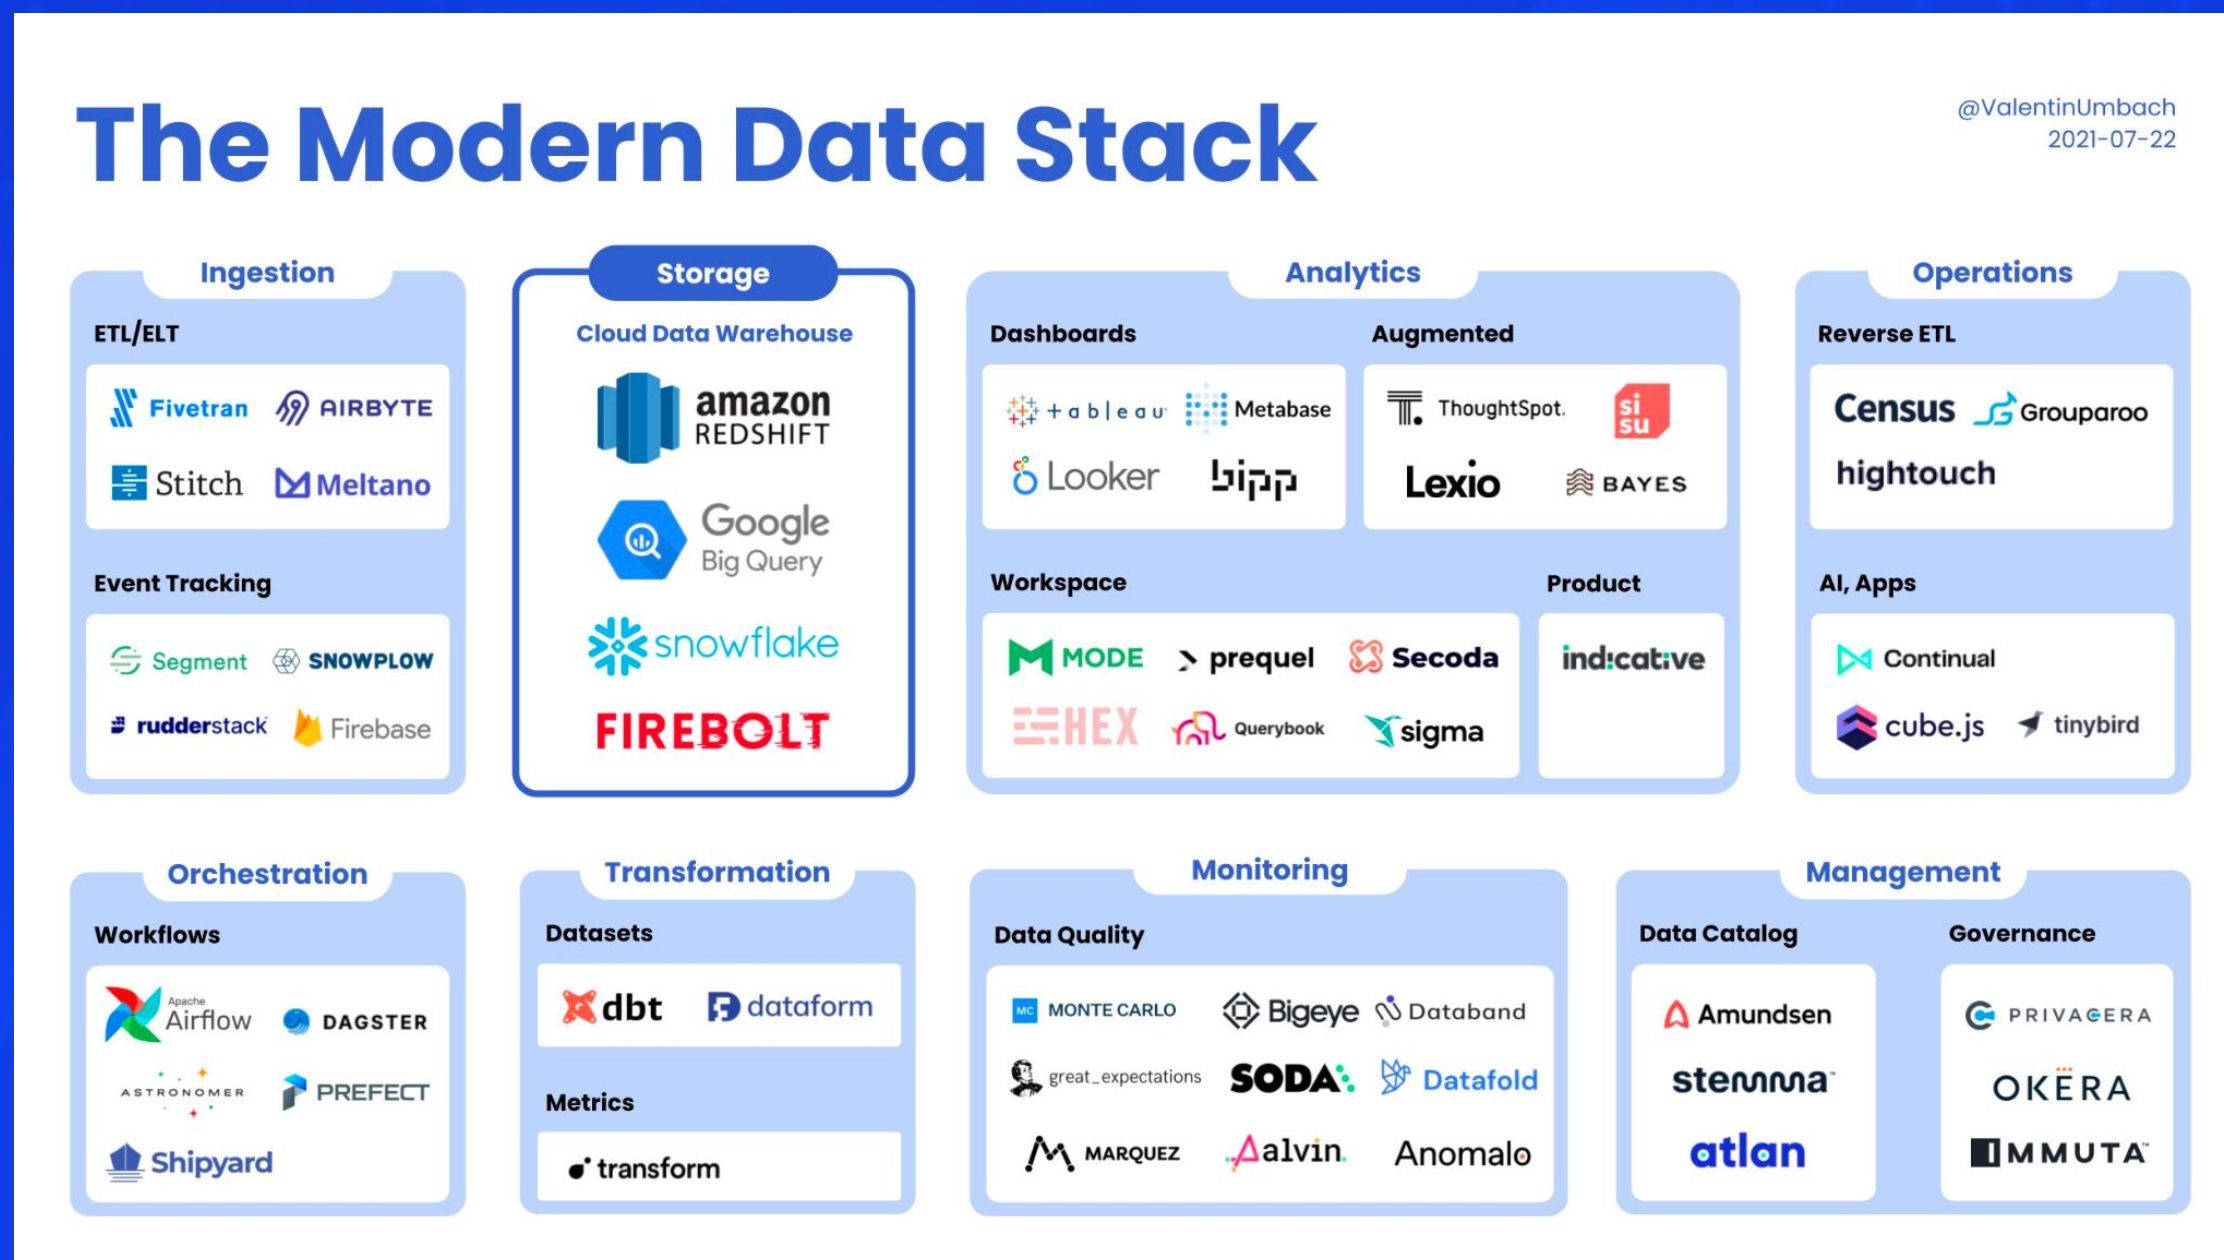


**Big Data**  
The Mod  
[bigdatadn.com/intelligence/t...](http://bigdatadn.com/intelligence/t...)



# What is it?

- Composable pieces
- Centered around the Cloud Warehouse
- Easy to implement and plug together



# Hype Cycle



# Goodbye, World



data

tech

## Is the Modern Data Stack dead already ?

Huge hype, an ocean of critics. Are we done with the modern data stack ?

🕒 Jul 17, 2022

### Apple Podcasts Preview



48 min

PLAY ▶

## Ep.99 'Is the Modern Data Stack Dead' with Special Guest Chris Tabb

Drill to Detail

Technology

[Listen on Apple Podcasts ↗](#)

[Episode Website ↗](#)

[More Episodes](#)

Rittman Analytics 2023

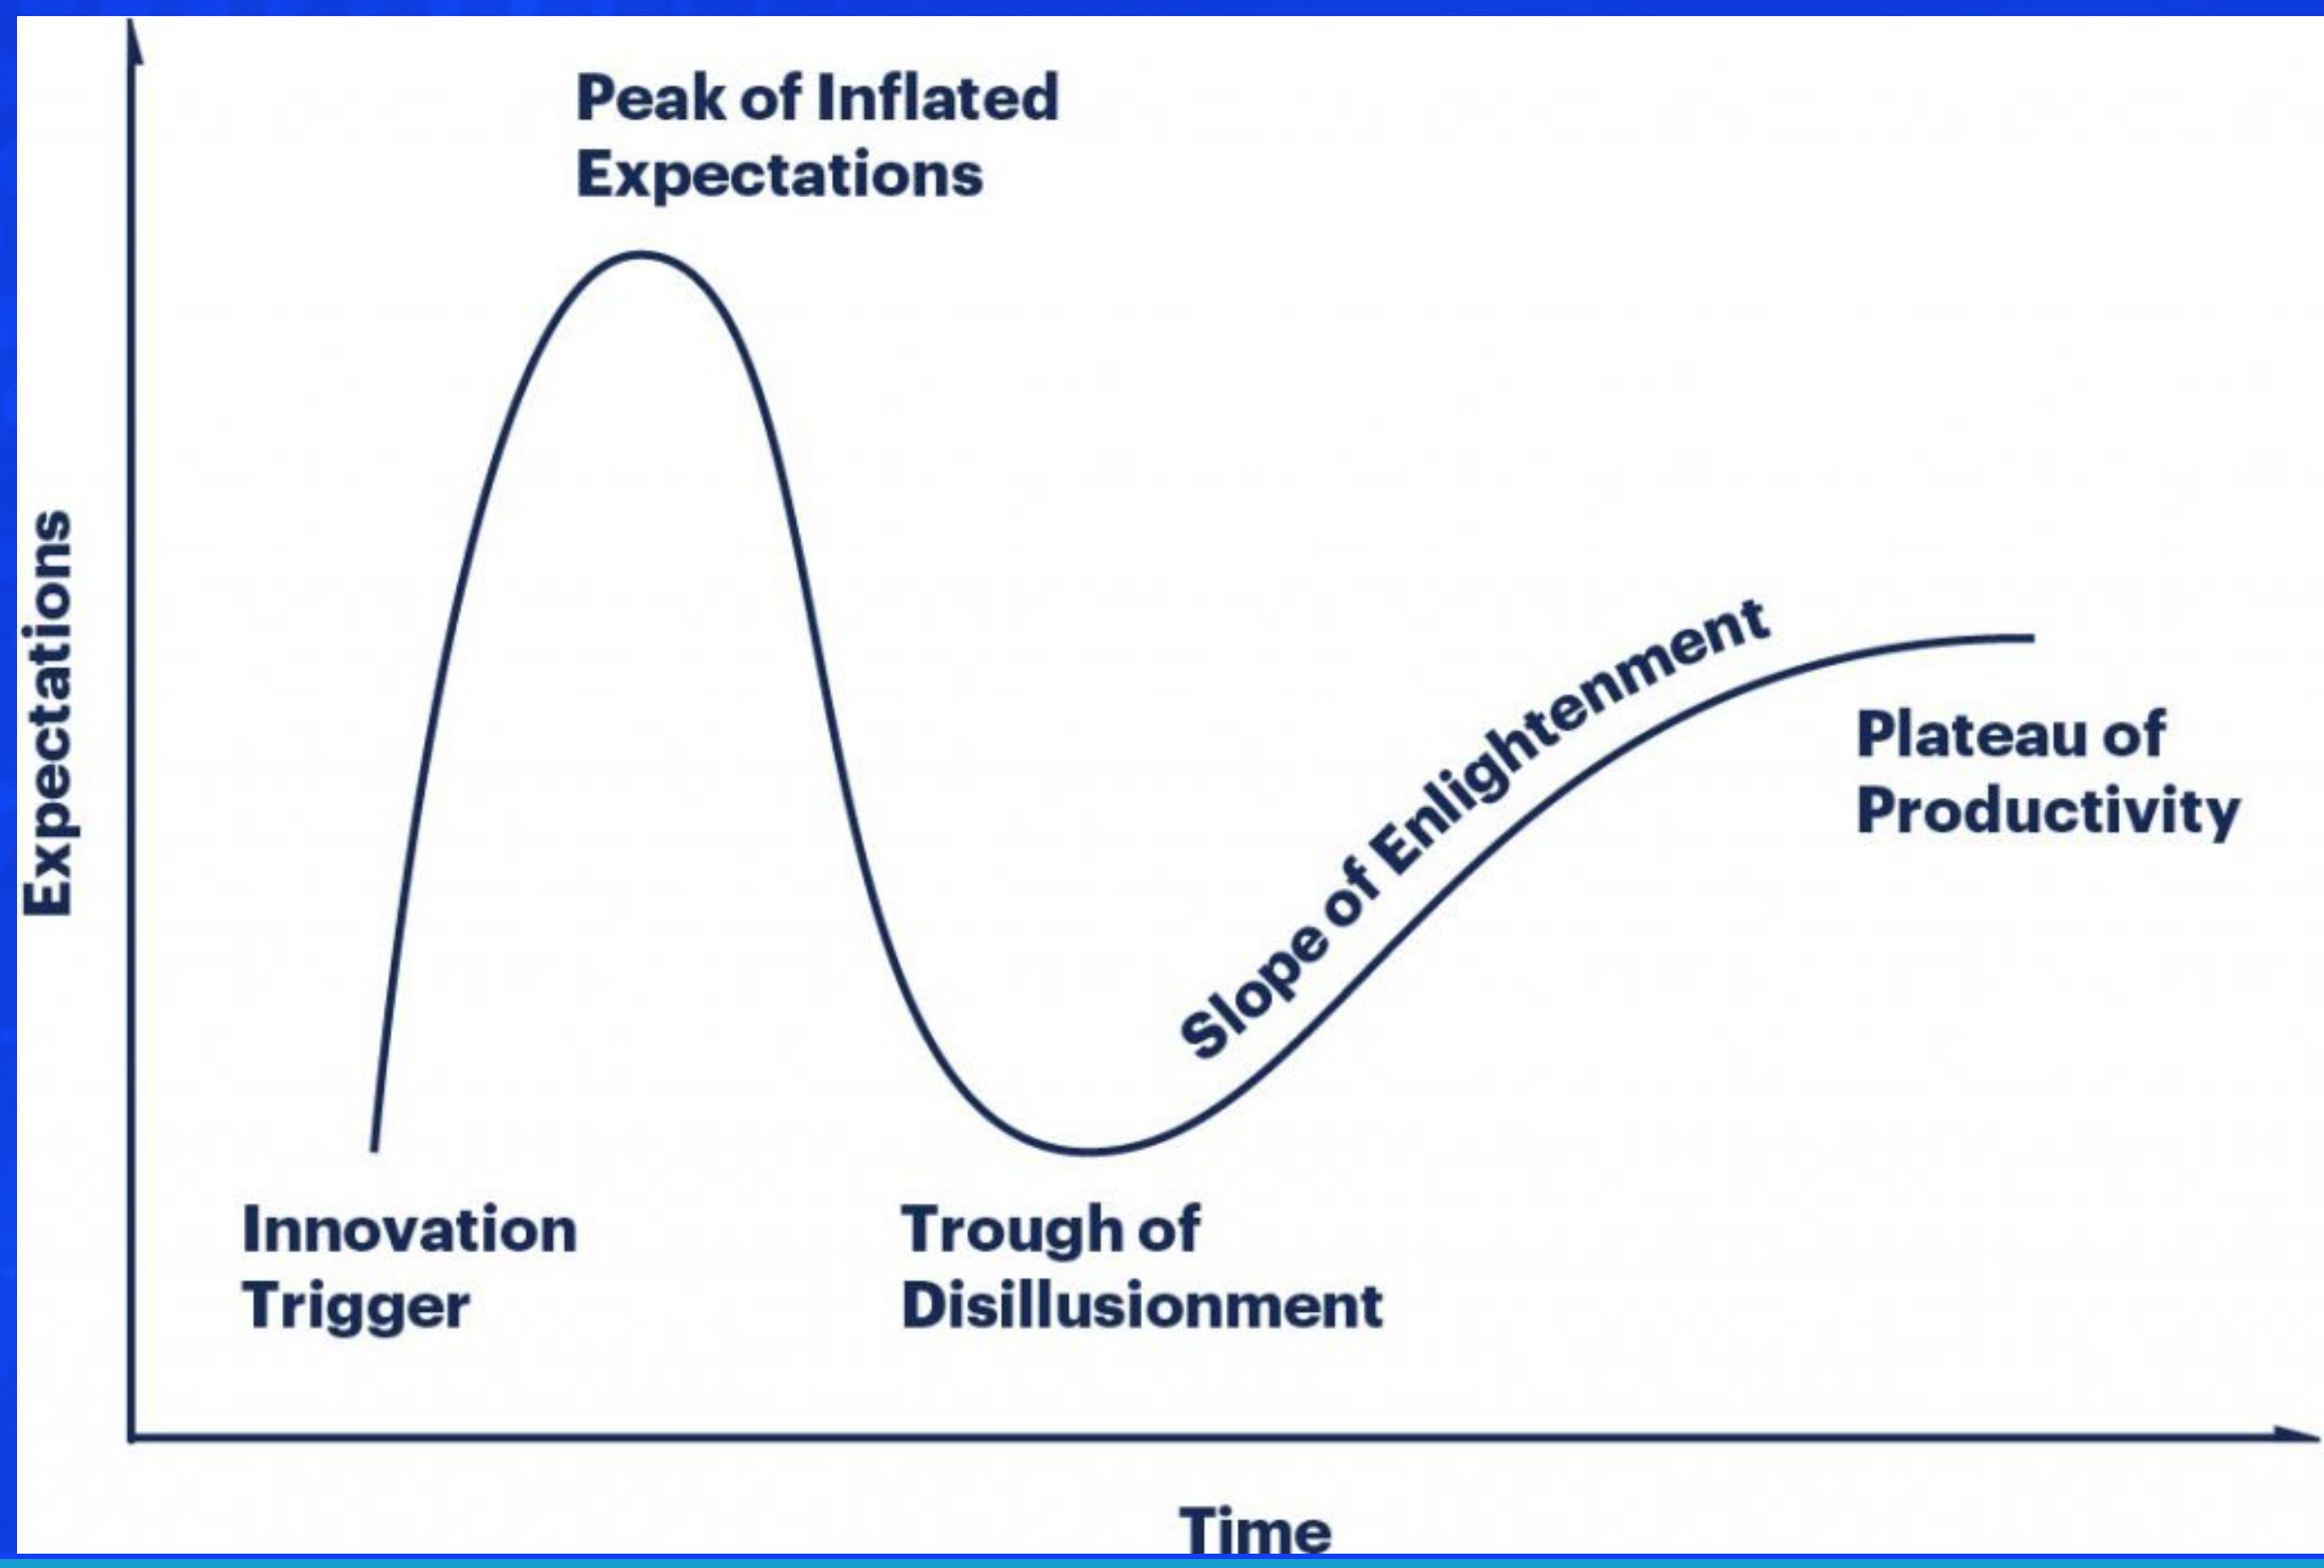
## The Modern Data Stack is Dead... Long Live the Modern Data Stack - Part 1

Where we've come from, lest we forget



DAVID JAYATILLAKE

JUL 19, 2023





# Nail in the Coffin

Blog > Updating dbt Cloud pricing to support long-term community growth



## Updating dbt Cloud pricing to support long-term community growth





 Tristan Handy | December 15, 2022 | Product News











Blog > Consumption-based pricing and the future of dbt Cloud






## Consumption-based pricing and the future of dbt Cloud

 Tristan Handy | August 8, 2023 | Company News

 **Josh Devlin (he)**  6:08 PM

  Tristan on consumption based pricing   (edited)

 21  2   11  4  2  4  2  2 

     73 replies Last reply 1 day ago

# Nostalgia Time



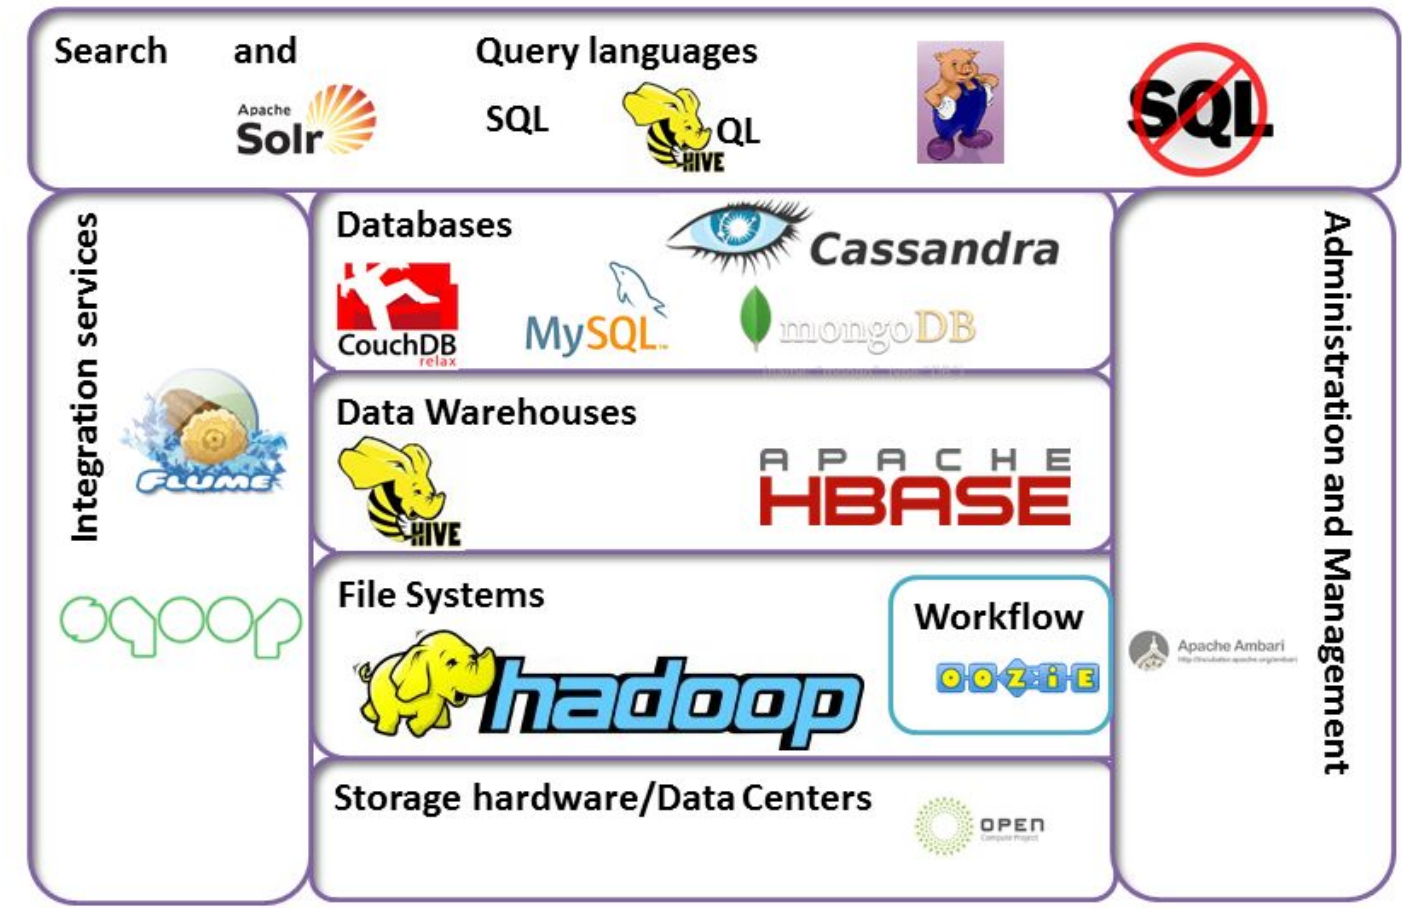


# The Technology Stack of Big Data

This is the wider Big Data landscape.

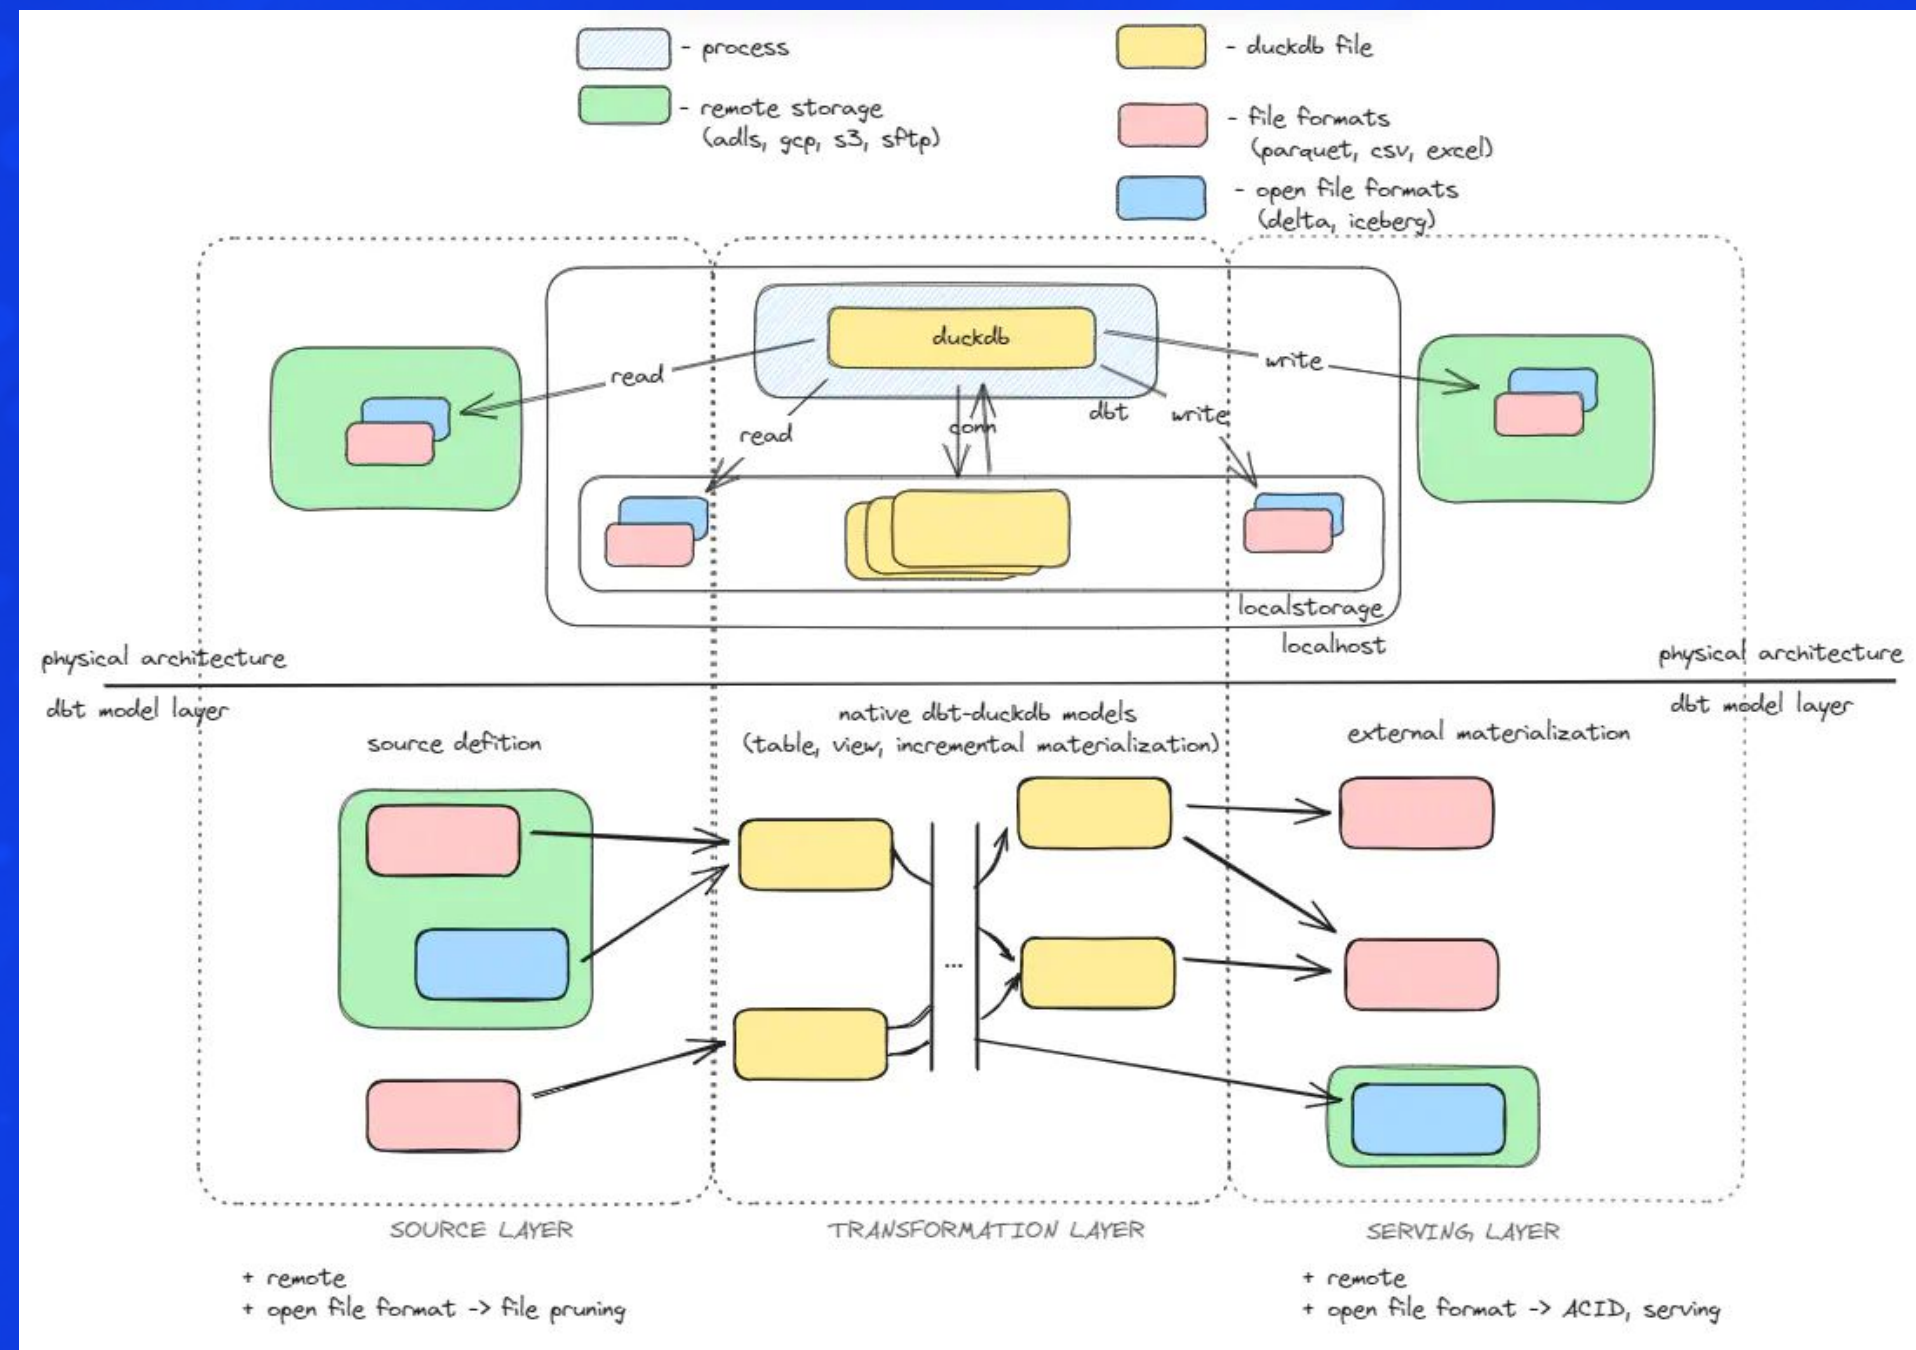


# The Big Data Open Source Technology Stack



Created by Pravi Solutions


# The Open Source Data Stack Architecture



<https://georgheiler.com/2023/12/11/dagster-dbt-duckdb-as-new-local-mds/>

# Data Pipelines

- Ingestion
- Storage
- Transformation
- Orchestration
- Activation
- Visualization



# Ingestion

- Plain Ol' Python Scripts
- Meltano / Alto
- Airbyte
- dlt
- Steampipe
- Sling
- pg2parquet
- And many many others



# Storage

- Postgres
- Clickhouse
- DuckDB
- Parquet
- Arrow
- Deltalake



# Transformation

- dbt
- dbt-duckdb
- polars
- Arrow DataFusion



 **Katie Bauer** @imightbemary · Aug 21  
Software engineers hate no code until someone asks them to write a SQL query

11 9 104 9,446

 **pedram** 🐥   
@pdrmnvd

software engineers will literally invent another dataframe framework before writing one line of sql

9:43 PM · Aug 21, 2023 · 1,348 Views

# Orchestration

- Dagster
- Prefect
- Airflow

## Conclusion

Regardless of whether you're an Airflow stan or one of the 500 data influencers who are being held at gunpoint to promote Mage (sorry Tommy, I still love you), I think we can all agree that if you can avoid Airflow, you'll be better off doing so.



# Visualization

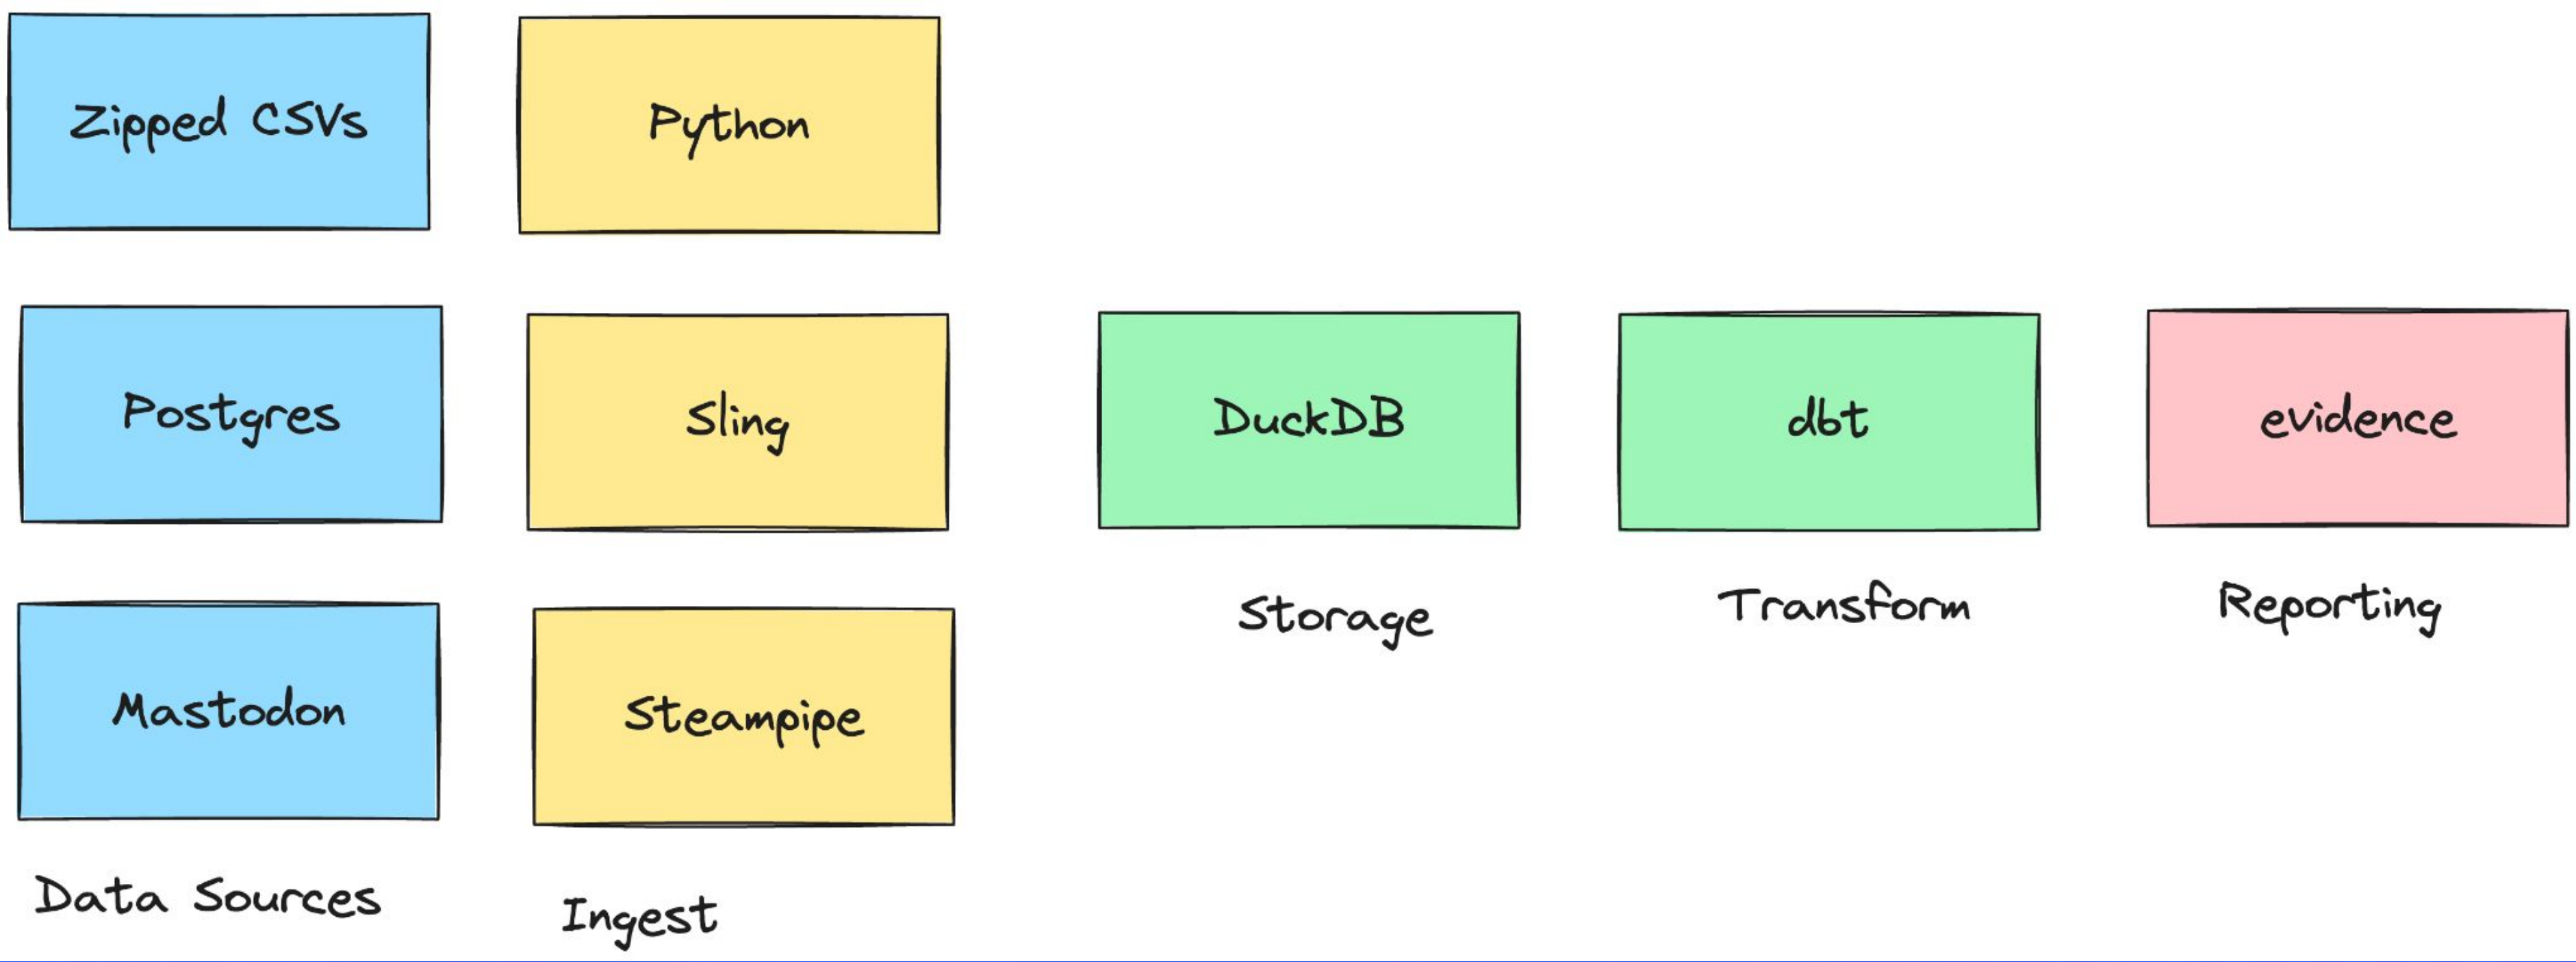
- Superset
- Evidence.dev
- Lightdash
- Metabase







# Let's Build a Stack



# Demo Time

Github Repo:

<https://github.com/dagster-io/mdsfest-opensource-mds>

Source Data:

<https://feederwatch.org/explore/raw-dataset-requests/>





# Q&A

## Follow Me

- ex-Twitter: @pdrmnvd
- LinkedIn: Pedram Navid
- Dagster Slack: [dagster.io/slack](https://dagster.io/slack)