

The background of the slide is a dark, atmospheric landscape at night. On the left, a large, craggy rock formation rises. A road or path, illuminated by a bright purple light, curves from the bottom right towards the center of the image. The sky is dark blue with visible stars. The Freewheel logo is in the top left corner.

FREOWHEEL
A COMCAST COMPANY

PROCESSING BILLIONS OF AD-TECH EVENTS IN REAL-TIME

MARGI DUBAL

March 30, 2023




WHO ARE WE?

FREEWHEEL A COMCAST COMPANY

COMCAST/NBCU IS #1 SELLER OF VIDEO ADVERTISING IN U.S.


2021 Total U.S. Video Ad Revenue
(\$ in billions, ex. Social)



FREEWHEEL POWERS ~90% OF PREMIUM VIDEO IN U.S. & EUROPE



BEESWAX AT A GLANCE

| | | | | |
|--------------|---|---|--|---|
| Who We Are | Industry experts who built a demand-side platform (DSP) purpose-built for premium programmatic | | | |
| What We Do | Solve programmatic challenges through bespoke in-housed implementations focused on transparency, flexibility, and control | | | |
| Who We Serve | Brands  Uber | Tech/Media Companies    | Agencies   | OEMs  |

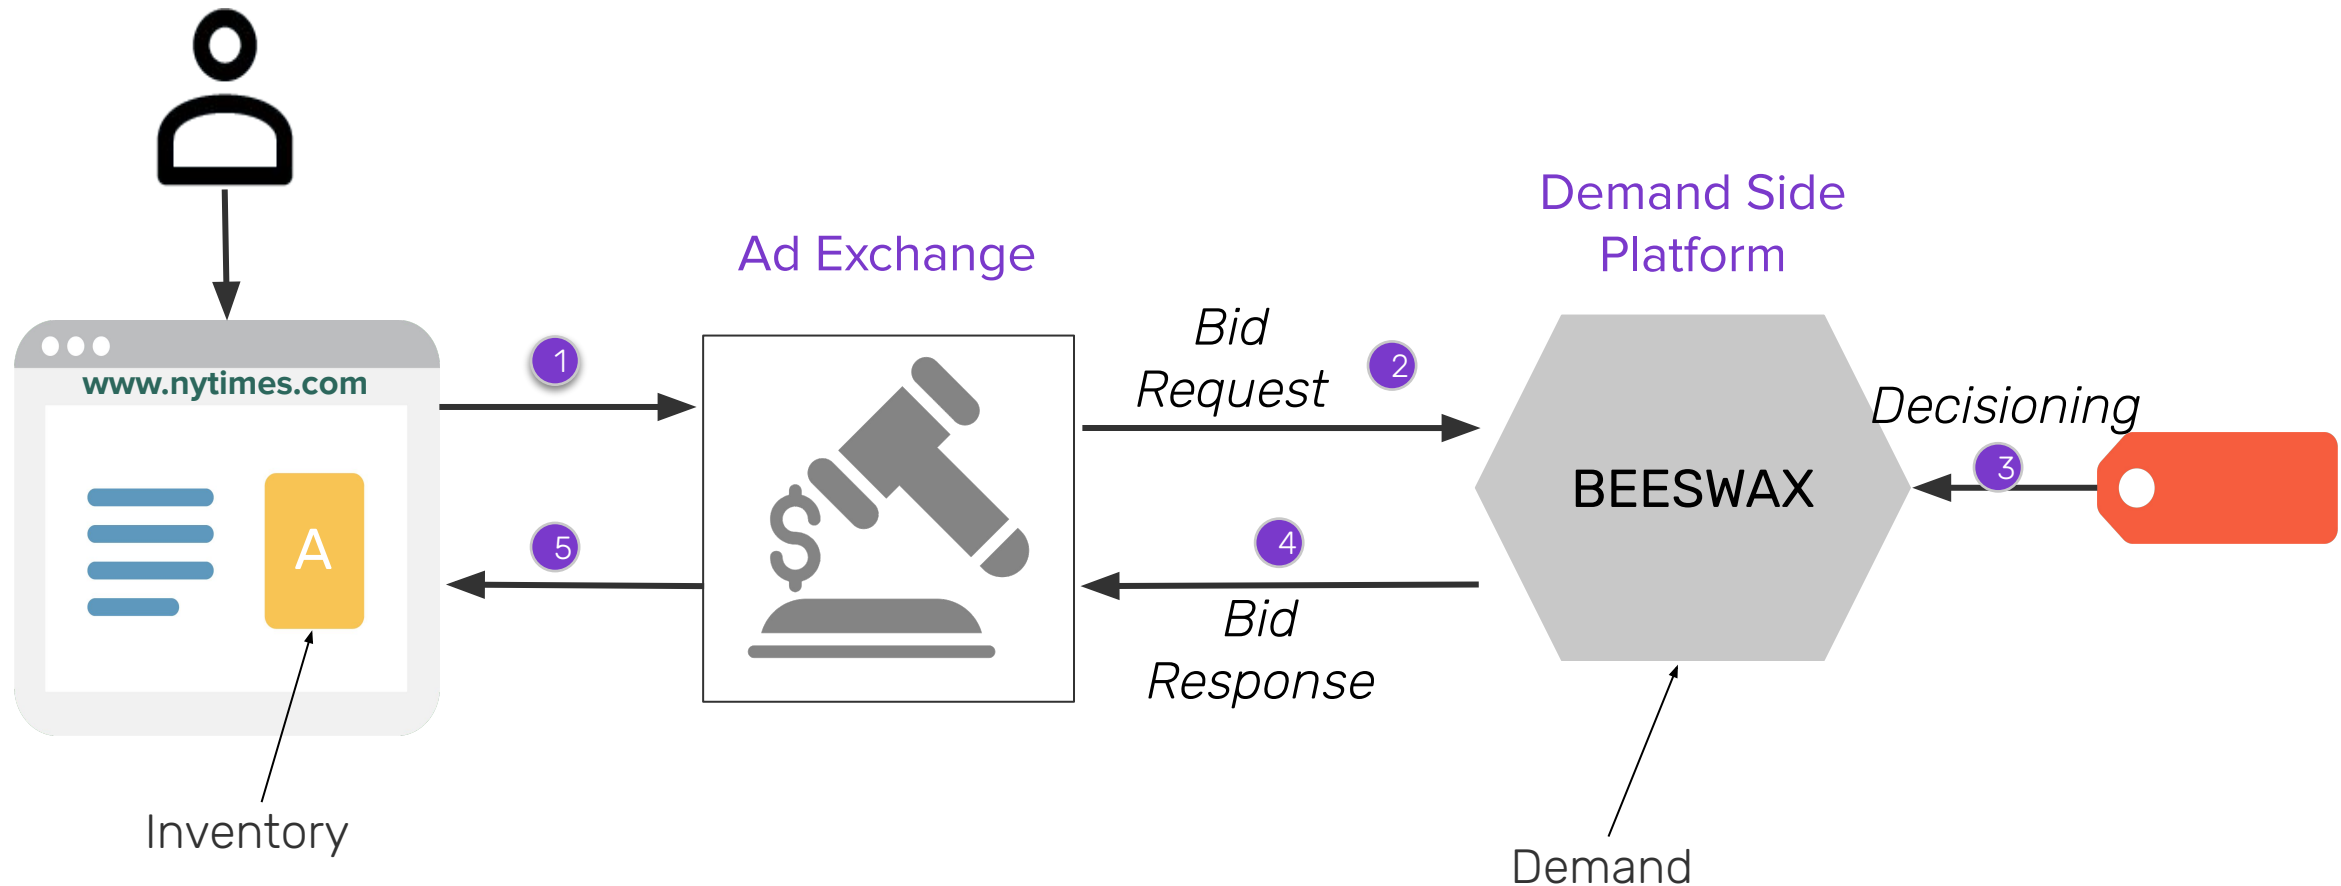


BASICS OF RTB (REAL TIME BIDDING)

Real-time bidding (RTB) is a method of transferring media on an individual impression basis, which is bid on in real time via an 'auction'.

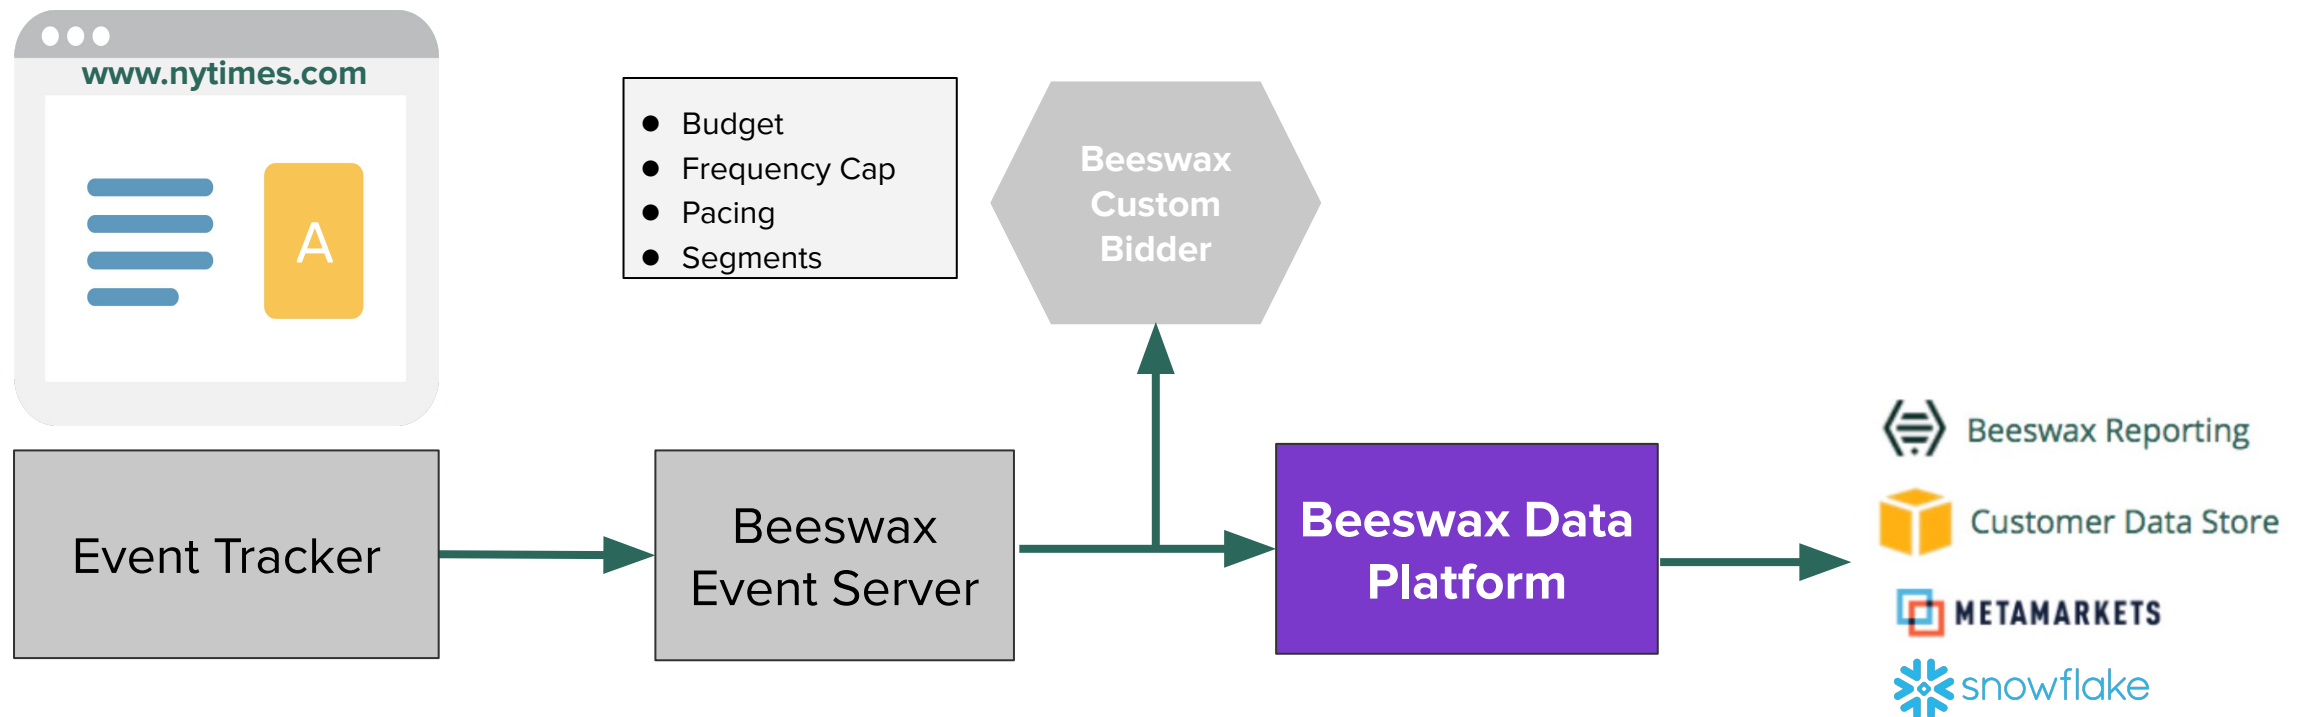
ADTECH 101 & REAL TIME BIDDING

Real time bidding basics



WHAT HAPPENS AFTER AD IS SERVED

Real time bidding basics



EVENT TYPES

| | | | | | | |
|------------------|--------|----------|-----------|-----------|------------|------------|
| Bid Request | Req ID | Creative | | | | |
| Win Notification | Req ID | | Win Price | | | |
| Event Tracker | Req ID | | | Click URL | Time stamp | Event Type |
| Complete Record | Req ID | Creative | Win Price | Click URL | Time stamp | Event Type |



INFRASTRUCTURE HIGHLIGHTS

DATA SCALE

Infrastructure Highlights



Every
second

2.5M

Concurrent
requests

200K – 500K

QPS to data
platform




Every
day

10B

Ad Events
logged

200TB

Data
processed



Every
year

400TB

Data for reporting &
analytics

6T

Bid requests
handled

DATA PLATFORM

COLLECTING, ORGANIZING & DELIVERING DATA

> Raw logs delivery & accurate campaign performance reports

> 50+ Data pipelines are scheduled using Airflow

> Elastically scalable infrastructure, optimized for cost

> Architected to run on AWS Spot instances and use autoscaling

> Diverse set of datastores

> Aerospike (KeyValue), Snowflake (Datawarehouse), Kinesis (Stream), S3 (Objects), RDS (MySQL/Postgres), InfluxDB (Timeseries)

DATA PRODUCT – DIVERSE DELIVERY METHODS

Infrastructure highlights

STREAMING LOGS

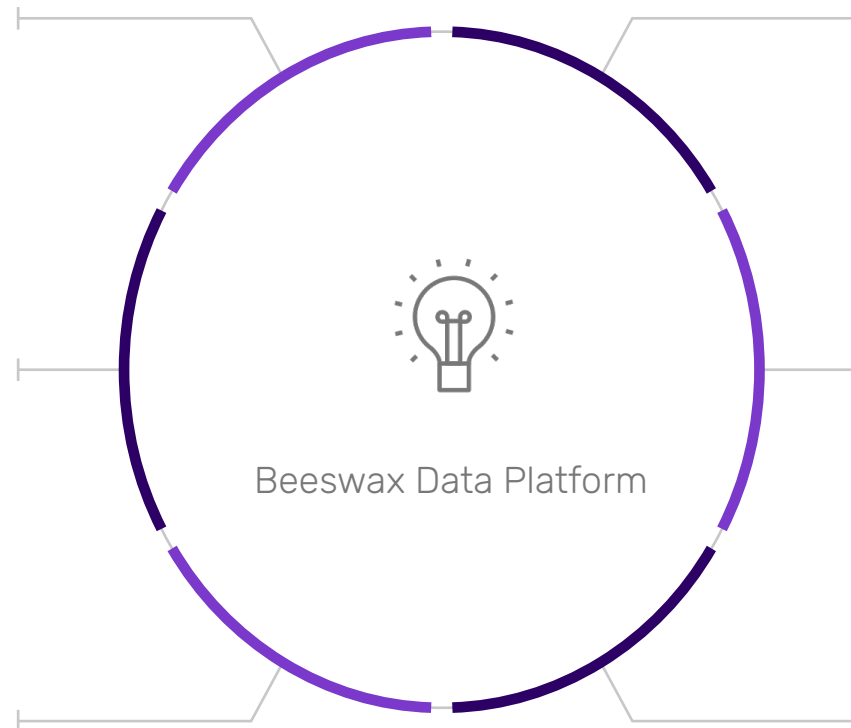
Real-time log delivery.
Support S3, Kinesis and HTTP integrations..

ANTENNA DATA SHARING

Allow customers to access raw log data using Snowflake.

BATCH LOGS

Deliver hourly and daily logs to customers' S3 buckets.



REPORTING UI

14 built-in reports.
Embedded Looker UI provides various visualizations and scheduled reports.

REPORTING API

Leverage Looker API to provide same capability as reporting UI.

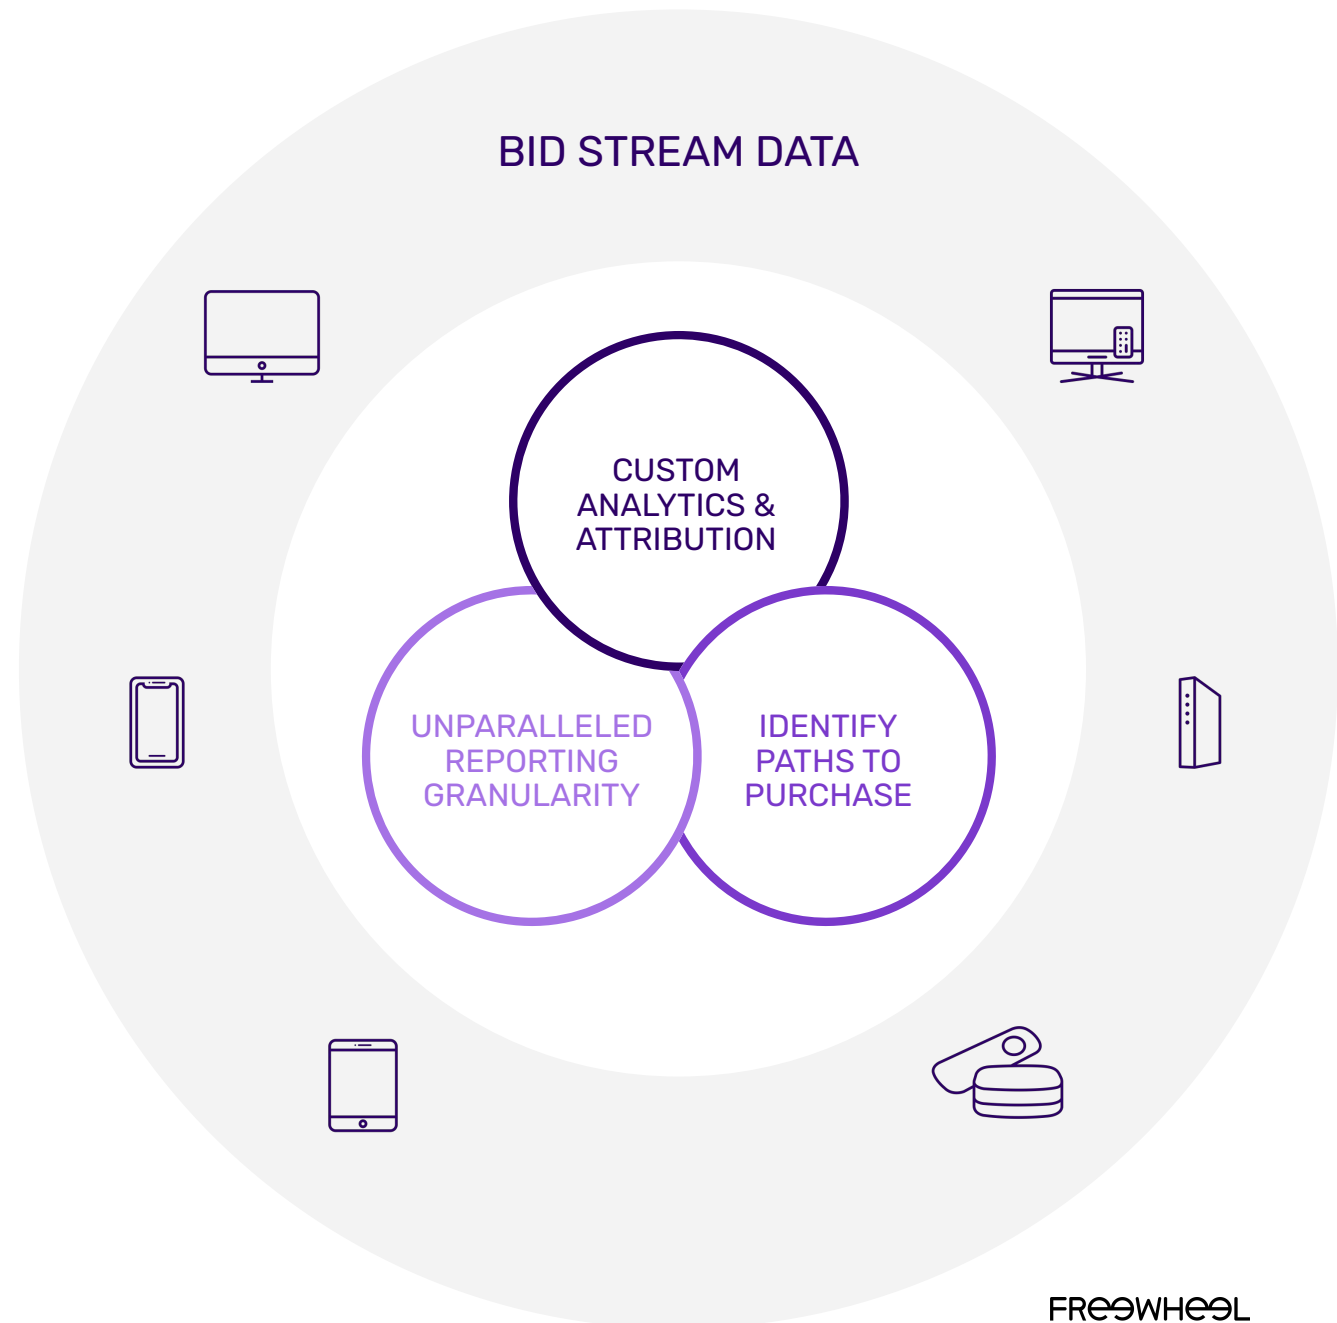
REAL-TIME BIDDING STATS

Visualize the real-time bidding/winning status.

DEEPER INSIGHTS WITH EVENT LOG DATA

Transparency into why campaigns perform and who drives performance

- Complete control and flexibility during reporting
- Join data any way you want to identify custom insights
- Make the most of your spend



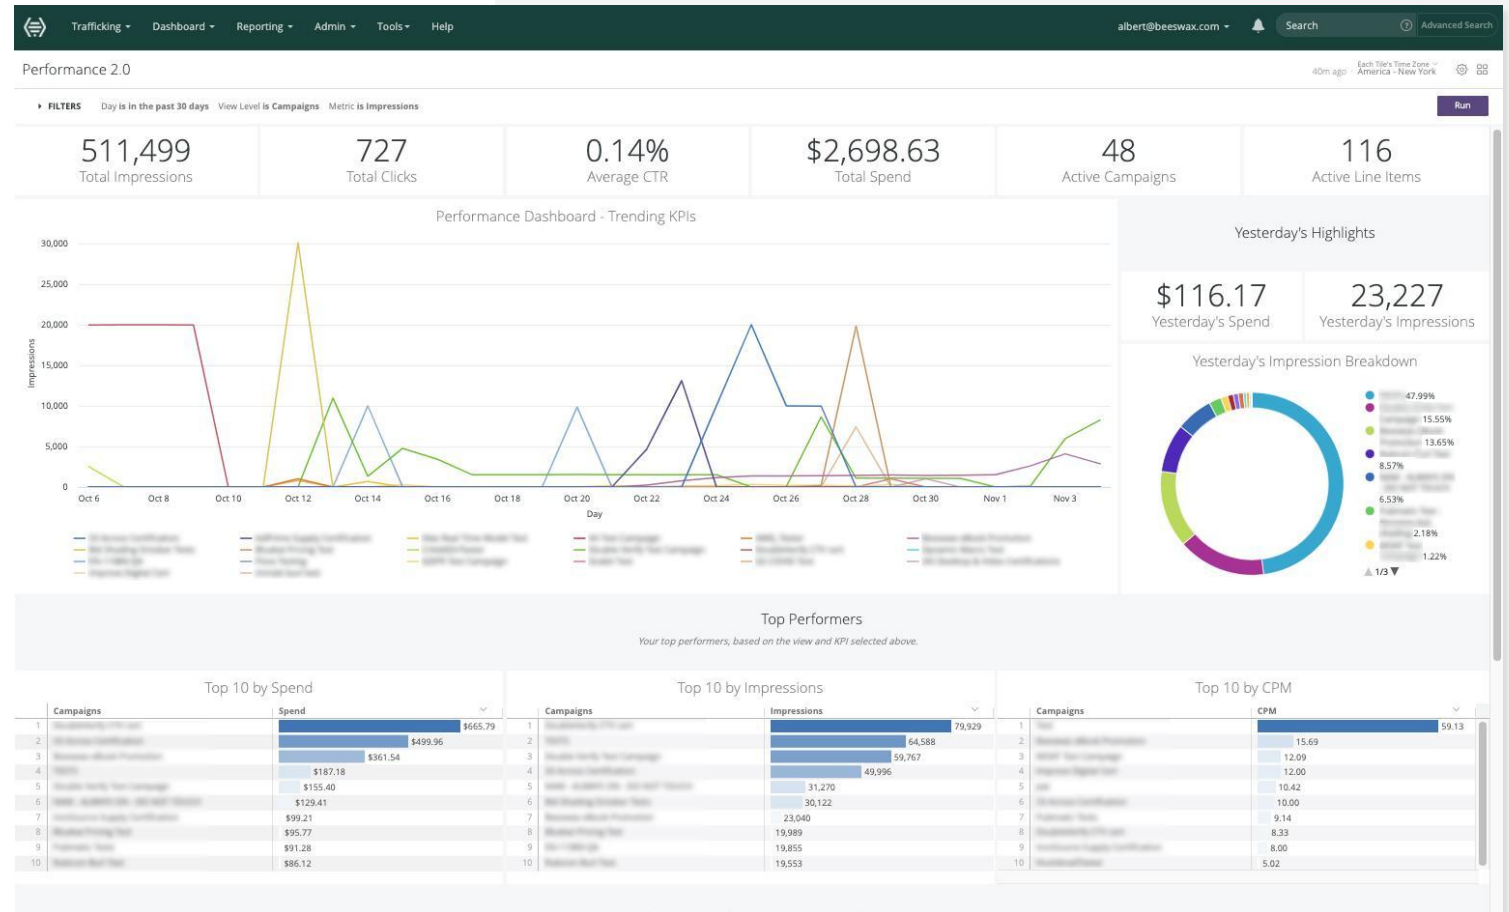
EVENT LOG SUMMARY

Different types of logs, format, delivery cadences

| CADENCE | LOG TYPE | DESCRIPTION | DELIVERY | FORMAT |
|------------------|----------------------------|---|-----------------|------------------------|
| Near real-time | Auctions | The auction request from exchange | S3 | Gzipped csv |
| Near real-time | Bids | Bids returned from Beeswax to exchange | S3 | Gzipped csv |
| Near real-time | Conversions | Conversion events recorded by Beeswax regardless of attribution | S3 | Gzipped csv or parquet |
| 8 hours delay | Attributed Conversions | Conversion events recorded by Beeswax | S3 | Gzipped csv or parquet |
| Hourly and Daily | Wins | The winning auctions, impressions, clicks and video activities | S3 | Gzipped csv or parquet |
| Realtime | Wins (Joined and Unjoined) | Stream of events associated with impression | HTTP or Kinesis | Protobuf or Json |
| 1 hour delay | Losses | Loss logs provided by limited set of exchanges | S3 | Gzipped csv |

REPORTING PLATFORM


- Intuitive data visualizations
- Customizable
- Cross-channel reporting



REPORTING PLATFORM

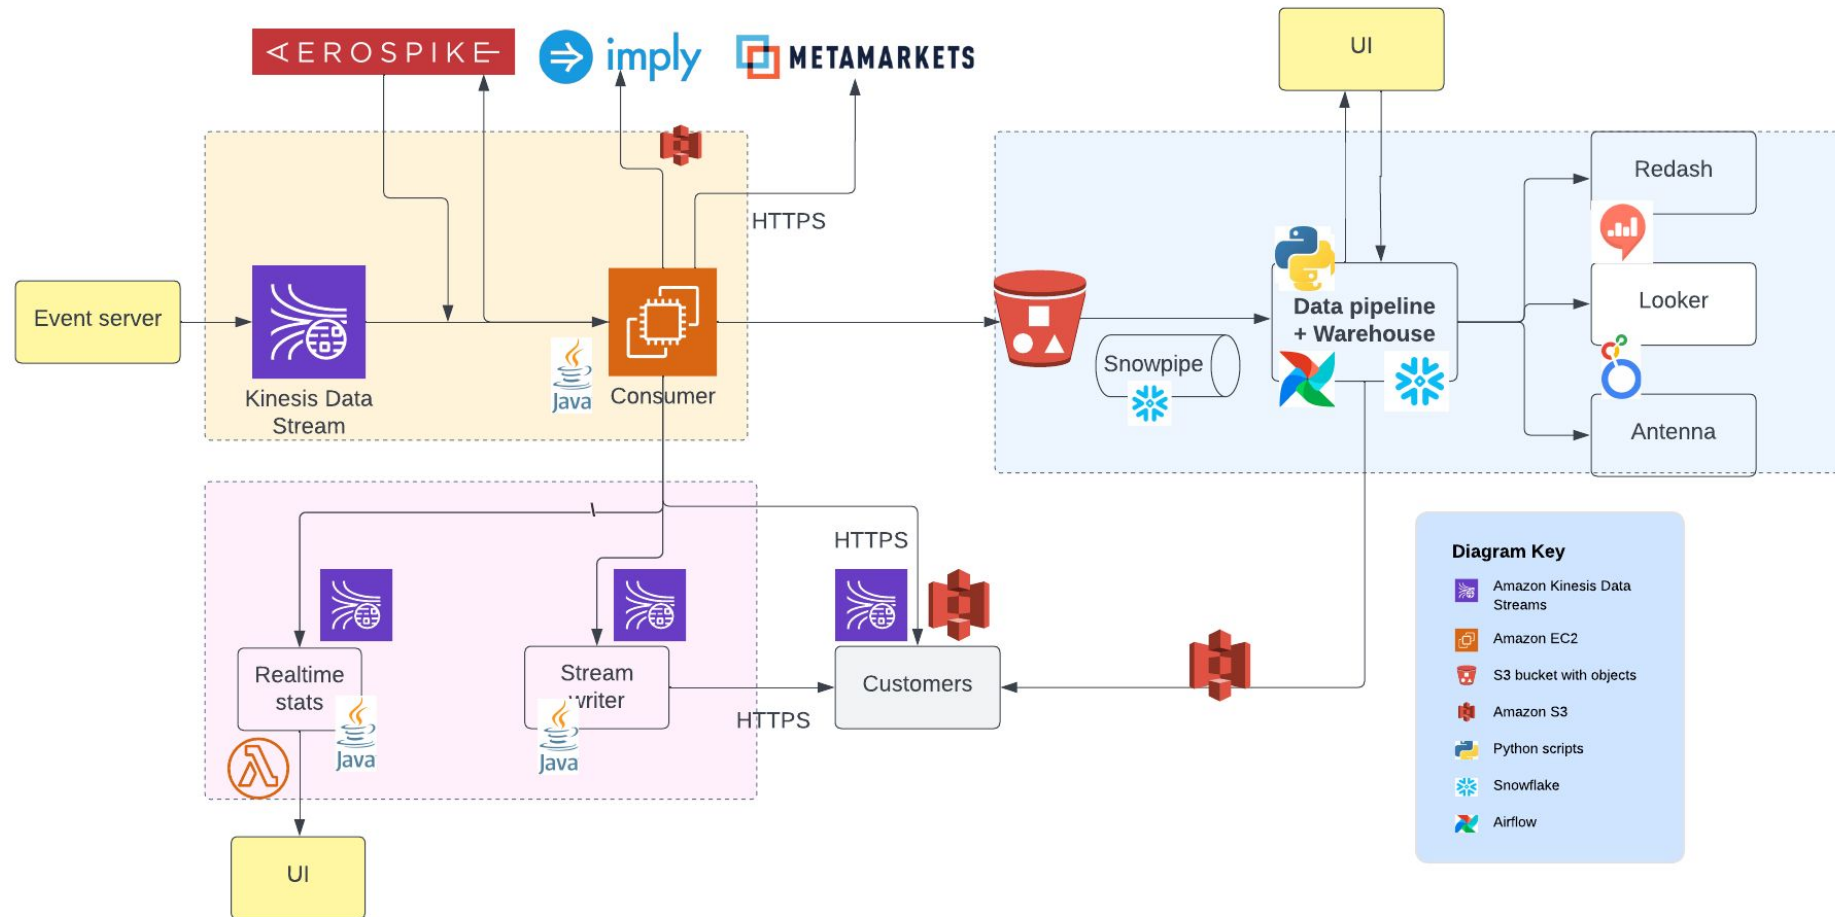
BID DECISION FLOW:

- Do we have enough supply?
- Are we winning enough?
- Are our bids priced right?
- How to improve my bidding strategy?



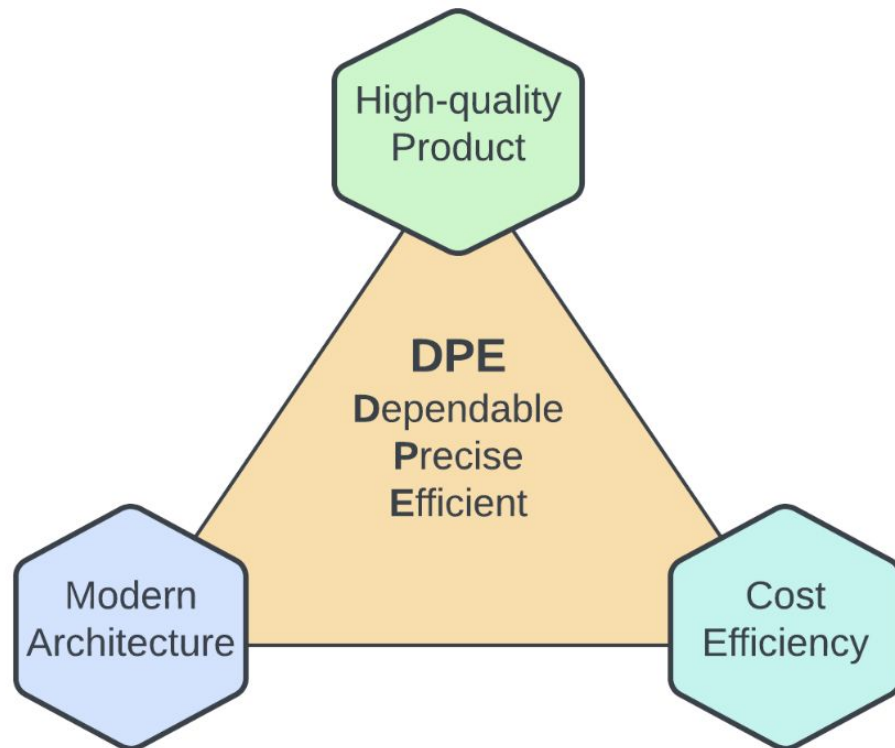
DATA ARCHITECTURE

ARCHITECTURE



PHILOSOPHY BEHIND THE TECHNOLOGY CHOICES

GUIDING PRINCIPLES



SCALABILITY & EFFICIENCY

Should be able to scale with growth without incurring major cost.

DEFENSIBLE DIFFERENTIATOR

Support customer needs

BUY TECH AND BUILD ADTECH

Focus engineering bandwidth on the differentiators.

TIME TO MARKET

TECHNOLOGY CHOICES

AWS KINESIS



- Fully managed service
- Easy to setup

AEROSPIKE



- Fast real-time NoSQL database
- Exceptional performance & high availability

SNOWFLAKE



- Internal warehouse
- Secure sharing
- Distinguishes compute from storage

LOOKER



- Analytics and reporting platform.
- Embedded in Buzz App UI.
- Easy to build reports

REDASH

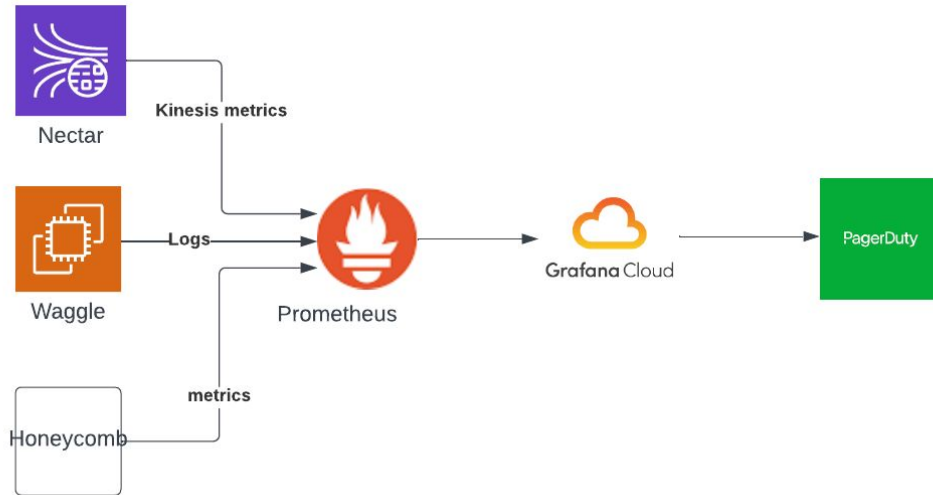


- Democratizing access to data
- Data query & visualization tool

A long, illuminated arch bridge spans a body of water at night. The bridge's arches are brightly lit from within, casting a warm glow. The sky above is dark and filled with stars, with the Milky Way galaxy clearly visible as a bright, hazy band of light. The bridge's reflection is visible in the calm water below. The overall scene is serene and majestic.

MONITORING & ALERTING


MONITORING & ALERTING



Alerts:

- Aging messages on Kinesis stream
- Capacity alerts
- Pipeline failure alerts

- 6.5 K metrics
- 8K GB logs
- 50+ alerts configured



ROADMAP

ROADMAP

- Increase scalability, performance & cost effectiveness
 - Bring the consumer applications onto Databricks Spark Streaming
 - Rebuild real-time bidding stats using Spark for faster performance
- Improve data quality
 - Great expectation
- Improve data observability
 - data cataloging and lineage



Building a solid data platform is
more about using the right
technology and focusing on
the differentiator for each use
case and having a right sized
team to support it.

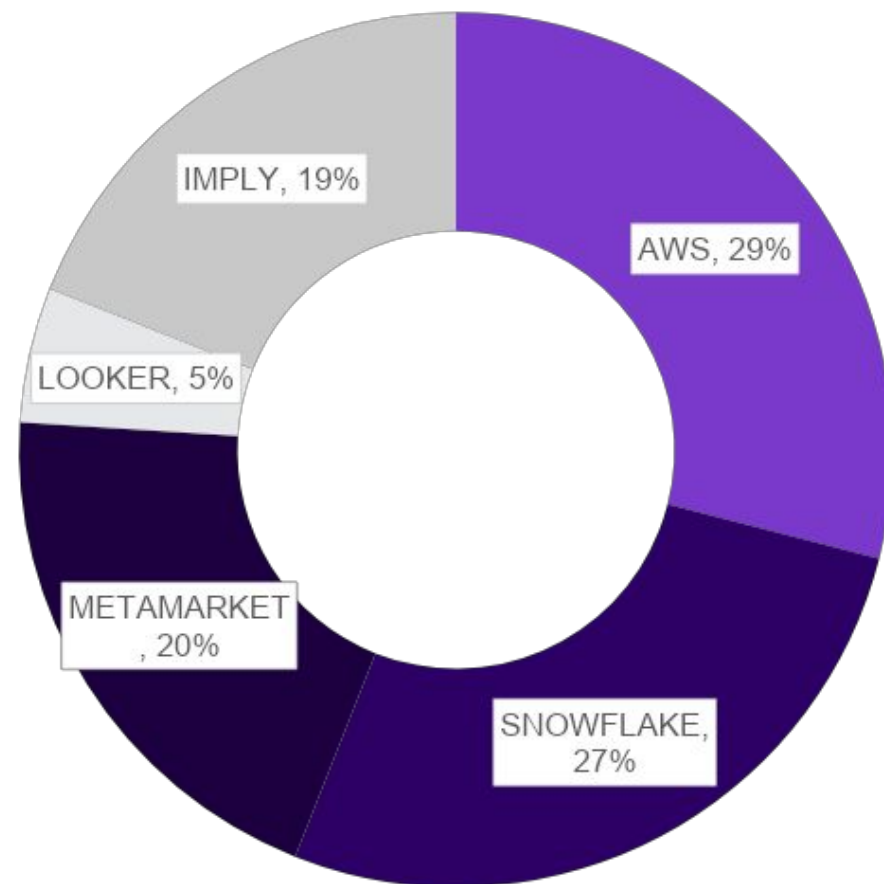


THANK YOU

FREOWHEEL
A COMCAST COMPANY

COST BREAKDOWN

■ AWS ■ SNOWFLAKE ■ METAMARKET ■ LOOKER ■ IMPLY



AGENDA

Table of Contents

1 Who we are?

2 Basics of bid request

3 Infrastructure Highlights

4 Data Architecture

5 Philosophy behind technology choices

6 Monitoring & Alerting

7 Cost Breakdown & Roadmap

8 Questions