



Gravy Analytics & Snowflake

Guy DeCorte, Founder & CTO

March 22, 2018

REAL-WORLD CONSUMER BEHAVIOR

Where we go is who we are.

The **events** consumers attend,
the **places they visit**,
where they **spend their time**,
translates into intelligence



INTERESTS



LIFE STAGES



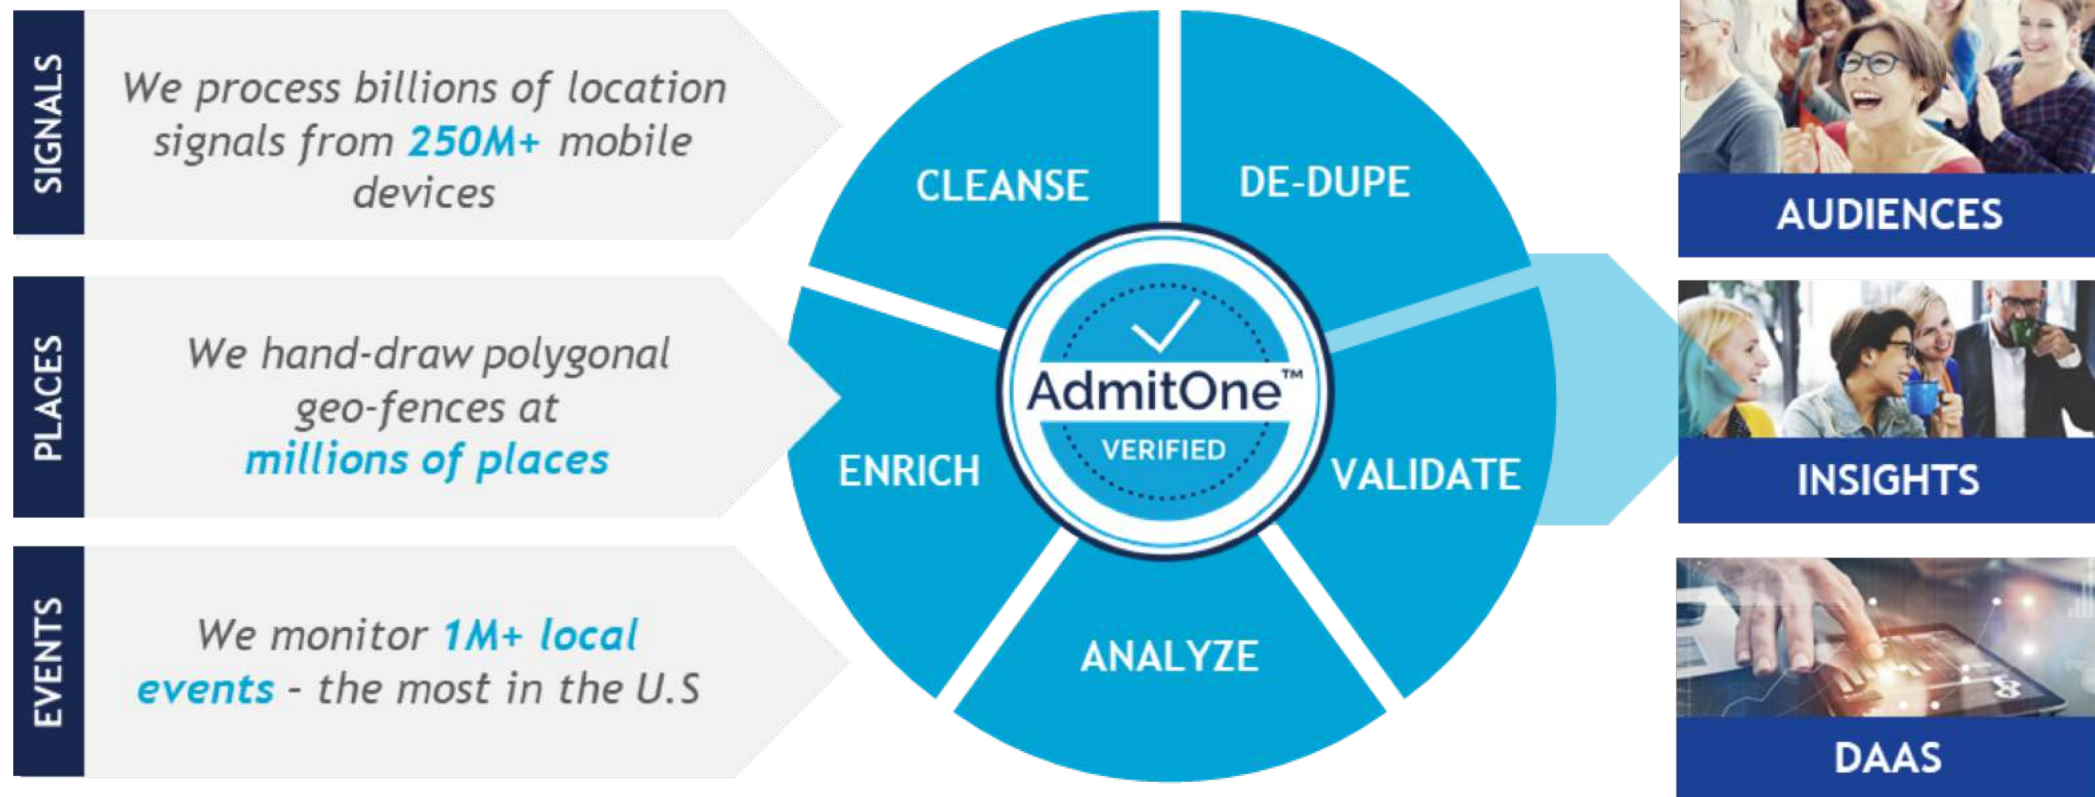
AFFINITIES



LIFESTYLES

INDUSTRY-LEADING CAPABILITIES

We translate the locations that consumers visit, the places they go, and the events they attend into real-world consumer intelligence



GRAVY SOLUTIONS

GRAVY AUDIENCES



Gravy Audiences let marketers reach engaged consumers based on what they do in real-life

- Lifestyle
- Enthusiast
- In-Market
- Branded
- Custom

GRAVY INSIGHTS



Gravy Insights provides brands with in-depth customer and competitive intelligence

- Foot Traffic
- Competitive
- Attribution

GRAVY DAAS



AdmitOne™ verified **Visitation, Attendance, Event** data and more for use in unique business applications

- Visitations
- Attendances
- Events
- Categories

THE GRAVY DIFFERENCE

REACH

Billions of daily location signals from **250M+ opted-in mobile devices**

EVENTS

The **largest events database** gives context to millions of places and POIs

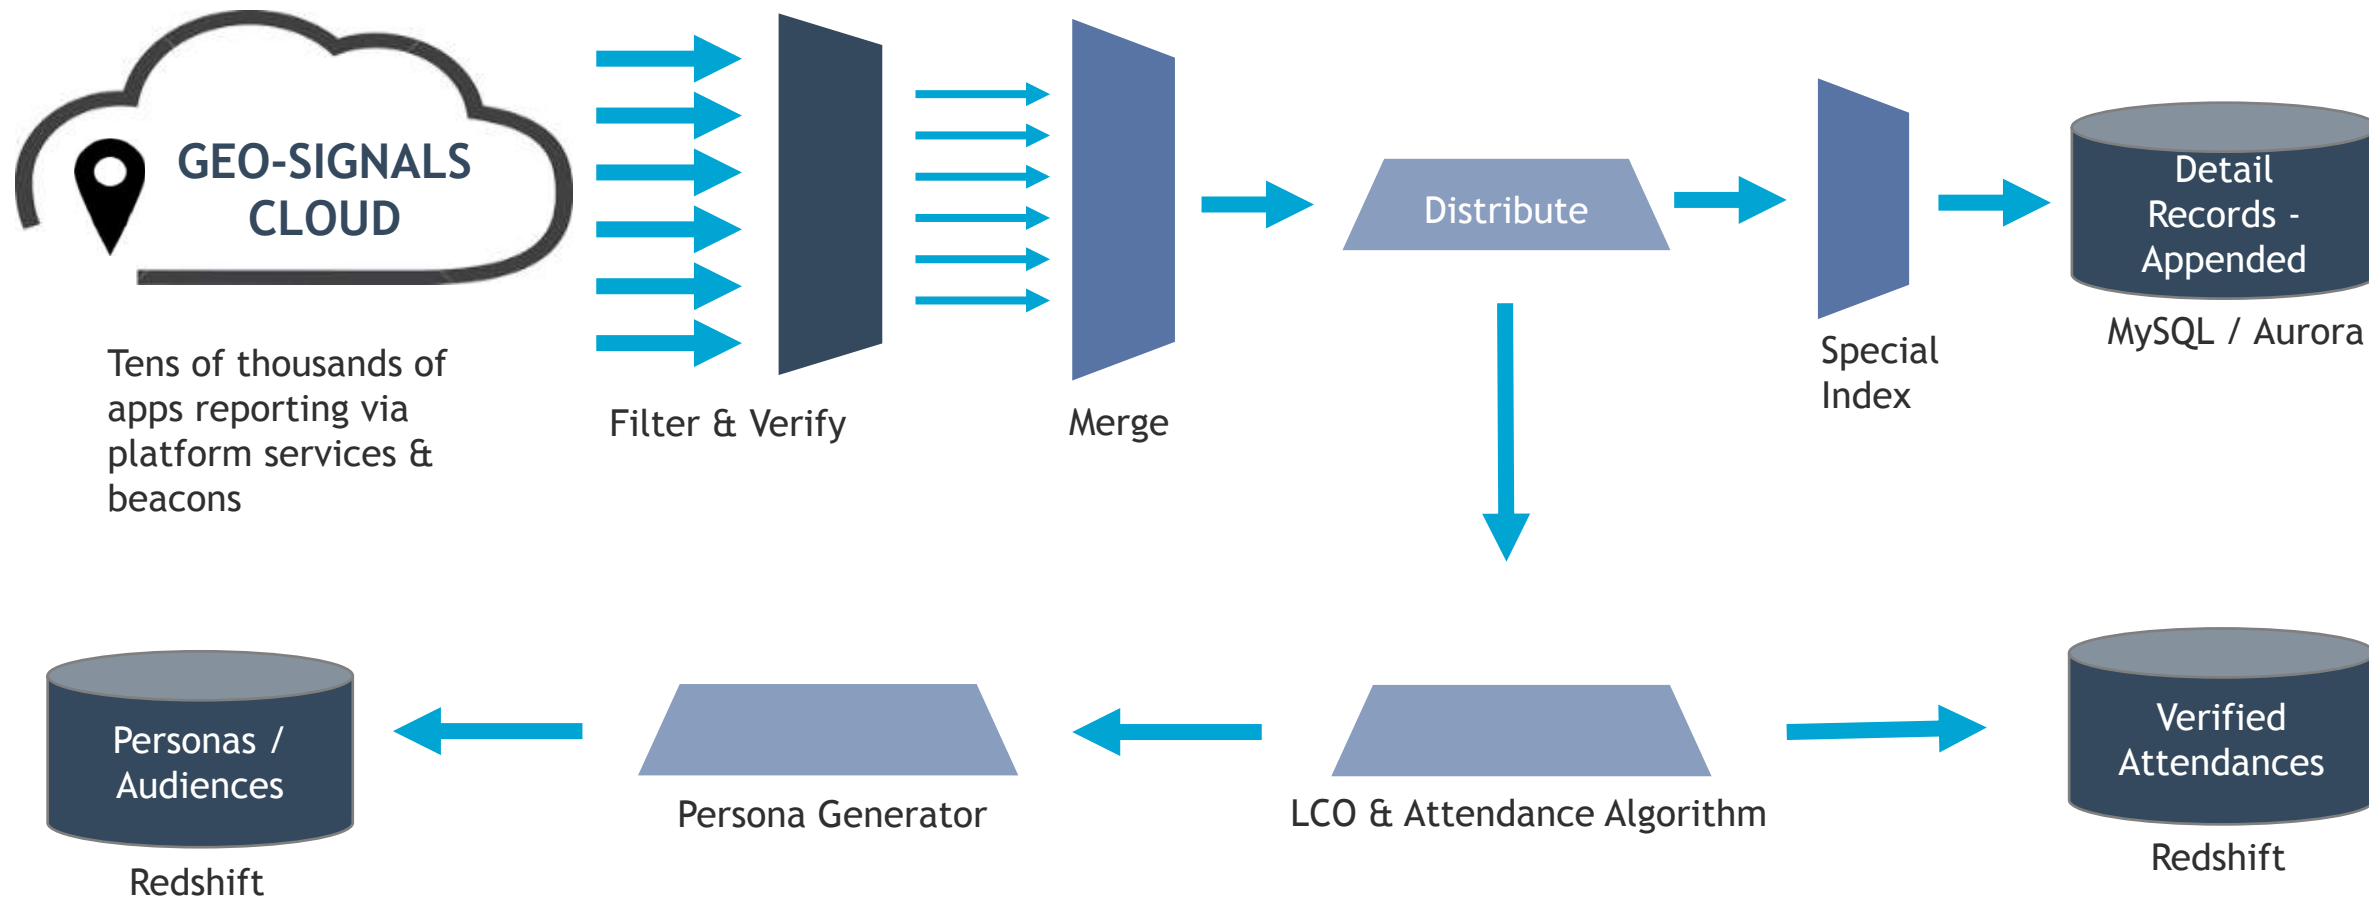
VERIFIED

Confirmed, deterministic consumer attendances at places and events.



Gravy's patented AdmitOne verification engine delivers the highest-quality location and attendance data in the industry

Processing Overview



CHALLENGES

1. Analytics on large data sets - billions of records

- **Initial Approach**

- Redshift - multi-node cluster

- **Issues**

- Constantly running out of space
- Slow performance
- Loss of team time

2. Loading and querying data sets - 300B-1T+ records

- **Initial Approach**

- Highly appended + compressed

- **Issues**

- Limited to query on keys
- Slow performance
- Data scalability
- Brittle loads

APPROACH

- **Objectives**

- No team time lost to data storage management
- No need to schedule load times separate from work times
- Faster load & query times
- Ability to store detail records un-appended

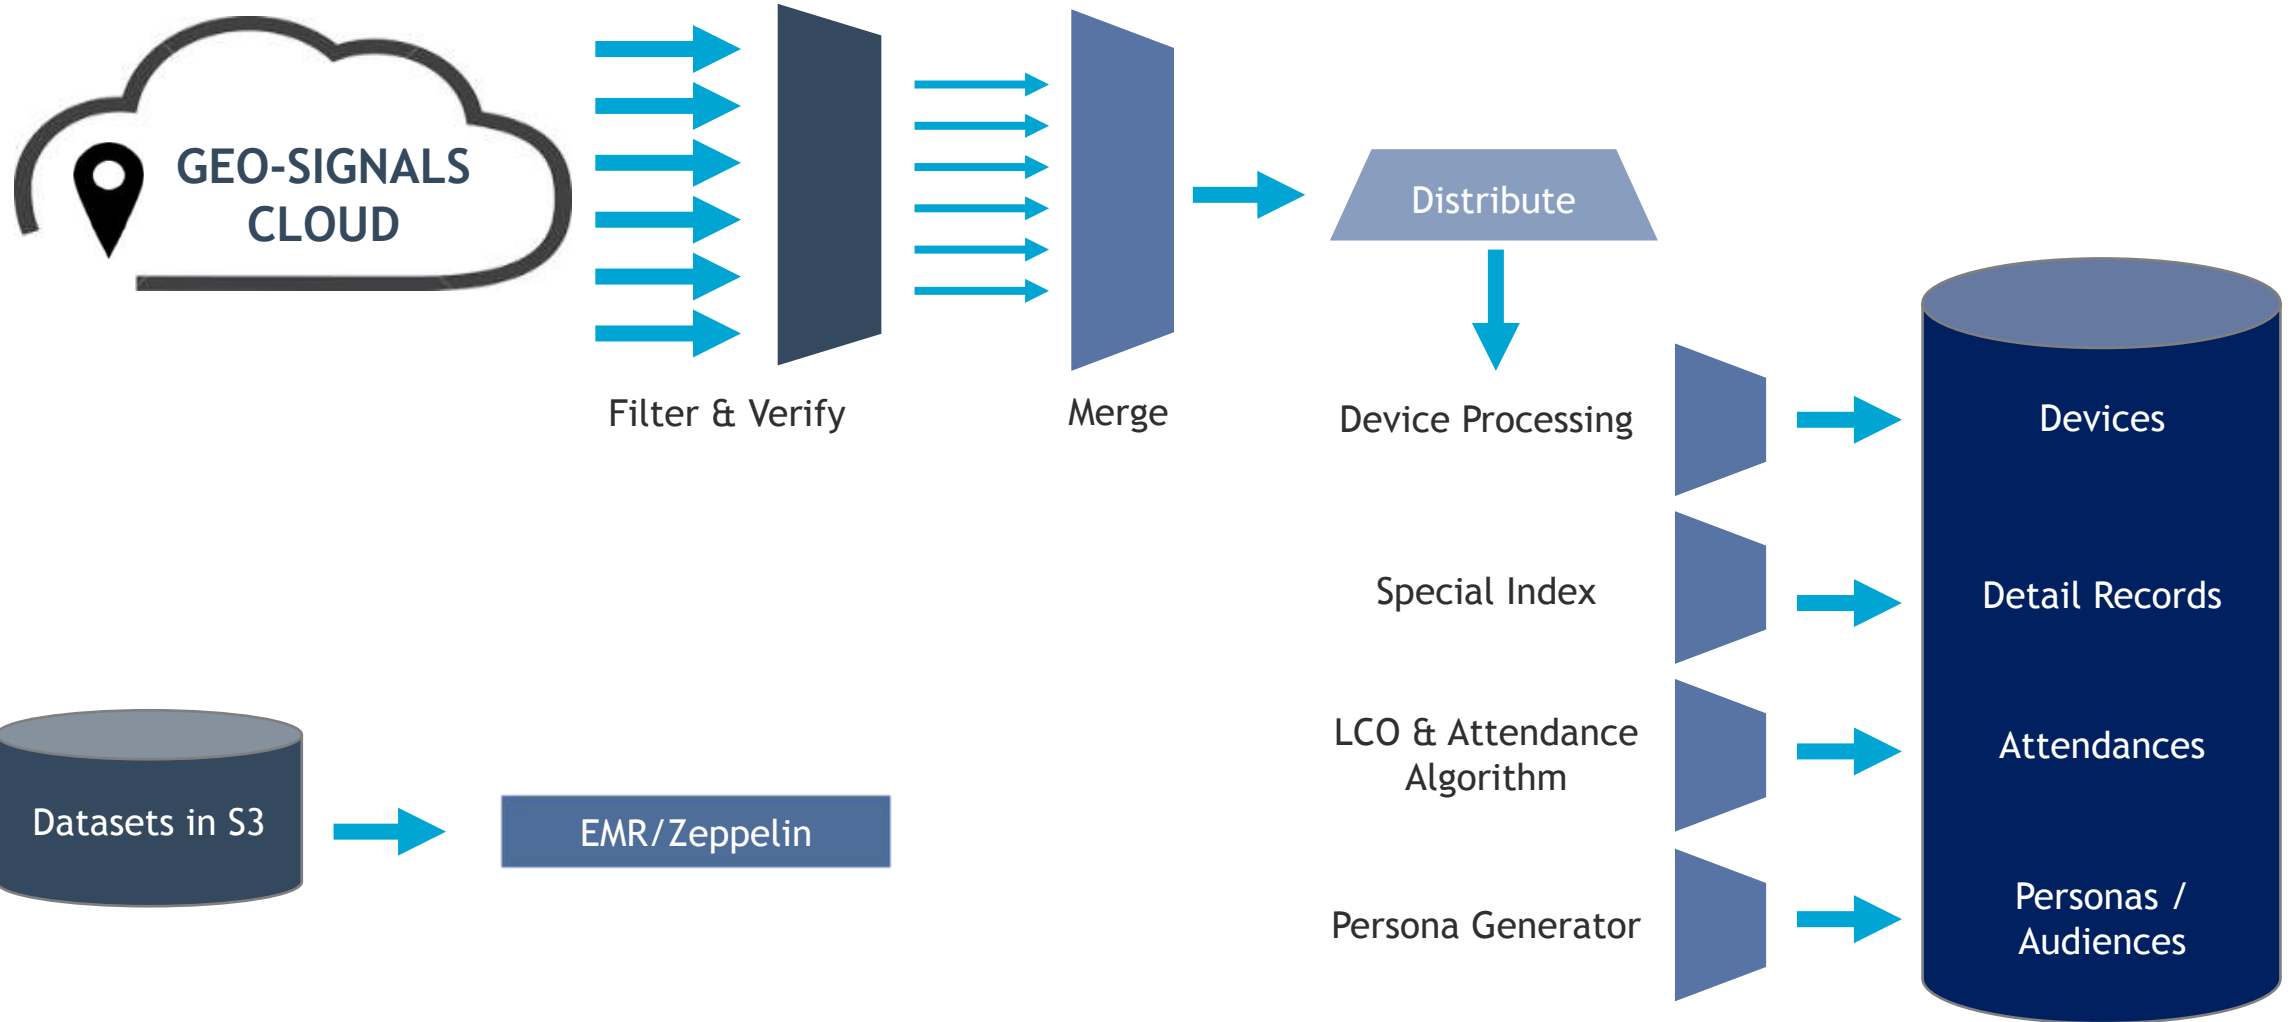
- **Decision Process**

- Evaluated & tested multiple products
- Pleased to find a single product to meet both challenges



- Redshift cluster
- Cassandra
- MySQL Cluster
- Redis clusters
- Snowflake

SOLUTION



RESULTS

- **Snowflake's architecture separates compute and storage...**
 - Faster load times - load and query efforts don't conflict
 - No conflict = no need for shift work
 - Faster, more efficient data queries
 - Lower storage costs enable 2-3x data storage
 - Better support for internal, external customers

COST SAVINGS



- **~25-30% cost savings** across both challenge areas
- **~10% time savings** across 3 team members (~48 hrs./mo.)

WHAT WE'D LIKE TO SEE



- **Explain Command/Cost Estimator**
 - For complex queries - understand query performance and costs upfront
 - Recommendations for clustering/organization of tables to help query performance
- **Queries that take advantage of cluster keys are fast; performance varies for others**
 - More options to increase performance - without duplicating data
 - Additional documentation on query performance

Guy DeCorte, CTO

703-840-8850

gdecorte@gravyanalytics.com



Yard House Happy Hour | 4:45 pm



Art in the Park | 3:30 pm



Monster Ja



8th Wonder Comedy Hour | 9:30 pm