Extending the Enterprise Data Warehouse with Hadoop

Robert Lancaster

Nov 7, 2012

ORBITZ

Who I Am

- Robert Lancaster
  - Solutions Architect, Hotel Supply Team
  - rlancaster@orbitz.com
  - @rob1lancaster
  - Organizer of Chicago Machine Learning Study Group
  - Co-organizer of Chicago Big Data.

ORBITZ

Launched in 2001

Over 160 million bookings

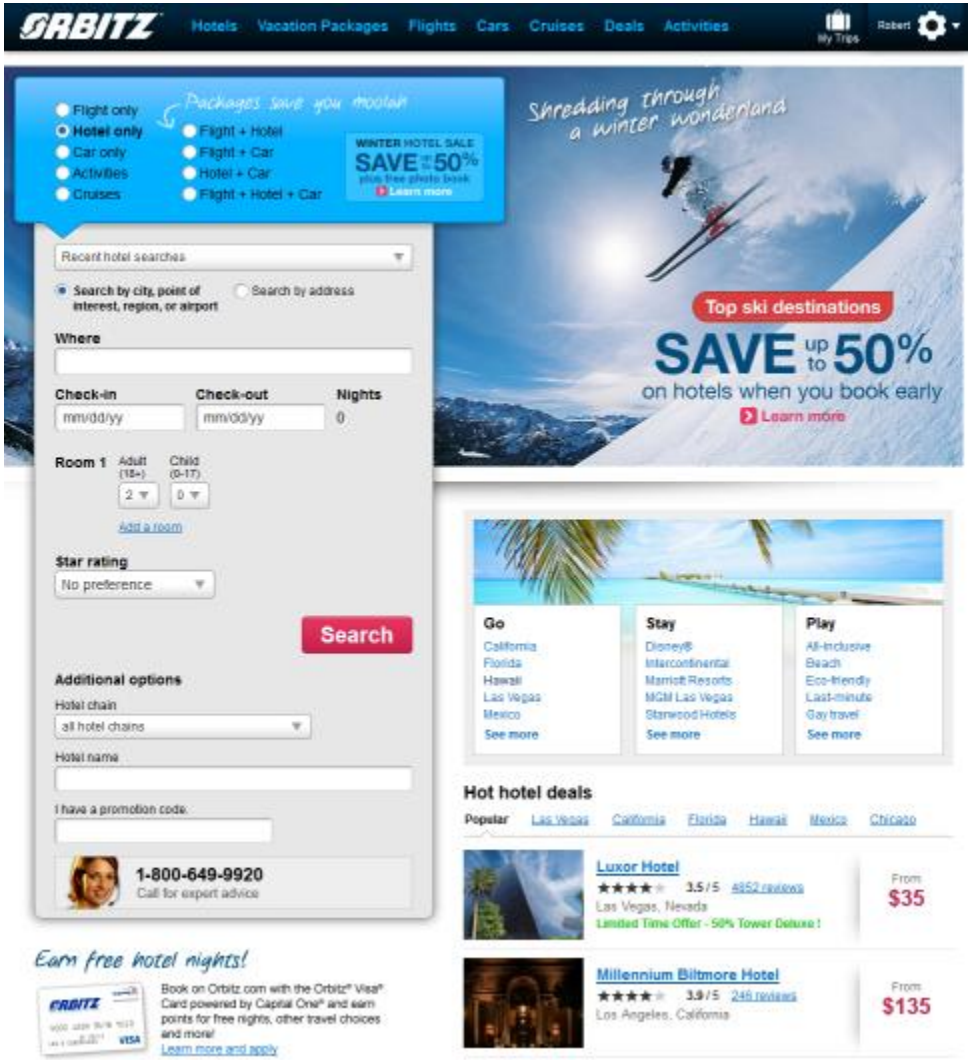

Some History…

San Francisco, CA, United States | Check-in 11/20/12 | Check-out 11/23/12 | Change search | Need help? 1-800-733-1297 (toll free)

Nights: 3 | Rooms: 1 | Guests: 1

WHY ORBITZ? Guaranteed low prices • Easy-to-use planning tools › Learn more

Map Expand map

Clear Selections

Hotel name

Areas

Our most booked areas:

- Any
- Union Square-Nob Hill (78)
- Financial District - South of Market (16)
- Fisherman's Wharf (18)
- Pacific Heights-Marina (19)
- Civic Center (10)
- San Francisco Airport (61)

Any

- ★★★★★ (4)
- ★★★★☆ (40)
- ★★★☆☆ (116)
- ★★☆☆☆ (121)
- ★☆☆☆☆ (13)

Reviewer score

1.0 to 5.0

294 matching hotels Previous | Next

Sort by: Best Values | Lowest Price | Distance | Star Rating | Reviewer Score

Parc 55 Wyndham San Francisco - Union Square

★★★★☆ 4.0 / 5 286 reviews

Union Square-Nob Hill

32 stories of luxury and style within steps of Union Square

This luxurious 32-story hotel is situated less than two blocks from shopping in San Francisco's... More

There are 3 people viewing this hotel.

Last booked 4 hours ago.

$112
Avg/night
✓ Price Assurance
Book

Save 20% on your stay

Omni San Francisco Hotel

★★★★☆ 4.6 / 5 152 reviews

Union Square-Nob Hill

Old and new pair with elegance and prime steakhouse up on Nob Hill

The past is new again at this circa-1926 hotel atop Nob Hill. Originally built as a bank, the... More

Last booked 4 hours ago.

~~$199~~
$159
Avg/night
✓ Price Assurance
Book

[+] Site Feedback

★★★★☆ 4.2 / 5 171 reviews

Union Square-Nob Hill

Urban elegance with Asian flair and an indoor pool near Union...

Hotel Nikko feels more like Hong Kong than San Francisco. The lobby is spacious and stark, and... More

There are 2 people viewing this hotel.

Last booked 3 hours ago.

~~$169~~
$143
Avg/night
✓ Price Assurance
Book

In 2009…

- The Machine Learning team is formed to improve site performance. For example, improving hotel search results.
- This required access to large volumes of behavioral data for analysis.
  - Fortunately, the required data was collected in session data stored in web analytics logs.

ORBITZ

The Problem…

- The only archive of the required data went back about two weeks.

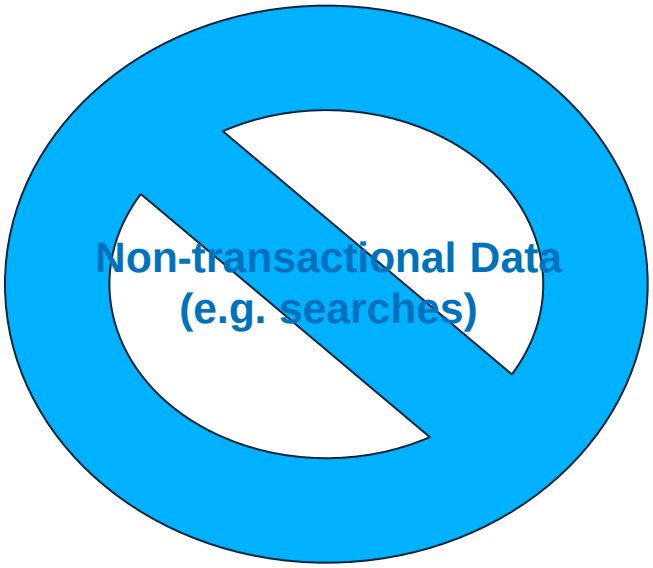

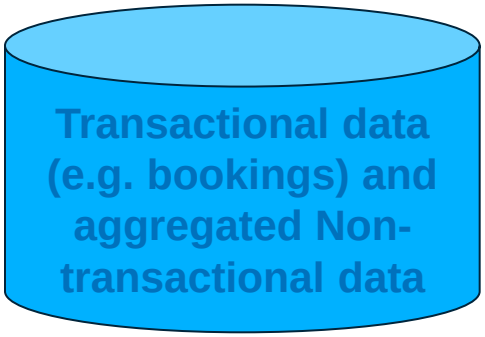

Data Warehouse

Hadoop Provided a Solution…

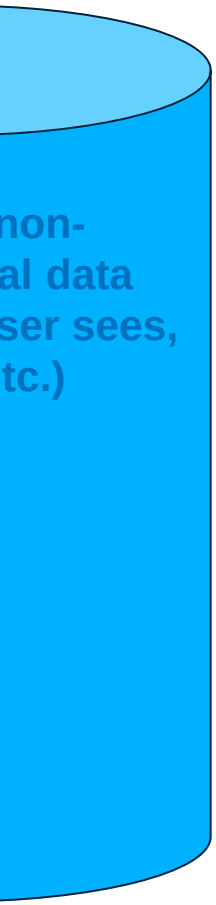

Hadoop

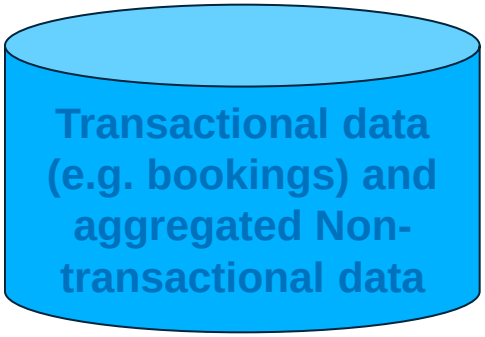

Data Warehouse

What is Hadoop?

- Distributed file system and parallel processing platform.
- Open source Apache project created by Doug Cutting.
- Modeled on papers published by Google on the Google File System and MapReduce.
- Intended to run on a cluster of relatively inexpensive machines (aka commodity hardware).
- Bring processing to the data.

ORBITZ

The Hadoop Ecosystem

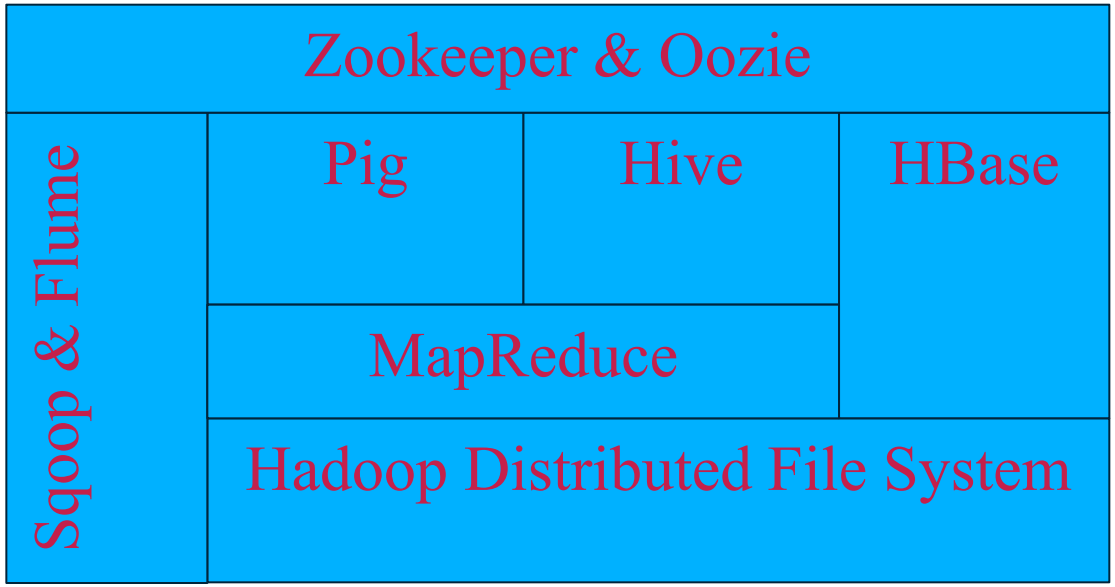

ORBITZ

Deploying Hadoop Enabled Multiple Applications…

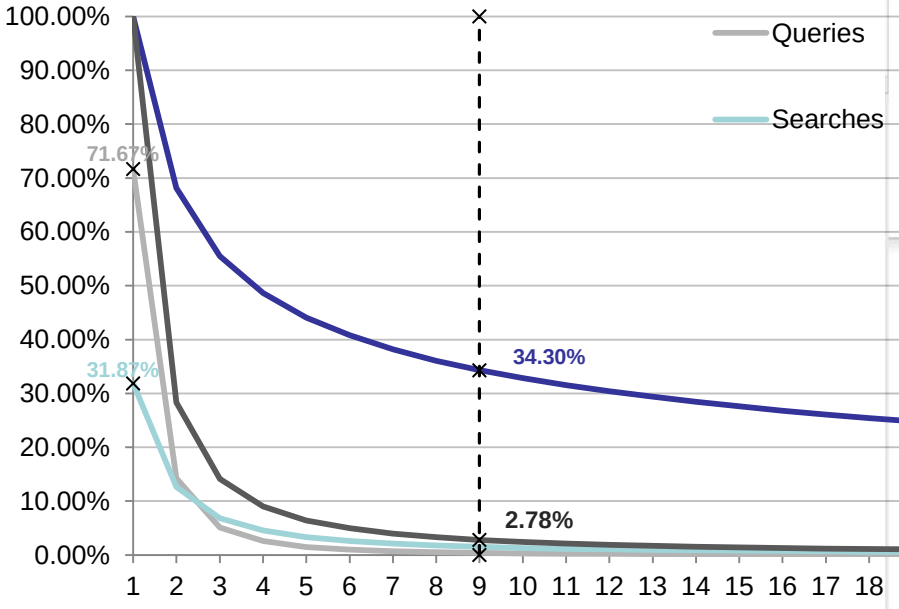

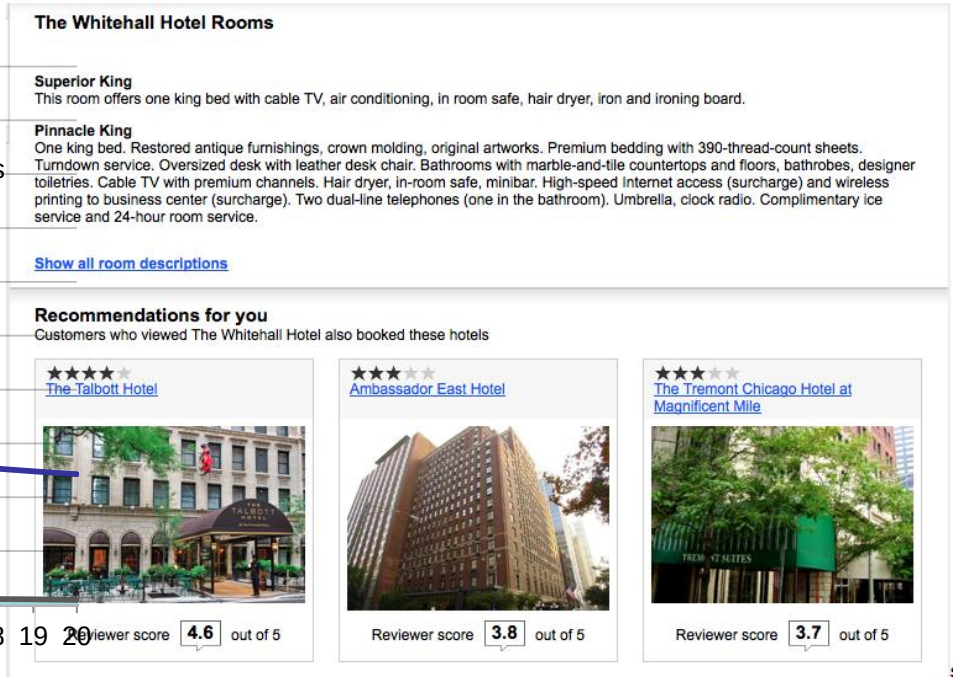

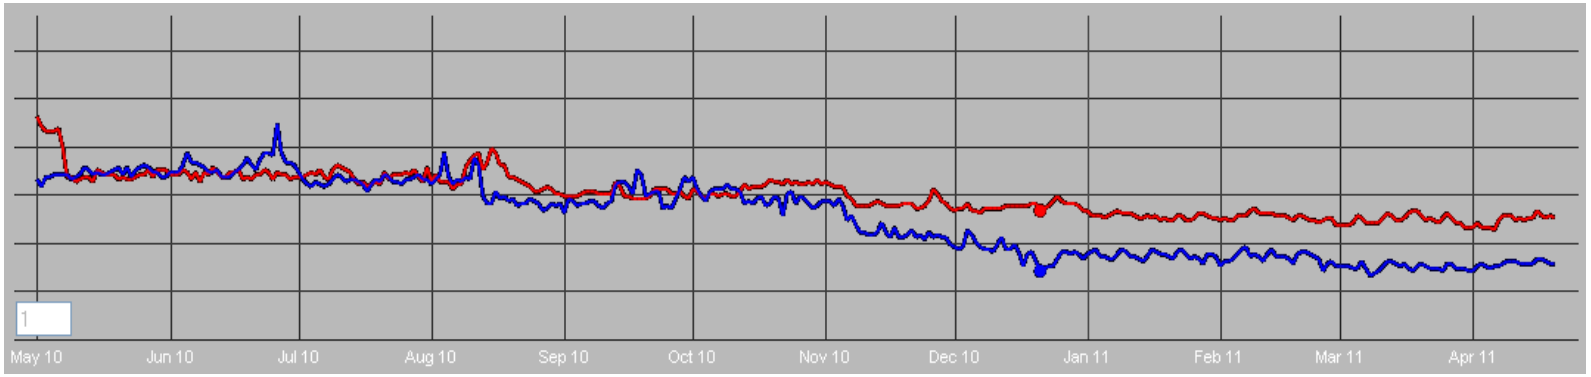

And Useful Analyses…

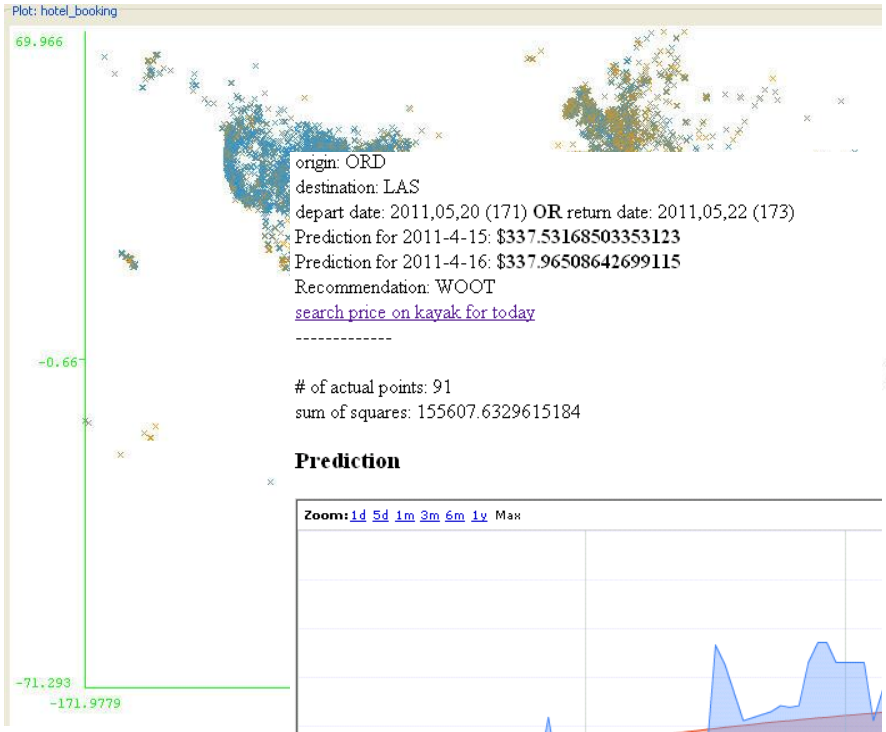

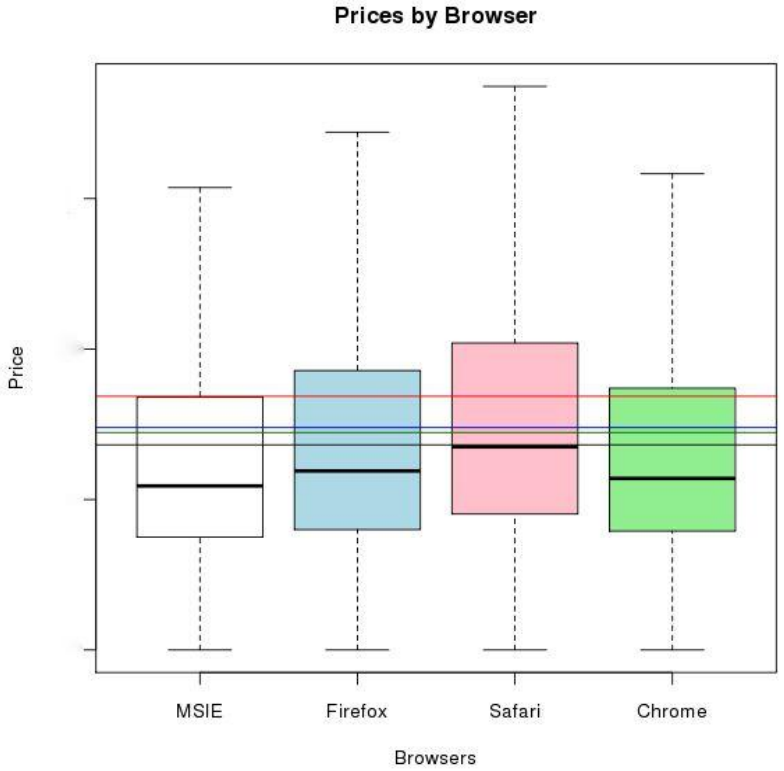

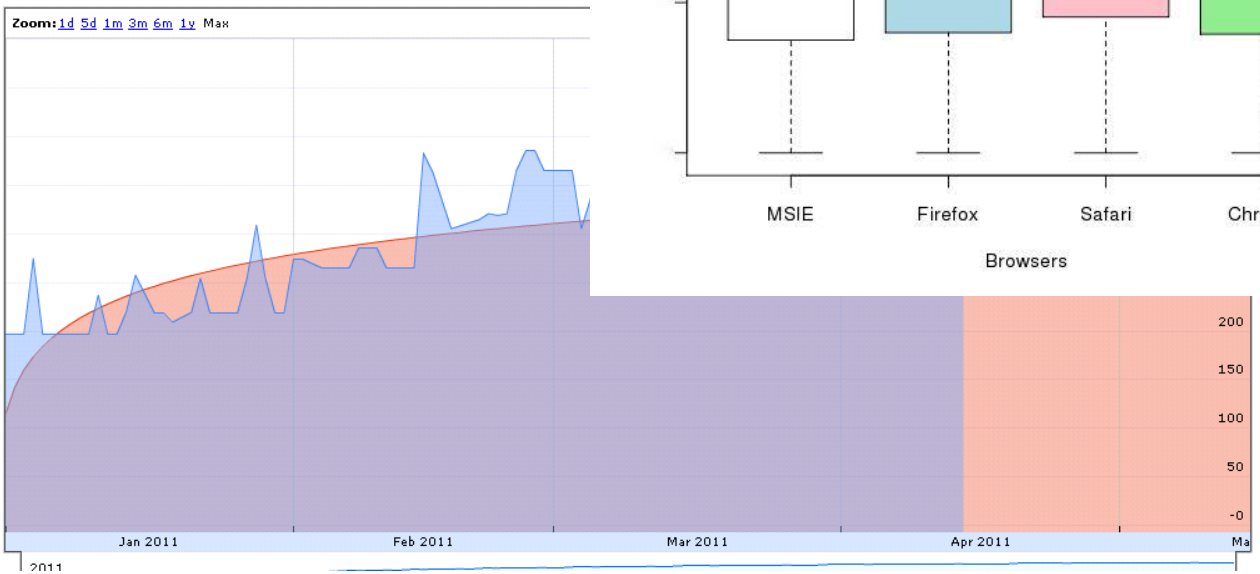

But Brought New Challenges…

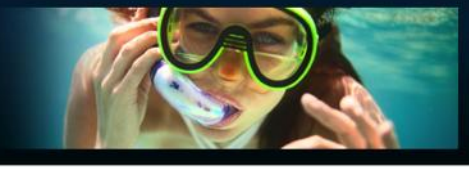

- Most of these efforts are driven by development teams.
- The challenge now is unlocking the value of this data for non-technical users.
- Support for Hadoop via traditional BI/reporting tools still meager.

ORBITZ

BI Vendors Are Working on Hadoop Integration

Both big (relatively)…

INFORMATICA
The Data Integration Company™

MicroStrategy

QlikView

GREENPLUM®

ORBITZ®

And small…

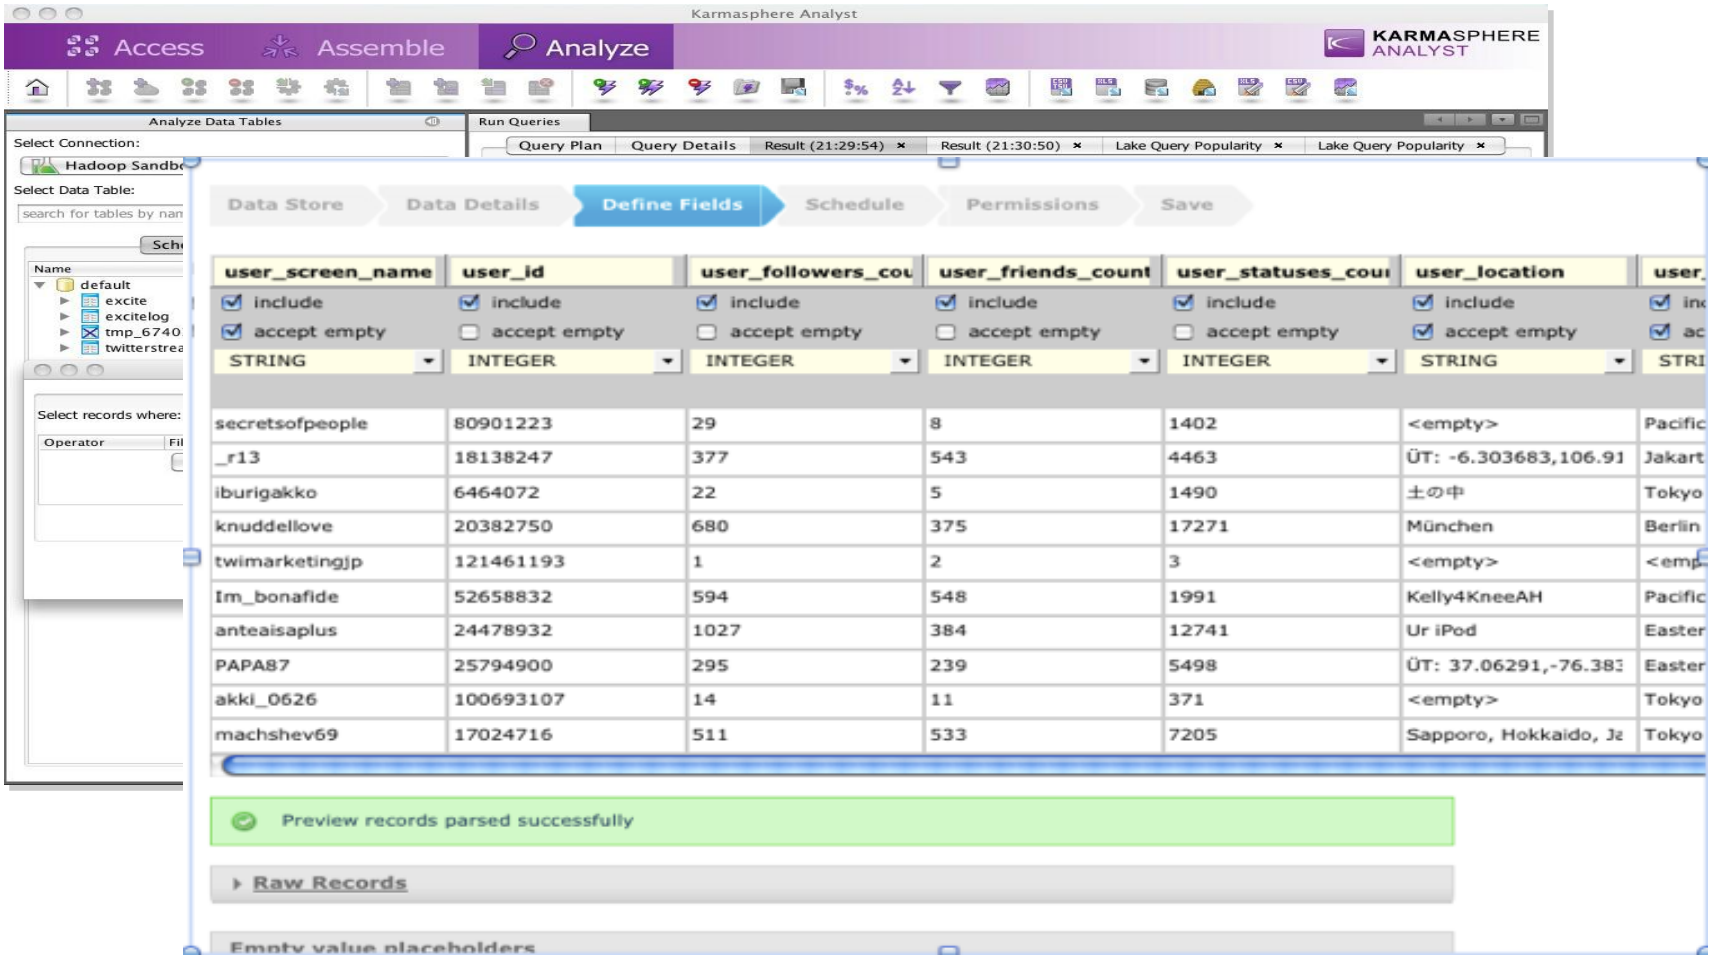
Karmasphere Analyst
Access
Assemble
Analyze
KARMASPHERE ANALYST
Analyze Data Tables
Select Connection:
Hadoop Sandbo
Select Data Table:
search for tables by nan
Name
default
excite
excitelog
tmp_6740
twitterstrea
Select records where:
Operator
Run Queries
Query Plan
Query Details
Result (21:29:54)
Result (21:30:50)
Lake Query Popularity
Lake Query Popularity
Data Store
Data Details
Define Fields
Schedule
Permissions
Save
user_screen_name
user_id
user_followers_cou
user_friends_count
user_statuses_cou
user_location
include
accept empty
STRING
INTEGER
INTEGER
INTEGER
INTEGER
STRING
secretsofpeople 80901223 29 8 1402 <empty> Pacific
_r13 18138247 377 543 4463 ÜT: -6.303683,106.91 Jakart
iburigakko 6464072 22 5 1490 土の中 Tokyo
knuddellove 20382750 680 375 17271 München Berlin
twimarketingjp 121461193 1 2 3 <empty>
Im_bonafide 52658832 594 548 1991 Kelly4KneeAH Pacific
anteaisaplus 24478932 1027 384 12741 Ur iPod Easter
PAPA87 25794900 295 239 5498 ÜT: 37.06291,-76.383 Easter
akki_0626 100693107 14 11 371 <empty> Tokyo
machshev69 17024716 511 533 7205 Sapporo, Hokkaido, Ja Tokyo
Preview records parsed successfully
Raw Records
Empty value placeholders

In 2011& 2012

- Big Data team is formed under Business Intelligence team at Orbitz Worldwide.
- Allows the Big Data team to work more closely with the data warehouse and BI teams.
- Reflects the importance of big data to the future of the company.
- Our production cluster has grown 40-fold since it was launched.

A View Shared Beyond Orbitz…

"We strongly believe that Hadoop is the nucleus of the next-generation cloud EDW…"

"…but that promise is still three to five years from fruition."*

*James Kobielus, Forrester Research,
"Hadoop, Is It Soup Yet?"

Two Primary Ways We Use Hadoop to Complement the EDW

- Extraction and transformation of data for loading into the data warehouse – “ETL”.
- Off-loading of analysis from the data warehouse.

ORBITZ

ETL Example

Proposed Processing

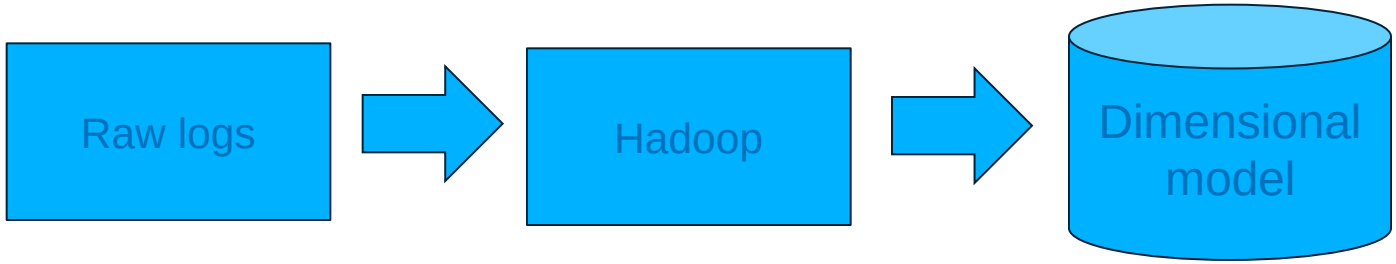

ETL Example: Click Data Processing

Previous Processing in Data Warehouse

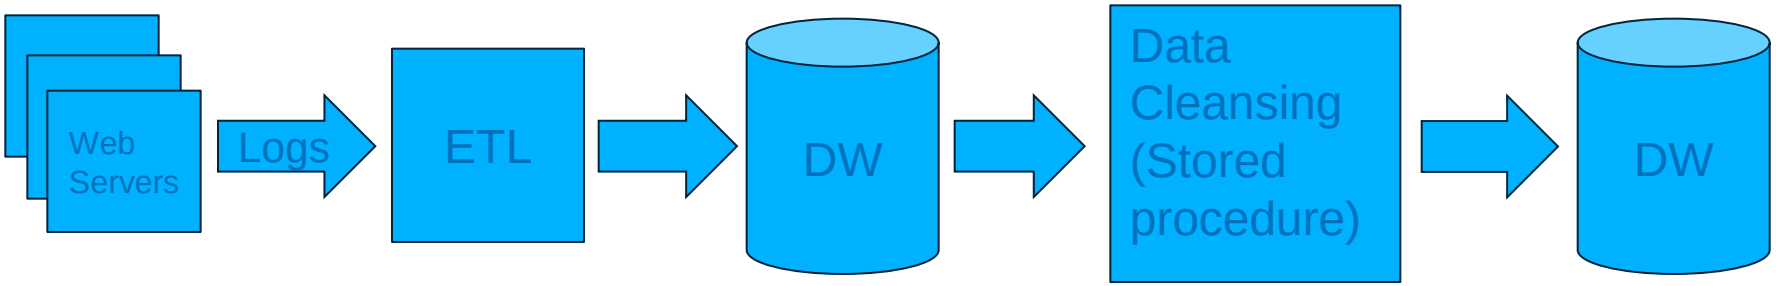

Several hours of processing

~20% original data size

ORBITZ

ETL Example: Click Data Processing

- Moving to Hadoop:
  - Removed load from the data warehouse.
  - Facilitated adding additional attributes for processing.
  - Allowed processing to be run more frequently.

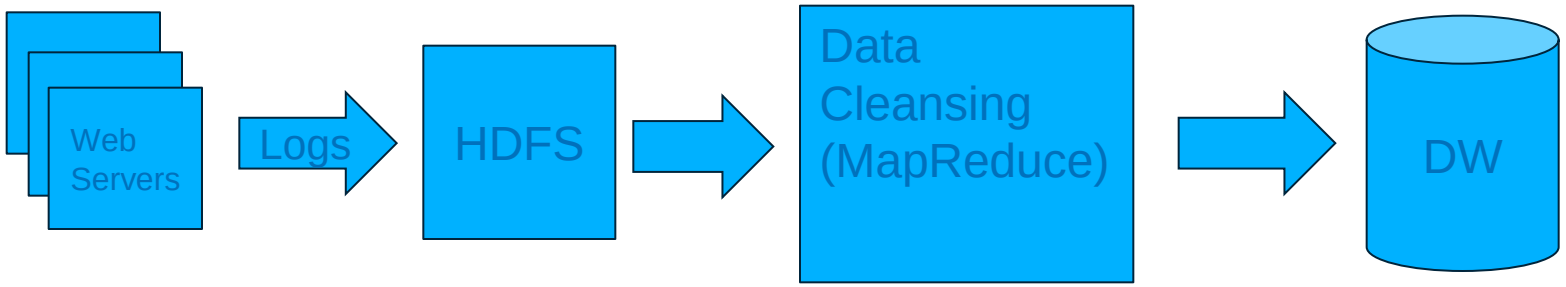

Processing in Hadoop

ORBITZ

Analysis Example: Geo-Targeting Ads

- Facilitated analysis that allows for more personalized ad content.
- Allowed marketing team to analyze over a years worth of search data.
- Provided analysis that was difficult to perform in the data warehouse.

Example Processing Pipeline for Web Analytics Data

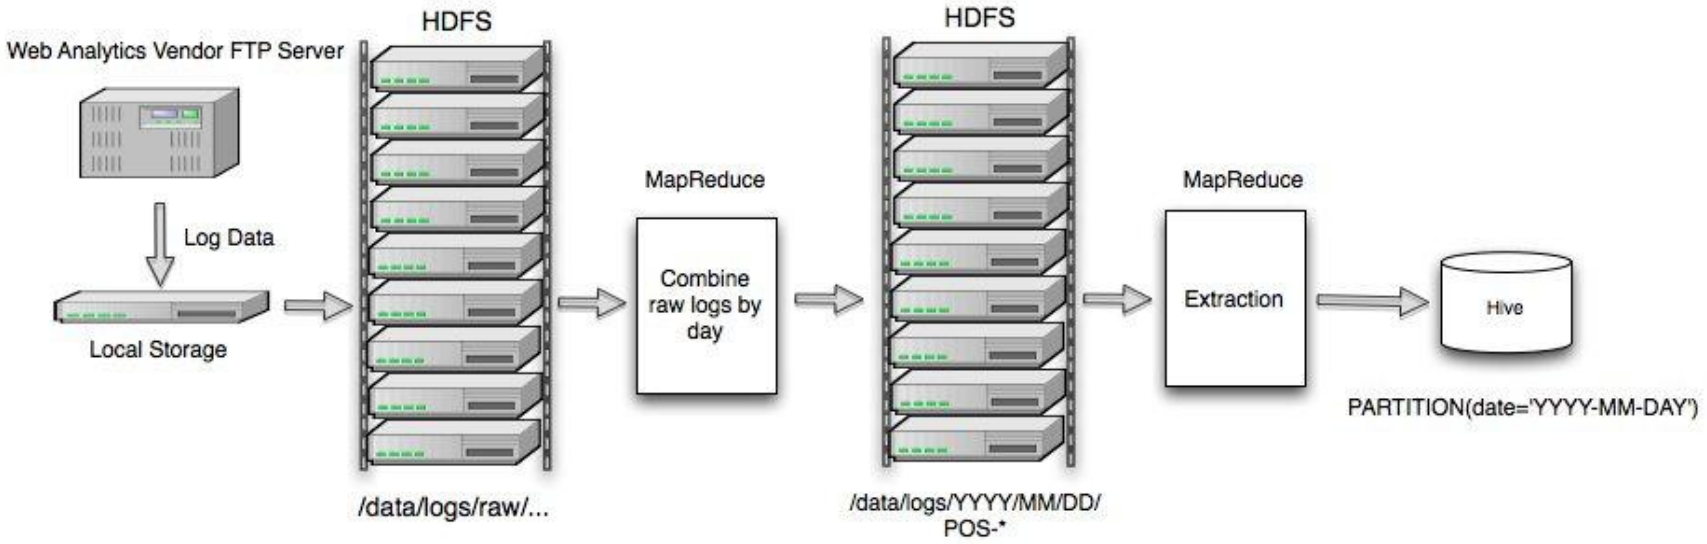

Example Use Case: Selection Errors

Use Case – Selection Errors: Introduction

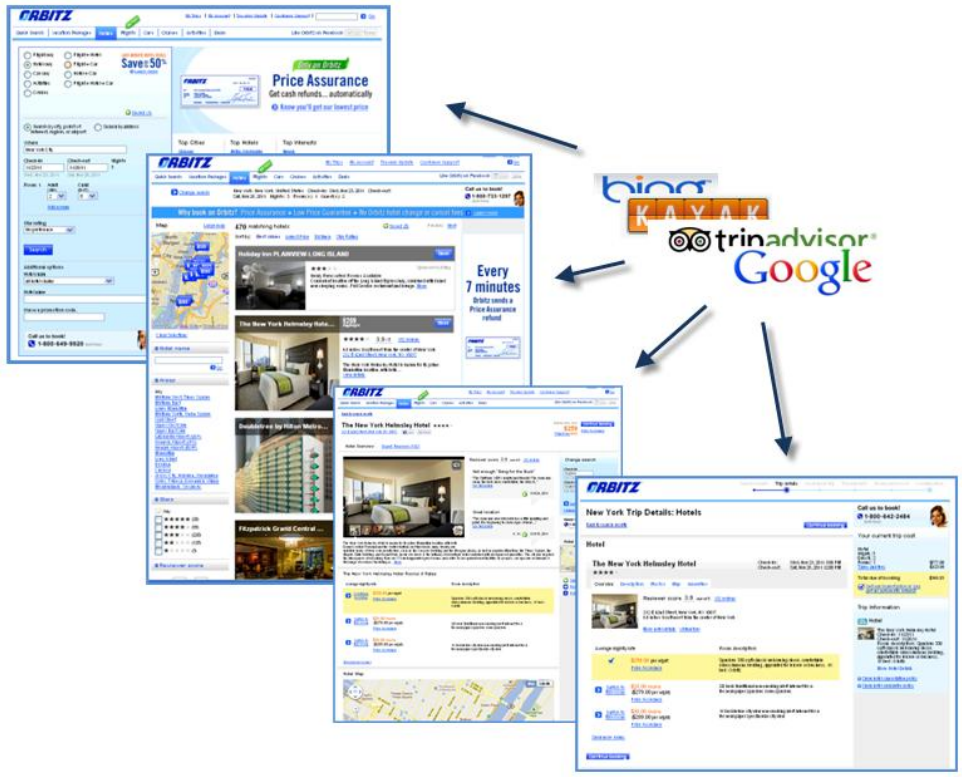

- Multiple points of entry.
- Multiple paths through site.
- Goal: tie events together to form picture of customer behavior.

Use Case – Selection Errors: Processing

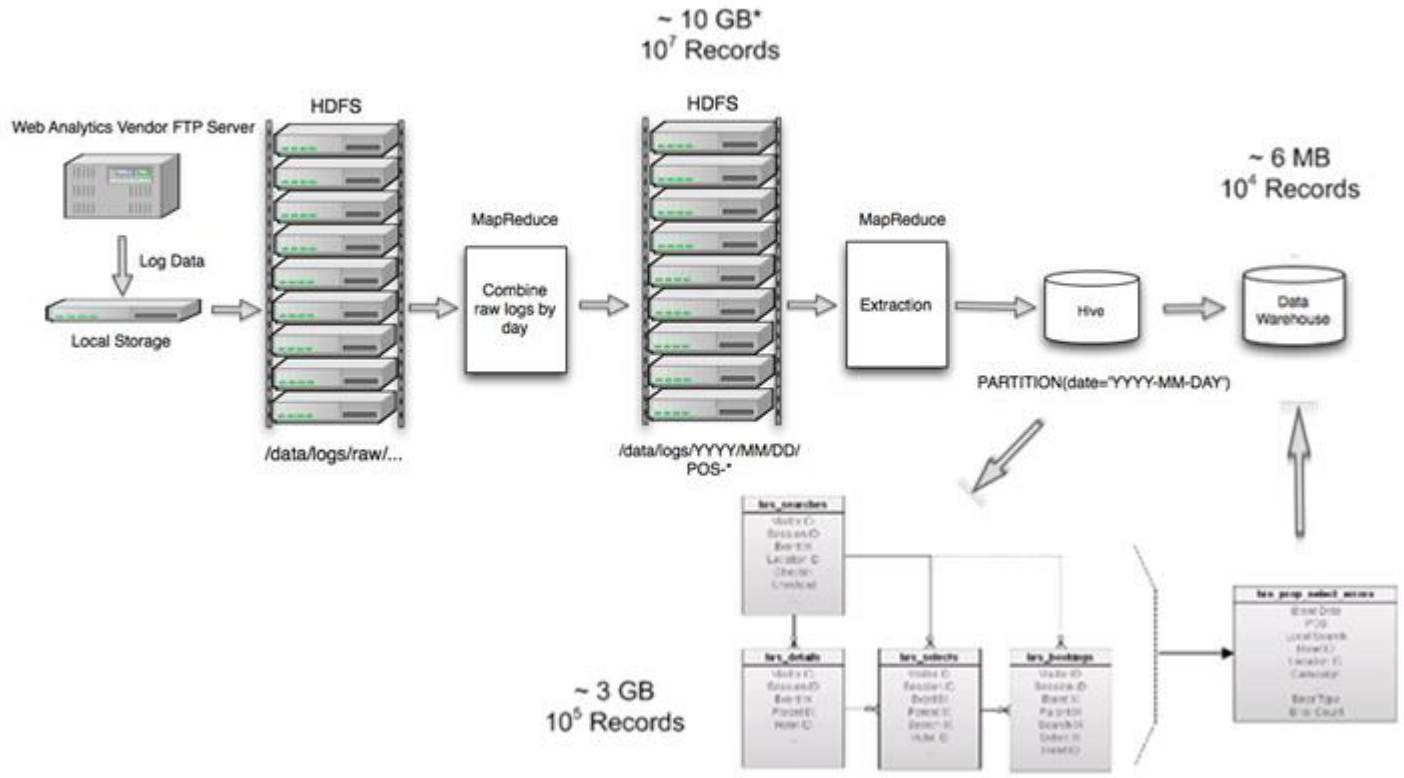

Use Case – Selection Errors: Visualization

Example Use Case: Beta Data

Use Case – Beta Data: Introduction

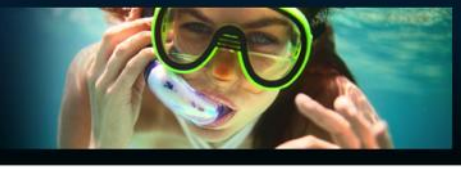

- Hotel Sort Optimization
- Compare A vs. B
- Web Analytics Data
  - What user saw.
  - How user behaved
- Server Log Data
  - Sorting behavior used.

Use Case – Beta Data Processing

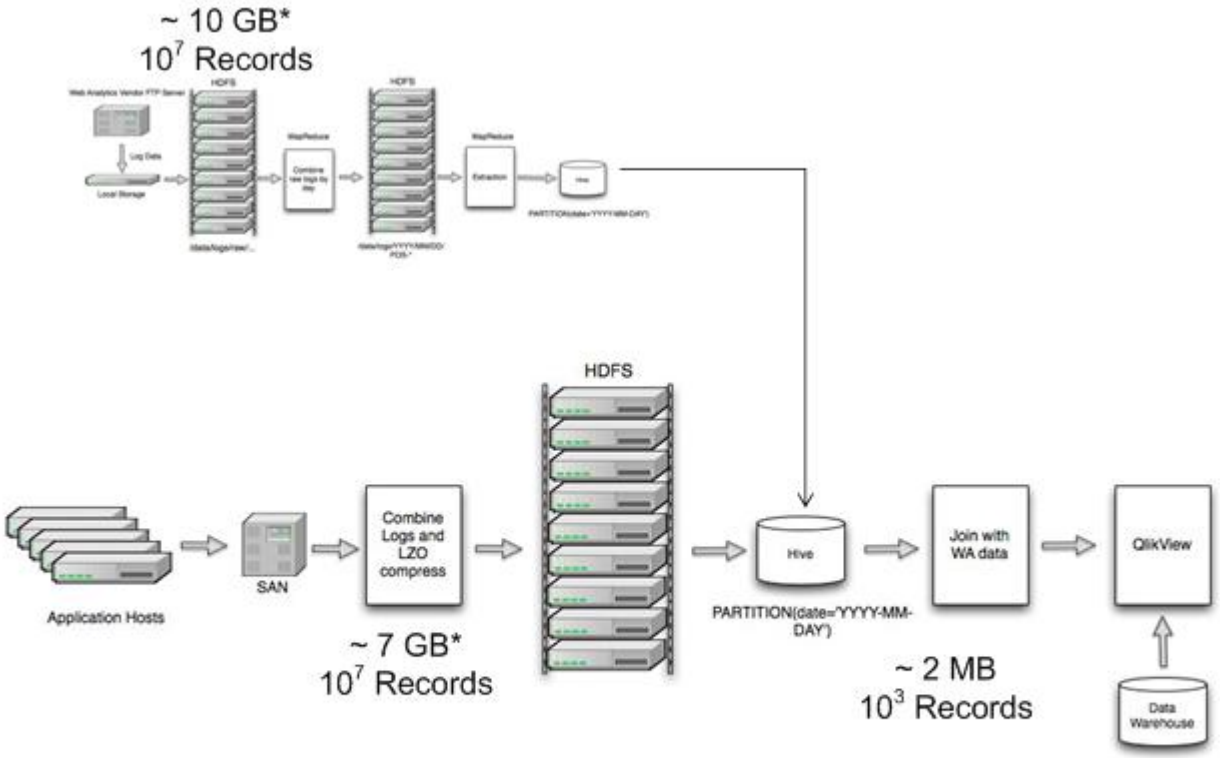

Use Case – Beta Data: Visualization

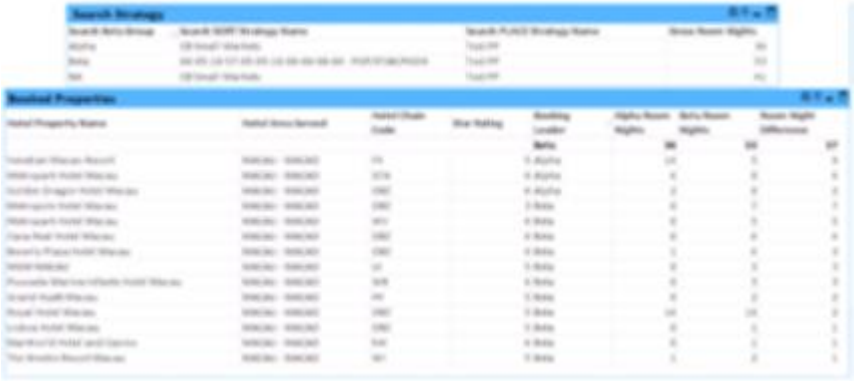

Example Use Case: RCDC

Use Case – RCDC: Introduction

- Understand and improve cache behavior.
- Improve “coverage”
  - Traditionally search 1 page of hotels at a time.
  - Get “just enough” information to present to consumers.
  - Increase amount of availability information we have when consumer performs a search.
- Data needed to support needs beyond reporting.

ORBITZ

Use Case – RCDC: Processing

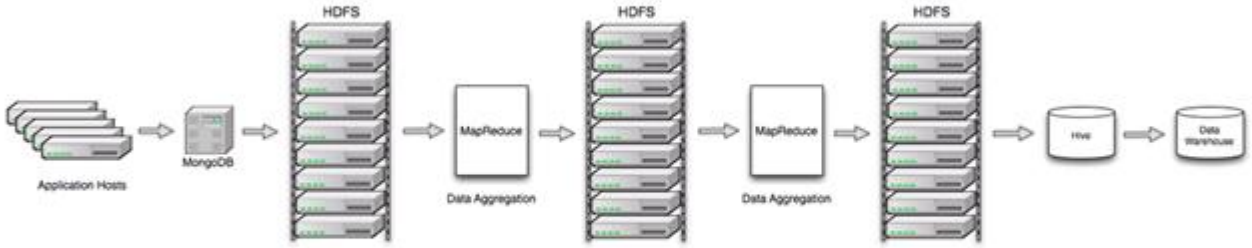

Use Case – RCDC: Visualization

Conclusions

- Hadoop market is still immature, but growing quickly. Better tools are on the way.
  - Look beyond the usual (enterprise) suspects. Many of the most interesting companies in the big data space are small startups.
- Hadoop won’t replace your EDW, but any organization with a large EDW should at least be exploring Hadoop as a complement to their BI infrastructure.

ORBITZ

Conclusions

- Work closely with your existing data management teams.
  - Your idea of what constitutes “big data” might quickly diverge from theirs.
- The flip-side to this is that Hadoop can be an excellent tool to off-load resource-consuming jobs from your data warehouse.

ORBITZ

ORBITZ

Thank you!

Questions?