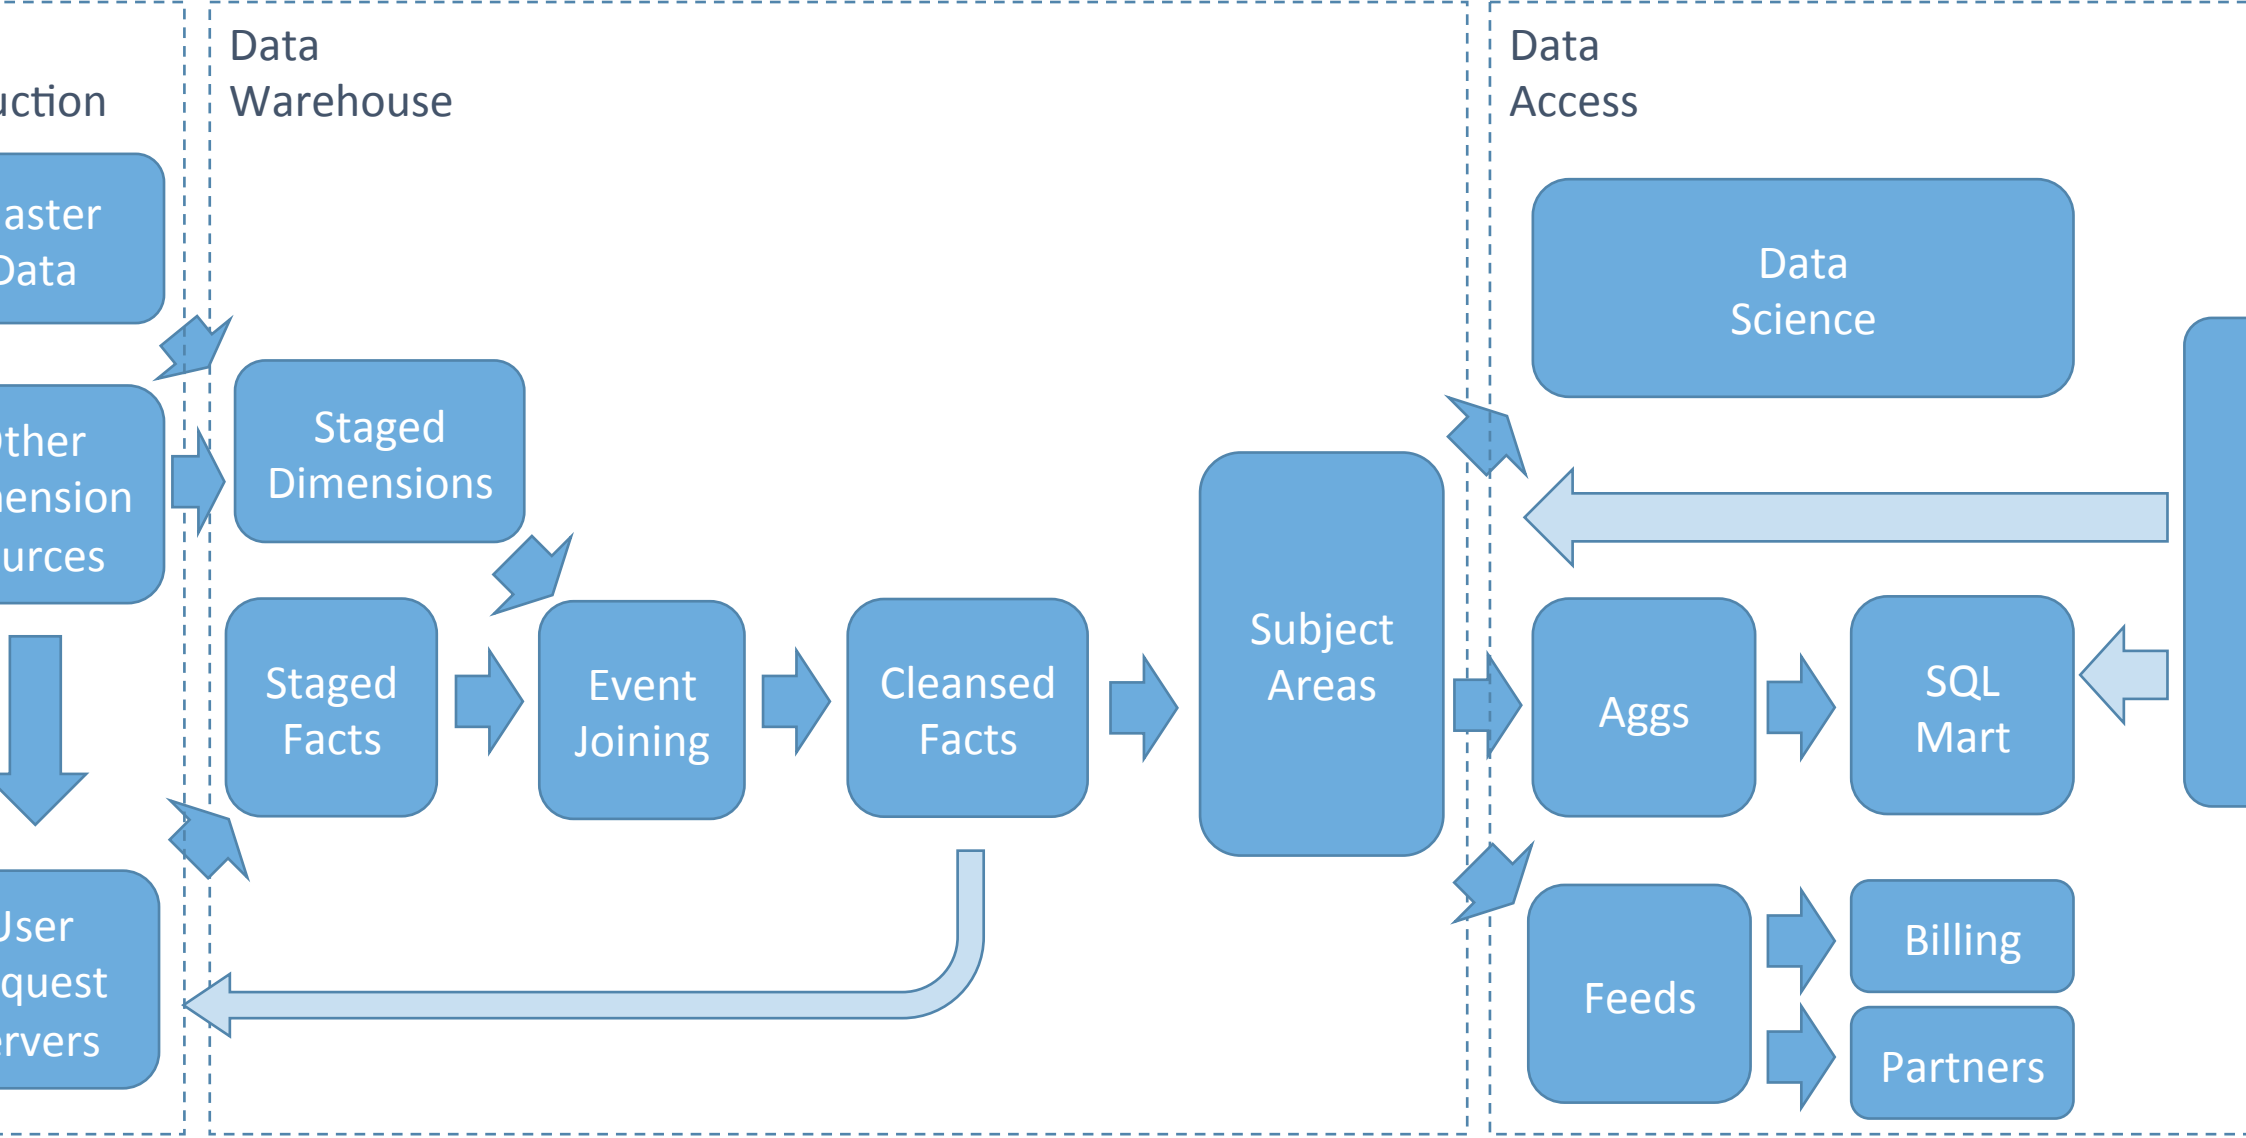


The Modern Giant-Scale Data Warehouse

Chad Walters
ASG Data Platform
Microsoft

HPTS
9/24/2013

Common Big Data Workflow



Data Production

Master Data:

- Interesting social challenges in large organization (100K+ employees)
- Bing Ads: Custom wrapper around SQL Server Master Data Services

Other Dimension Sources:

- Bing Ads: 10s of TBs of OLTP-based advertiser campaign data
 - Federated SQL Server farm
 - Business continuity requires at least 2 geo-replicated instances
- Google AdWords: 10s of TBs(?), F1/Spanner

User Request Servers:

- At giant-scale, these are globally distributed
 - Not necessarily close to the OLTP system
- At giant-scale, these are generally custom application services
- Bing Ads: Many 1000s of servers in several DCs around the globe, 100s of GBs of logs/hour
- Instrumentation done with Protocol Buffers (Google), Bond (Bing), Thrift (Facebook), etc.
 - Serialization/deserialization frameworks with forward/backward compat capabilities

Production

Master Data

Other Dimension Sources

User Request Servers

Data Warehouse

Dimension Staging:

Deltas moved to event joining fabric where snapshots are hosted, maintained

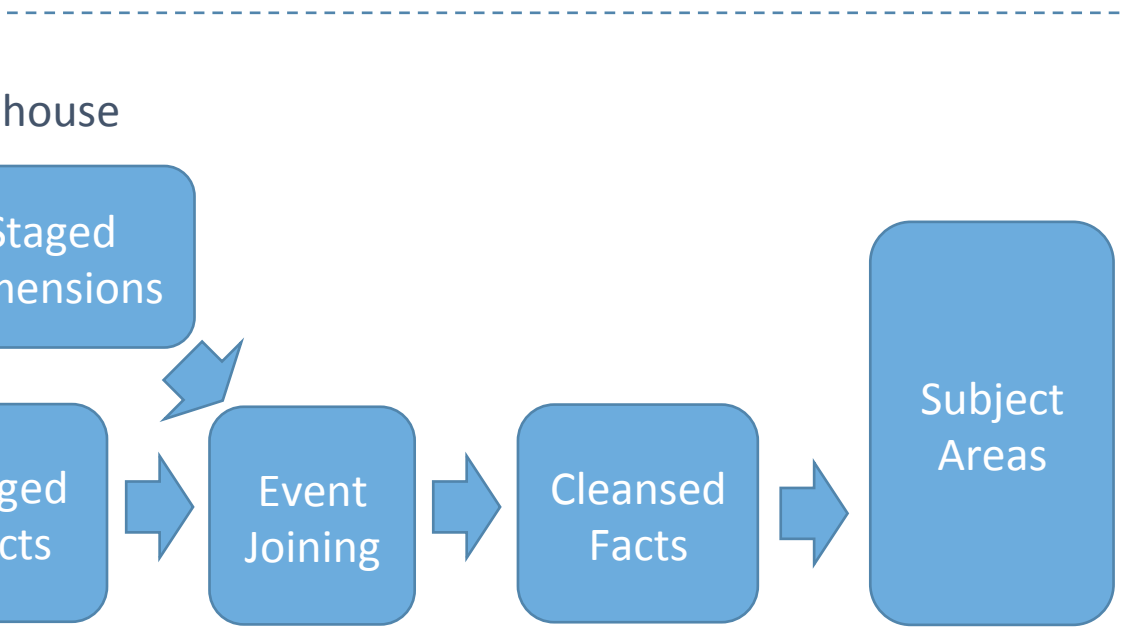
Bing Ads: SQL Server farm used to distribute deltas globally; deltas uploaded to Cosmos [production], deltas sent to Storm + Tempest [alpha]

Fact Staging:

- Instrumentation from user request servers move to event joining fabric
- Bing Ads: Log uploads to Cosmos [production], Kafka + Storm [alpha]

Event Joining:

- Bring together disparate pieces of same request as well as connecting downstream user actions to original request
- Bing Ads: Cosmos [production], Storm + Tempest (custom store) [alpha]
- Google AdWords: Photon
- Note that both Bing Ads and AdWords process data redundantly...



Data Warehouse (cont.)

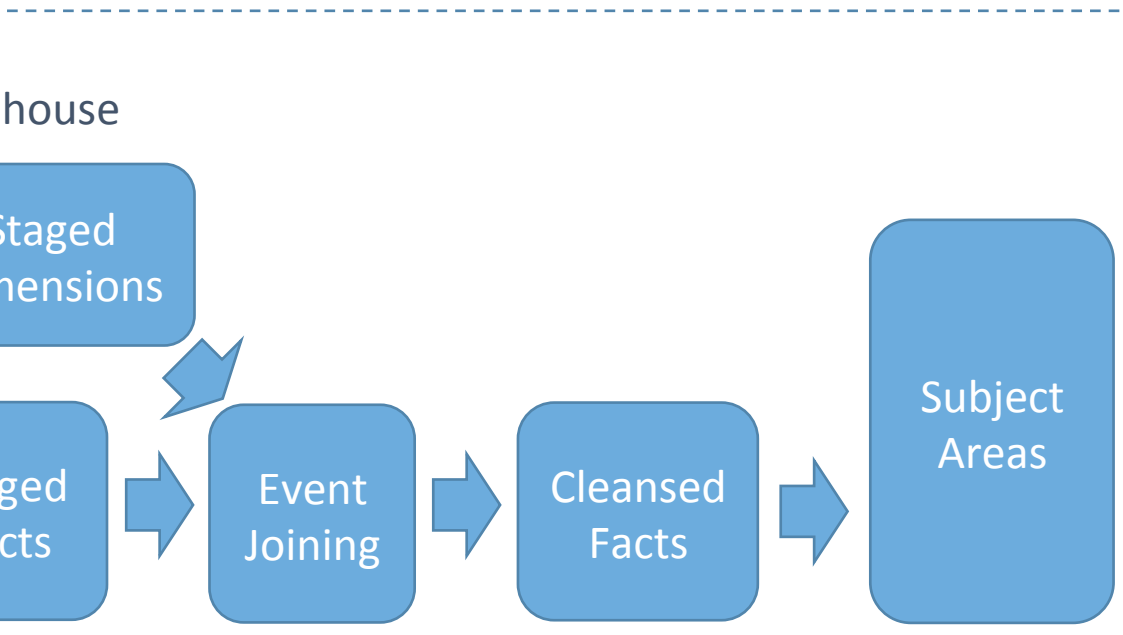
Used Facts:

Data quality at syntactic and semantic levels

Bing Ads: Fraud processing; Cosmos [production],
Storm + Tempest (recent) + Cosmos (archival)
[alpha]

Subject Areas:

- Curated views grouped around particular business needs; designed to be joined together
- At giant scale, this is too big to keep it all in memory cost effectively (maybe flash too)
- Bing Ads: 10s of PBs; Cosmos [production], Storm + Tempest (recent) + Cosmos (archival) [alpha]



N.B.:

- Use of flexible schema systems based around Parquet/Thrift/Avro/Bond is critical for these systems at scale to avoid syntactic breakage
- Centrally managed data model for “public” data with “private” pass-through blocks for production consumer velocity

Data Access

Most of these are various consumption patterns

Bing Ads uses federated SQL marts for advertiser reporting

- 100s of TBs, 20GB loaded per hour

- Dimensions provide via same system that feeds request servers

- Soon: ClusteredColumnStore in SQL Server 2014

- 3x storage decrease, 35+% better query performance

Google AdWords uses F1's OLAP capabilities

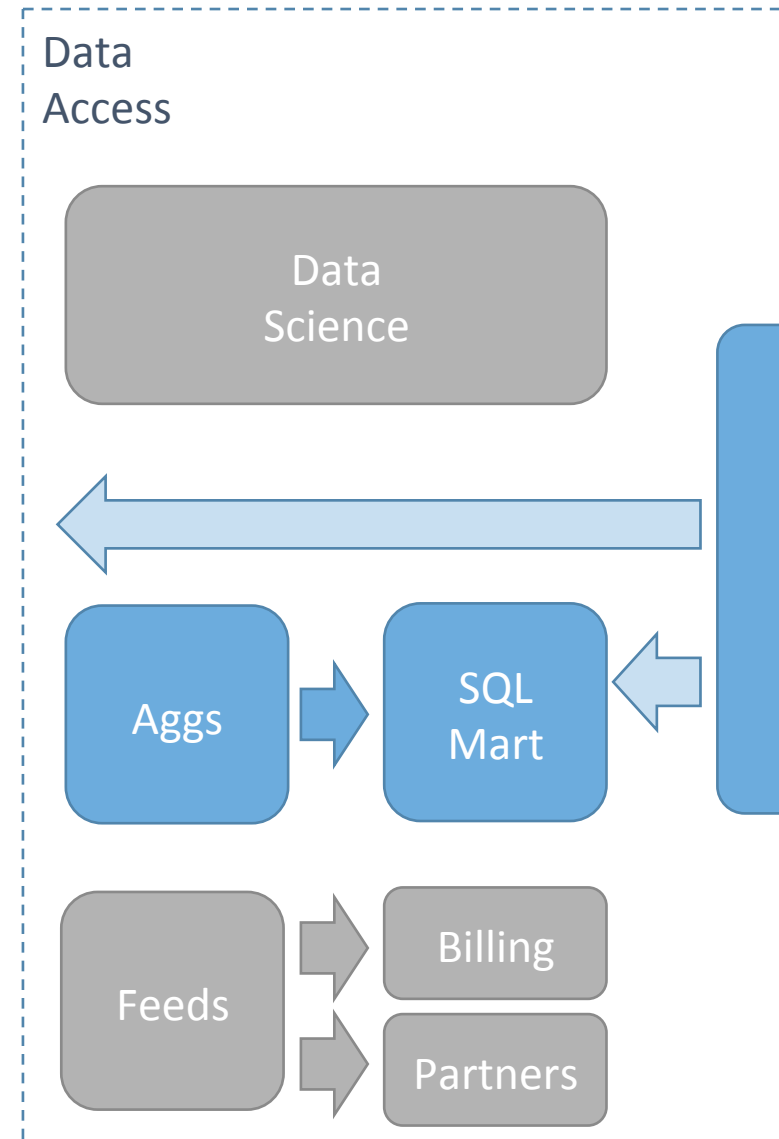
- Joins against OLTP tables for dimensions

Data Virtualization (Bing Platform):

- Holds sophisticated semantic model of data publications for purposes of reporting

- Can target multiple fabrics: SCOPE, SQL marts, etc.

- Can do some joins and aggregations across fabrics



OLTP vs. OLAP vs. Streaming

Can I do all my work in one system?

“I like columns so I put a column store in my OLTP system”

No need for streaming/CEP/ETL...

AWESOME!!

Breaks down at giant scale, though

Even F1/Spanner doesn't do it all