



# Big Data Platform at *interest*

*Mao Ye*



**Data Architecture**

**Design Choices for Hadoop Platform**

**Pinball for Workflow Management**



# **Data Architecture**

# Data at Pinterest


- 60 Billion Pins
- 1 Billion boards
- 100M MAU
- 60 PB of data on S3
- 3 PB processed every day
- 2000 node Hadoop cluster
- 250 engineers




# Pinterest Data Architecture



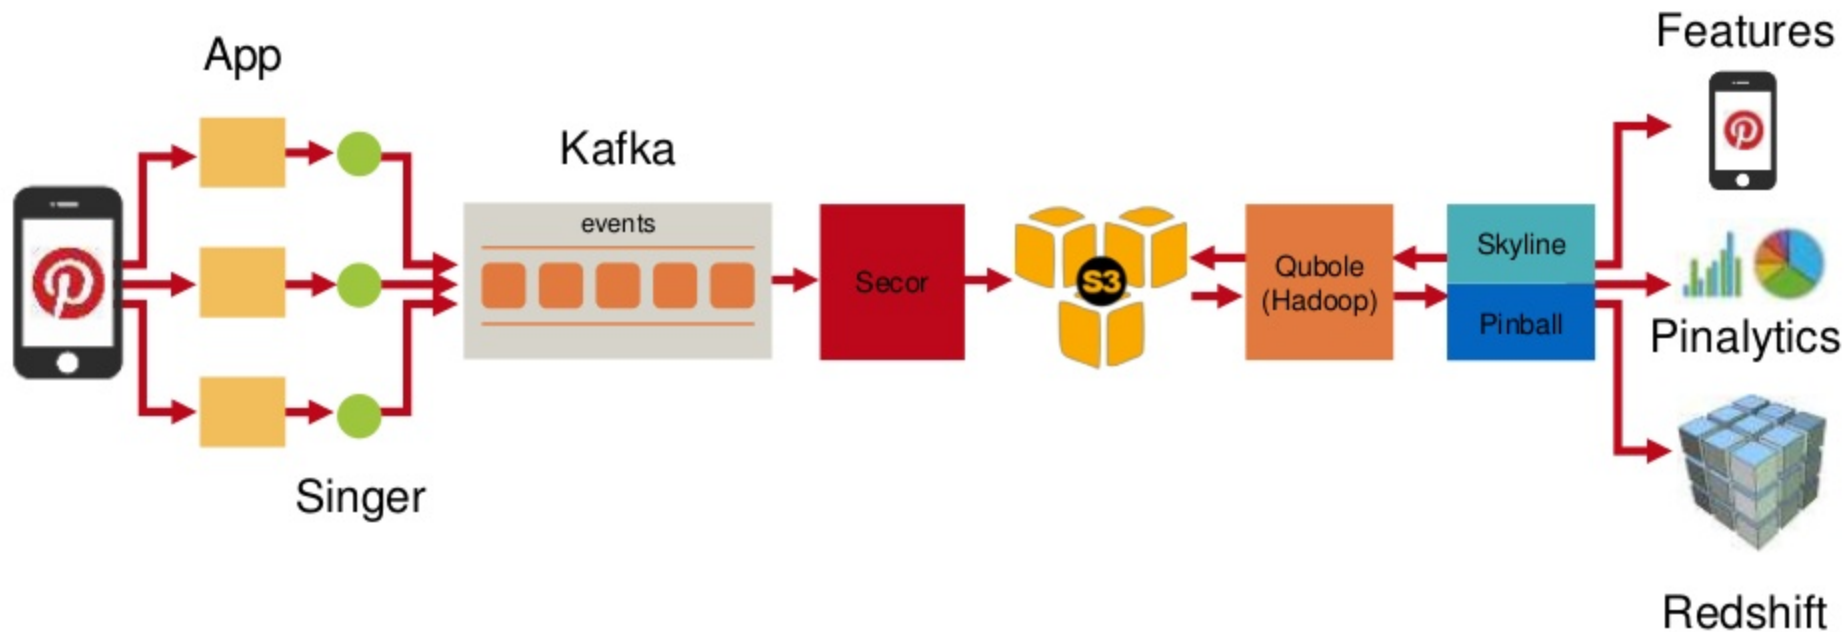
# Pinterest Data Architecture



# Pinterest Data Architecture



# Pinterest Data Architecture








# Design Choices for Hadoop Platform

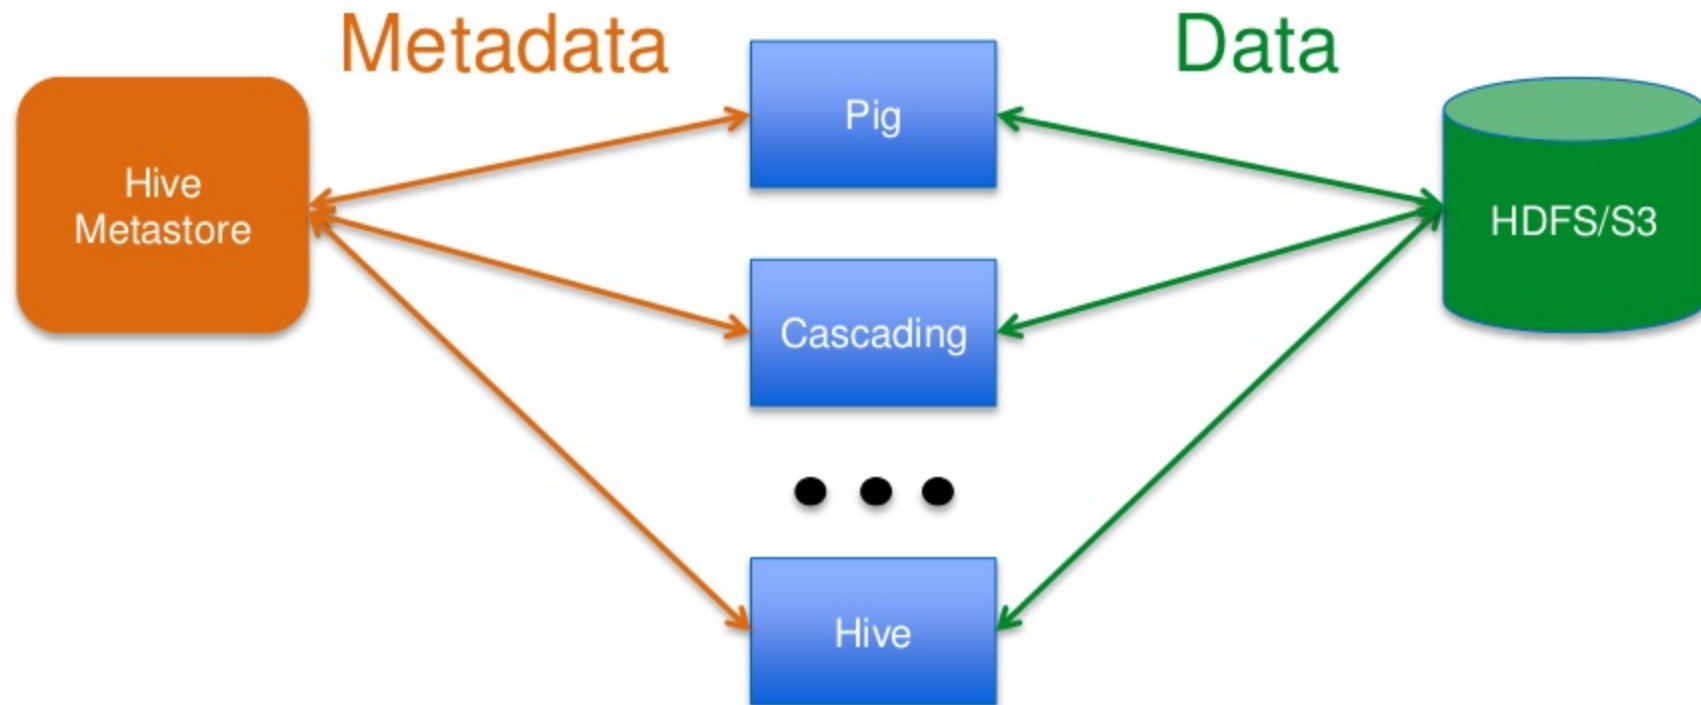
# Hadoop Platform Requirements

- Isolated multi-tenancy
- Elasticity
- Support multiple clusters
- Ephemeral clusters
- Access control layer
- Shared data store
- Easy deployment

# Decoupling compute & storage



# Centralized Hive Metastore



# Multi-layered Packaging

Runtime Staging  
(on S3)

Automated  
Configuration  
(Masterless Puppet)


Baked AMI

Mapreduce Jobs  
Hadoop Jars/Libs  
Job/User level Configs

Software Packages/Libs  
Configs (OS/Hadoop)  
Misc Sys Admin

OS  
Bootstrap Script  
Core SW

# Executor Abstraction Layer



# Why Qubole?

- Hadoop & Spark as managed services
- Tight integration with Hive
- Graceful cluster scaling
- API for simplified executor abstraction
- Advanced support for spot instances
- Baked AMI customization



# **Pinball for Workflow Management**




# Scale of Processing

- **Scale:**

- 60 Billion Pins
- Hundreds of workflows
- Thousands of jobs
- 500+ jobs in a workflow
- 3 petabytes processed daily

- **Support:**

- Hadoop, Cascading, Hive, Spark ...



# Why Pinball?


- **Requirements**

- Simple abstractions
- Extensible in future
- Reliable stateless computing
- Easy to debug
- Scales horizontally
- Can be upgraded w/o aborting workflows
- Rich features like auto-retries, per-job emails, overrun policies...

- **Options**

- Apache Oozie, Azkaban, Luigi

# Pinball Design



# Workflow Model

- **Workflow**
  - A directed graph of nodes called jobs
- **Edge**
  - Run after dependence
- **Node**
  - Job is a node

