



by Facebook

Rebuilding Atlas – Advertising at Scale at Facebook



by Facebook

Jason McHugh, Software Engineer | March, 2015

Outline

- Ad serving technology
- History of Atlas
- The acquisition
- Challenges
- Architecture
- Lessons learned

Advertising



#1 Reason for a Big Belly

 **\$10 OFF \$20 Coupon Code: SPRSALE2G**

MiniInTheBox [Shop now!](#)

 €10.34

 €12.87

[Learn More](#)



Orders Ship Same Day from Digi-Key



[Shop Now](#)

Advertising

- Advertising a huge industry with massive budgets
- Digital advertising is the fastest growing advertising medium
- More time spent with digital media than TV
- Time spent on mobile exceeded TV

Third Party Ad Serving




How Third Party Ad Serving Works

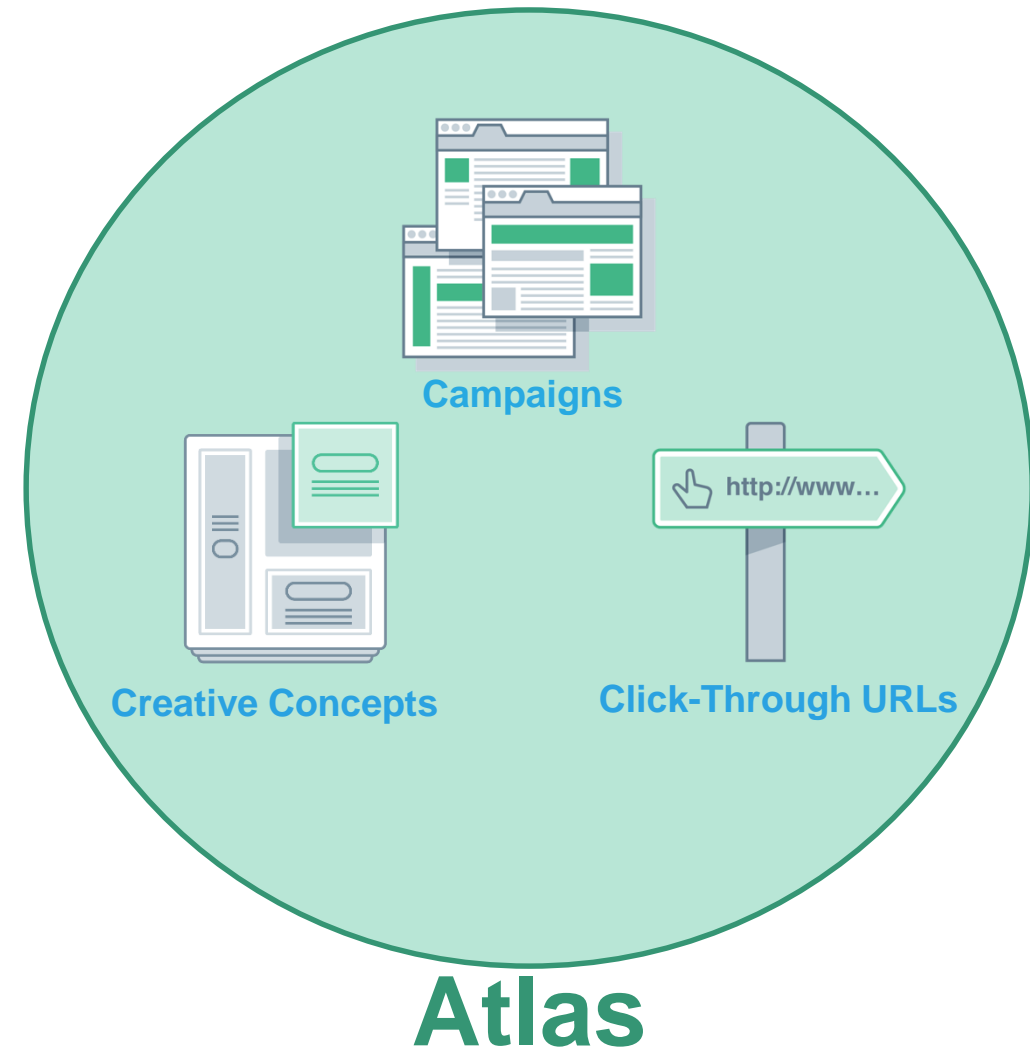


Advertiser

I want
to go to
there.



How Third Party Ad Serving Works



 DASHBOARD

 CAMPAIGNS

 ADVERTISERS

Roni Air

Manage Creatives

Set Up Targeting

Measure Actions

Manage Tags

 REPORTS

Advertisers > Roni Air >



Create a New Audience

Cancel


Save

Improve ad performance by reaching the people who matter to you most.

Name*

Description

Locations
United States, Nevada
Las Vegas
United States, Washington
Seattle

Home location or Device IP 

-

Age
The size of this audience will reflect the number of people on Facebook who match your age settings.

Gender All Men Women


Device Type All Mobile Desk/Laptop

Remarketing

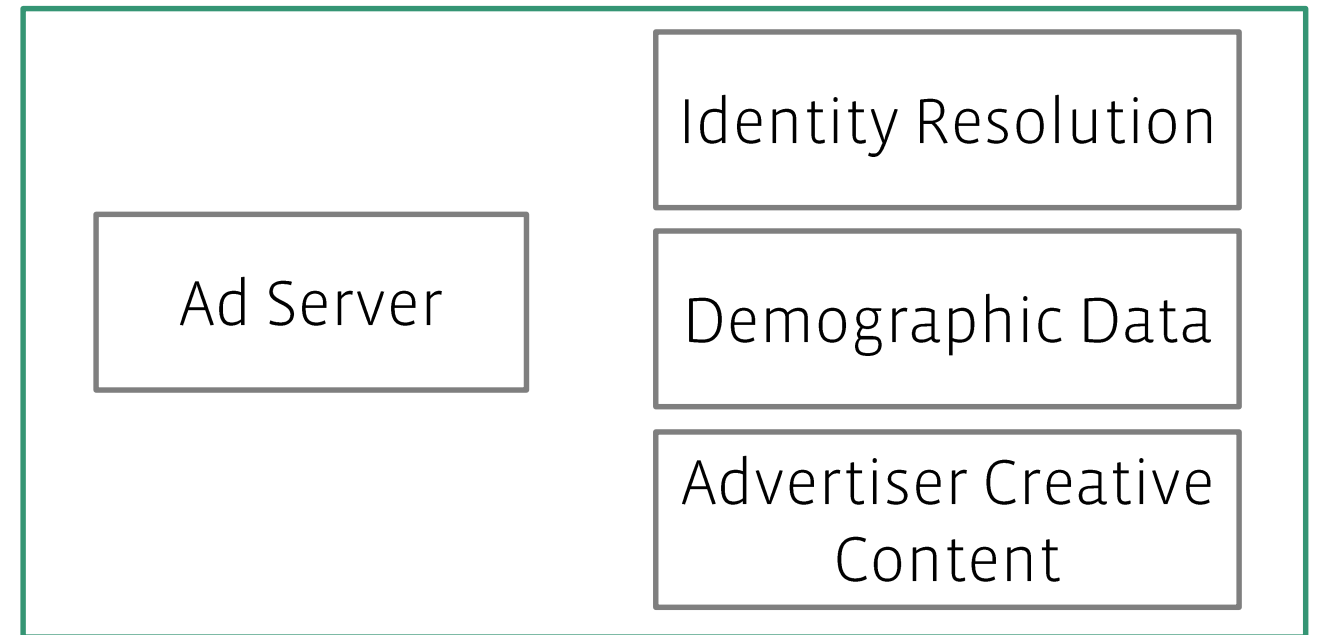
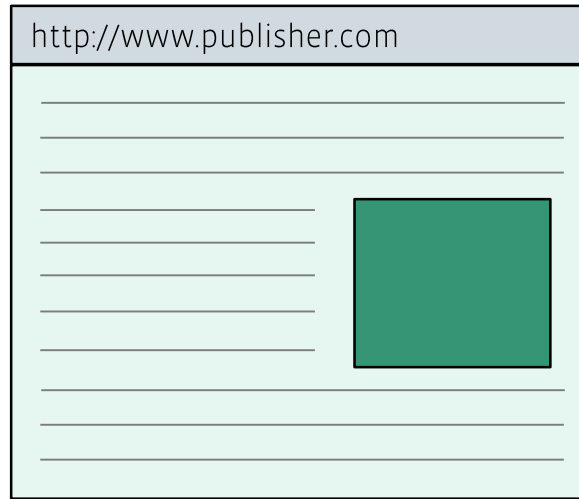
Remarket to your audience

Choose a segment of people who've already visited an advertiser site. You can target ads to them based on their behaviors.

Serving An Ad



You, me, or
other random
human



Atlas

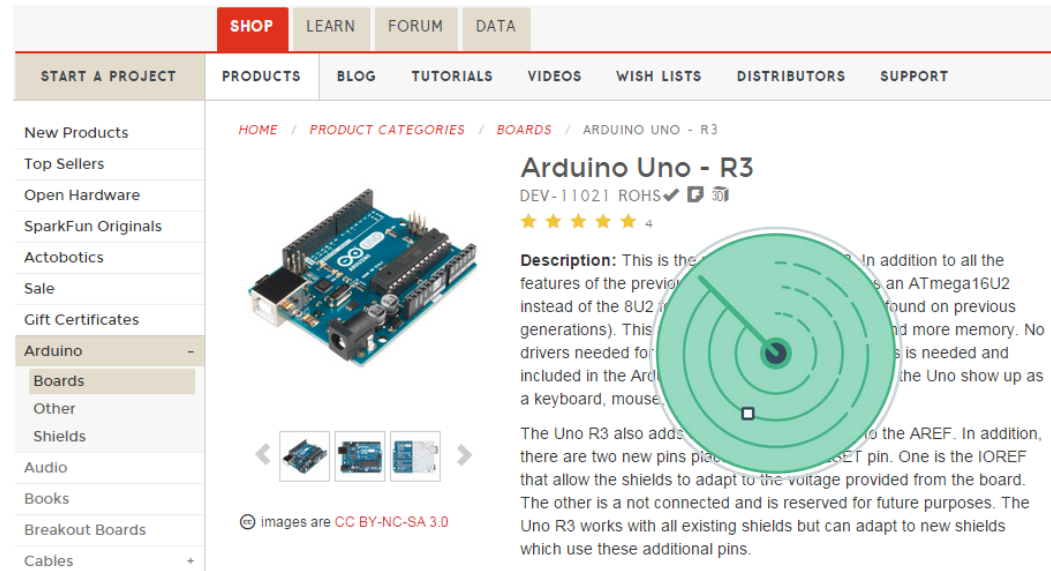
Ad Server

Identity Resolution

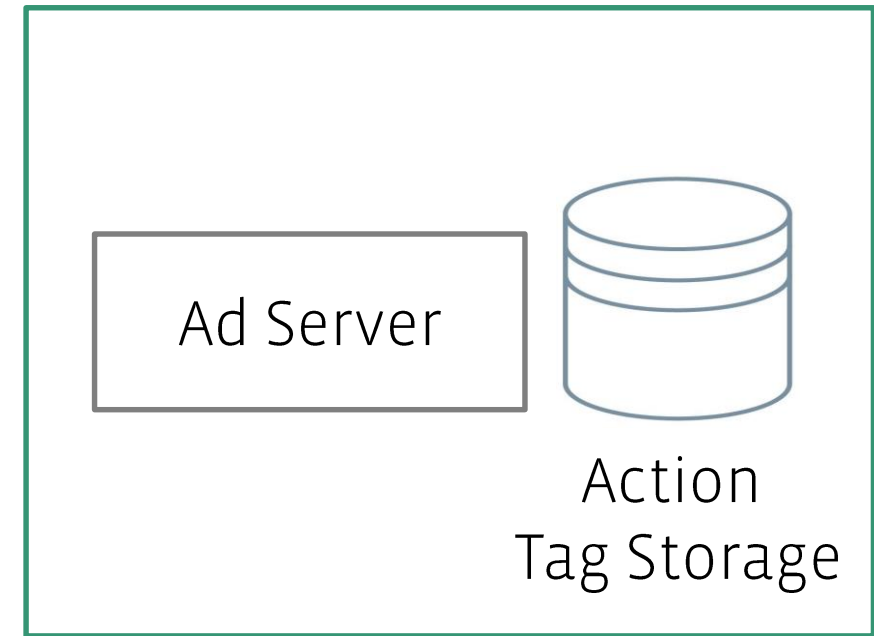
Demographic Data

Advertiser Creative
Content

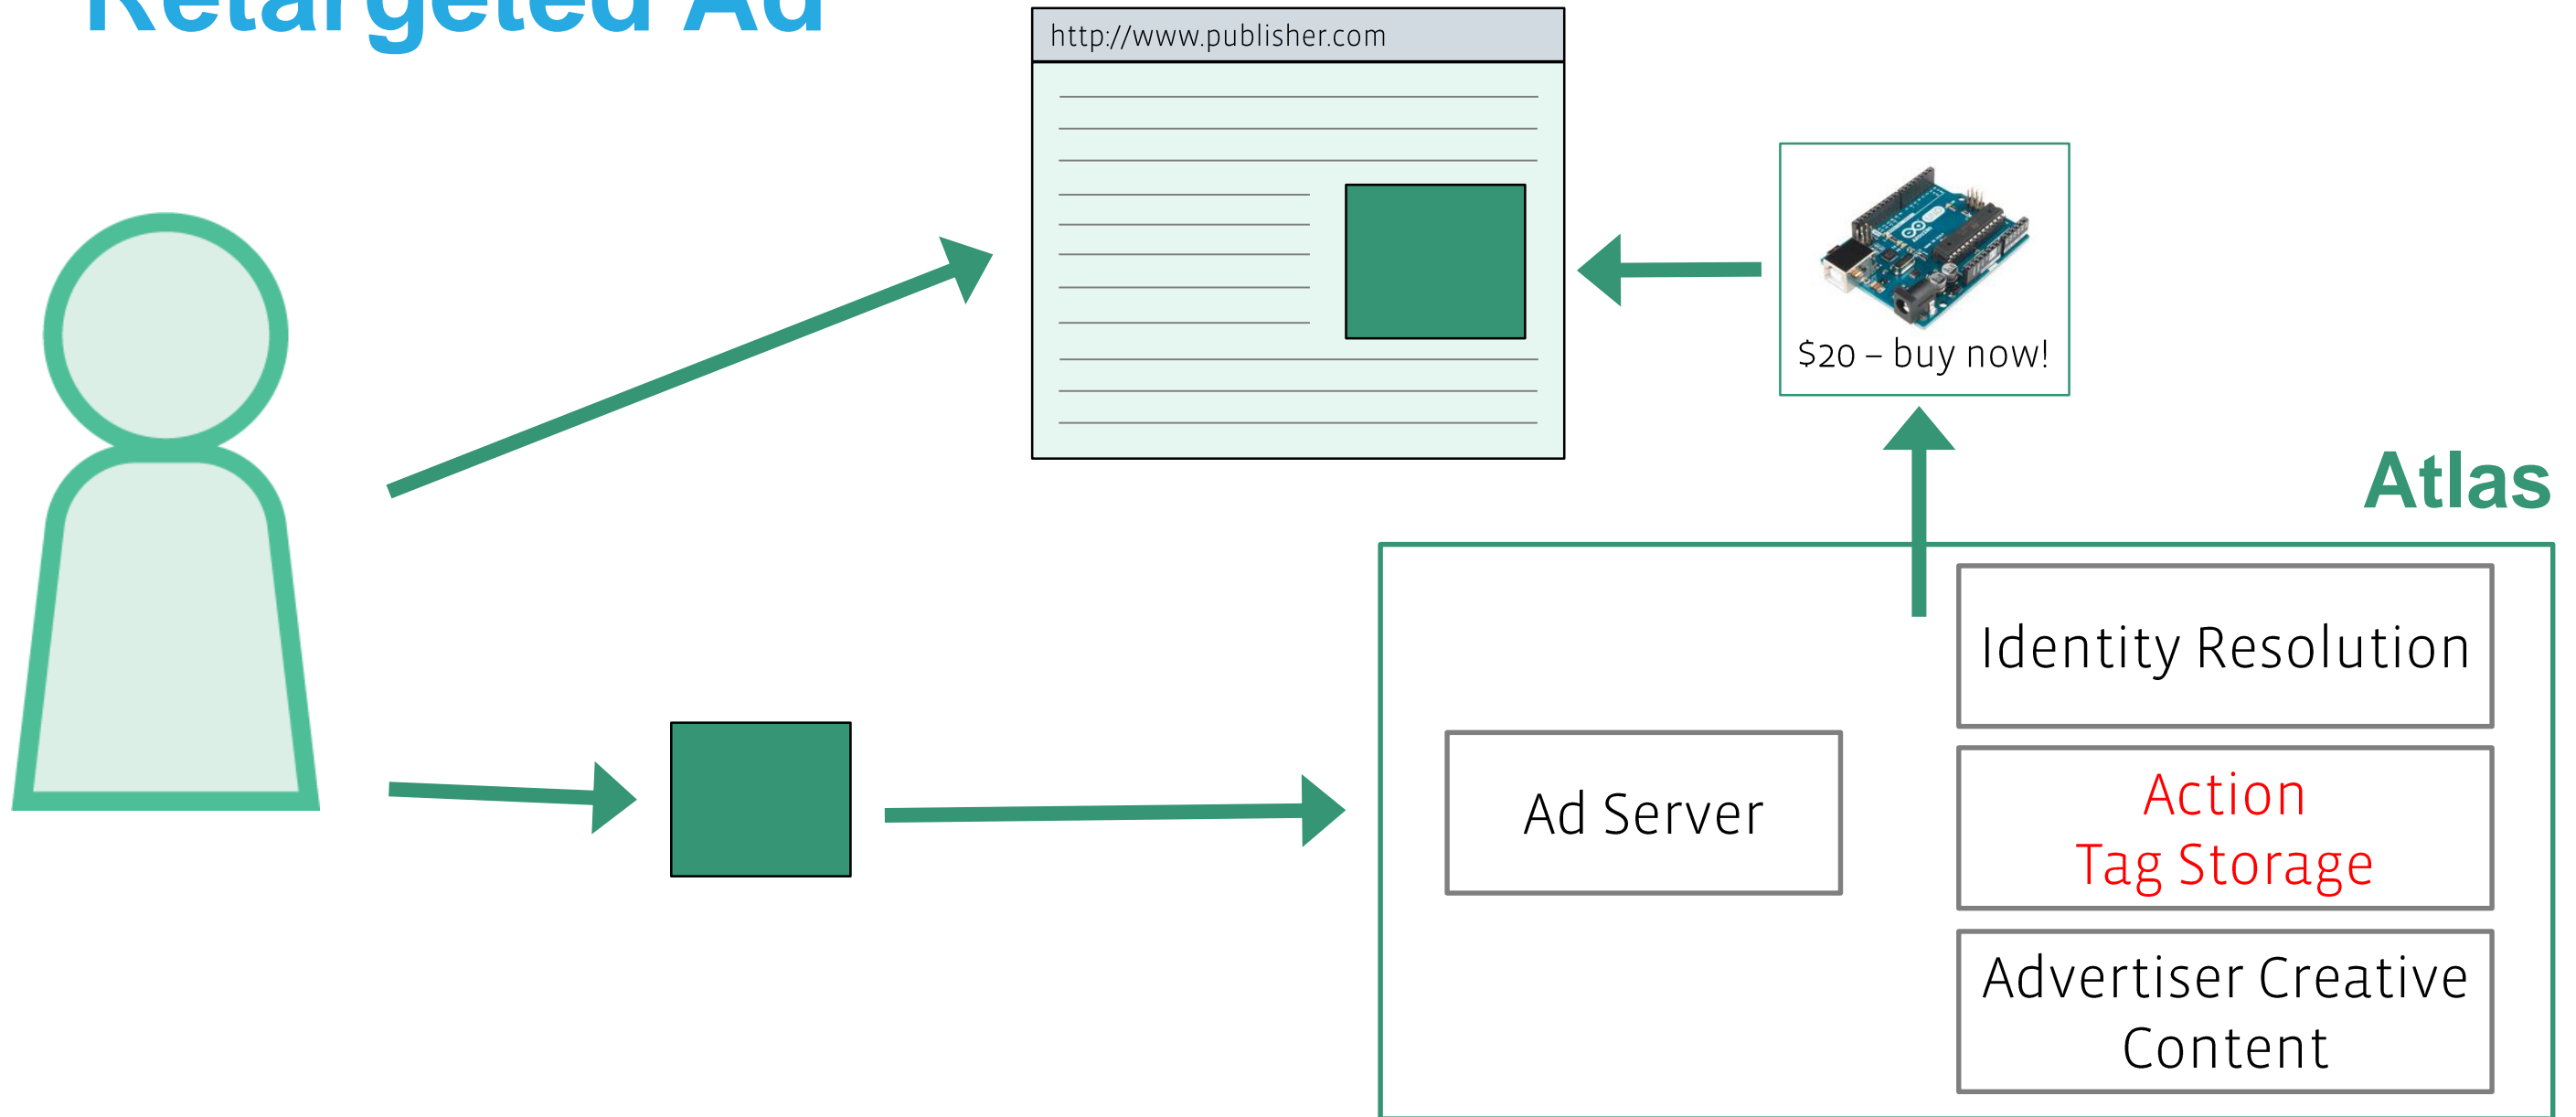
Retargeting



Atlas



Retargeted Ad



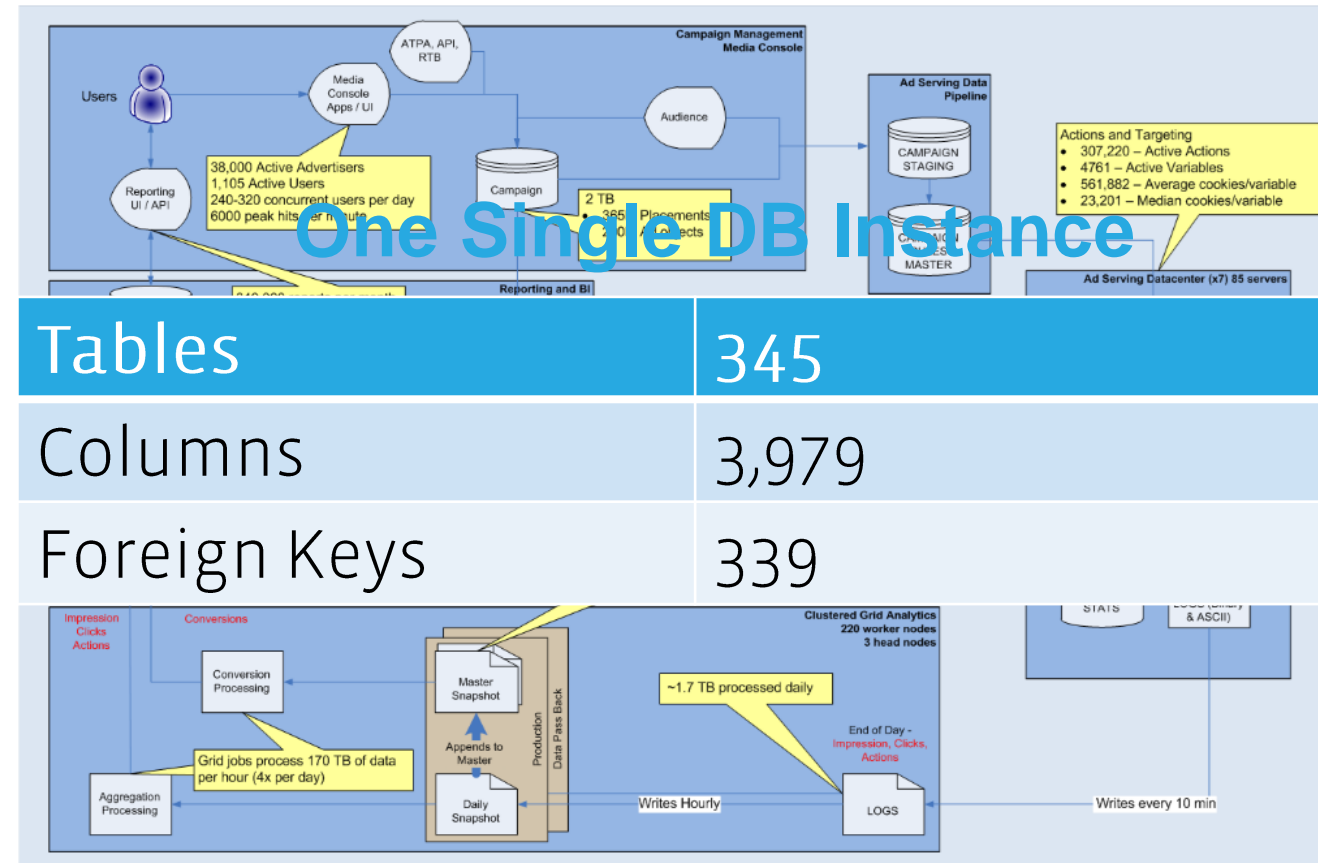
History of Atlas

- 1997 – Atlas started life as Avenue A
- 2003 – Avenue A becomes aQuantive
- 2007 – aQuantive was purchased by Microsoft for \$6 billion dollars
- 2012 – Microsoft took a \$6.2 billion dollar writedown
- 2013 – Facebook acquired Atlas in April
- 2014 – Atlas New publically launched at Ad Week in September

Challenges

Challenge: Understand the System

- Ad Tech Stack
- Architecture
- Data Model and databases
- Data flows
- Deployment



Challenge: Huge Product

- Third party ad server
- Advertiser and publisher negotiations – RFP
- Search management
- Search optimization
- Email tracking
- Custom analysis and reporting
- Rich media including video
- Franchise Management

Challenge: No Lift and Shift

- Lift and Shift
 - Common approach after acquisition
 - Take as much as exists at the time of the purchase
 - Move to your data centers and then evolve
- Not possible here
 - Non-open compute hardware
 - Usage of Microsoft close-source technologies

Architecture

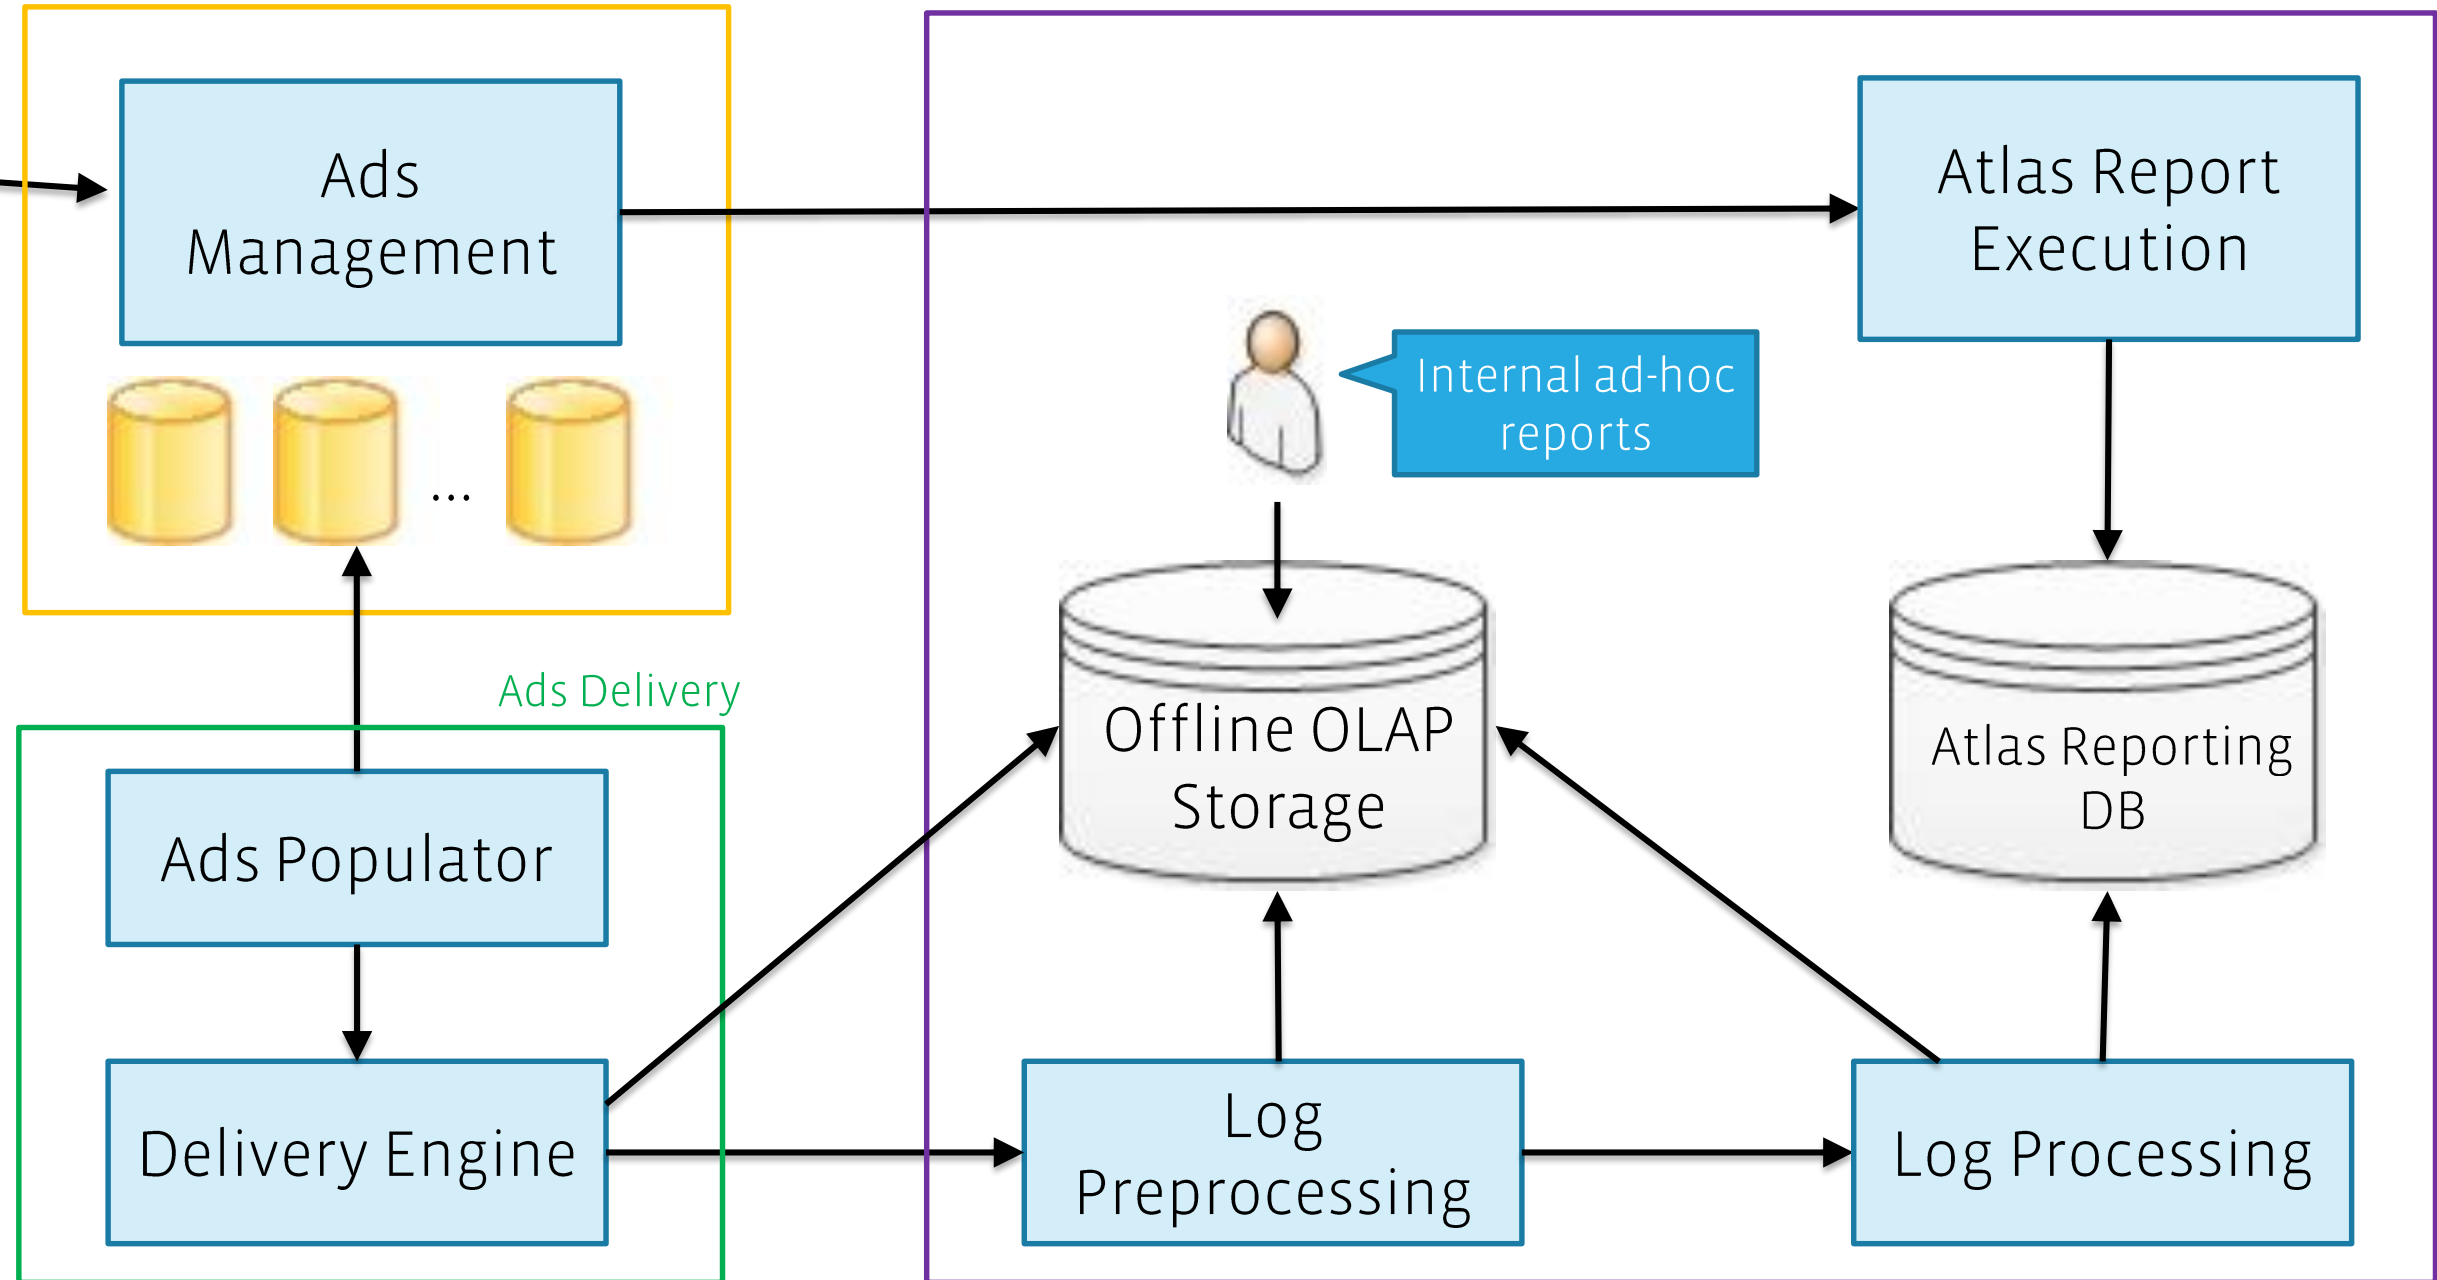
- Holistic view of the logical architecture
- Detail a piece of the physical architecture

Logical Architecture

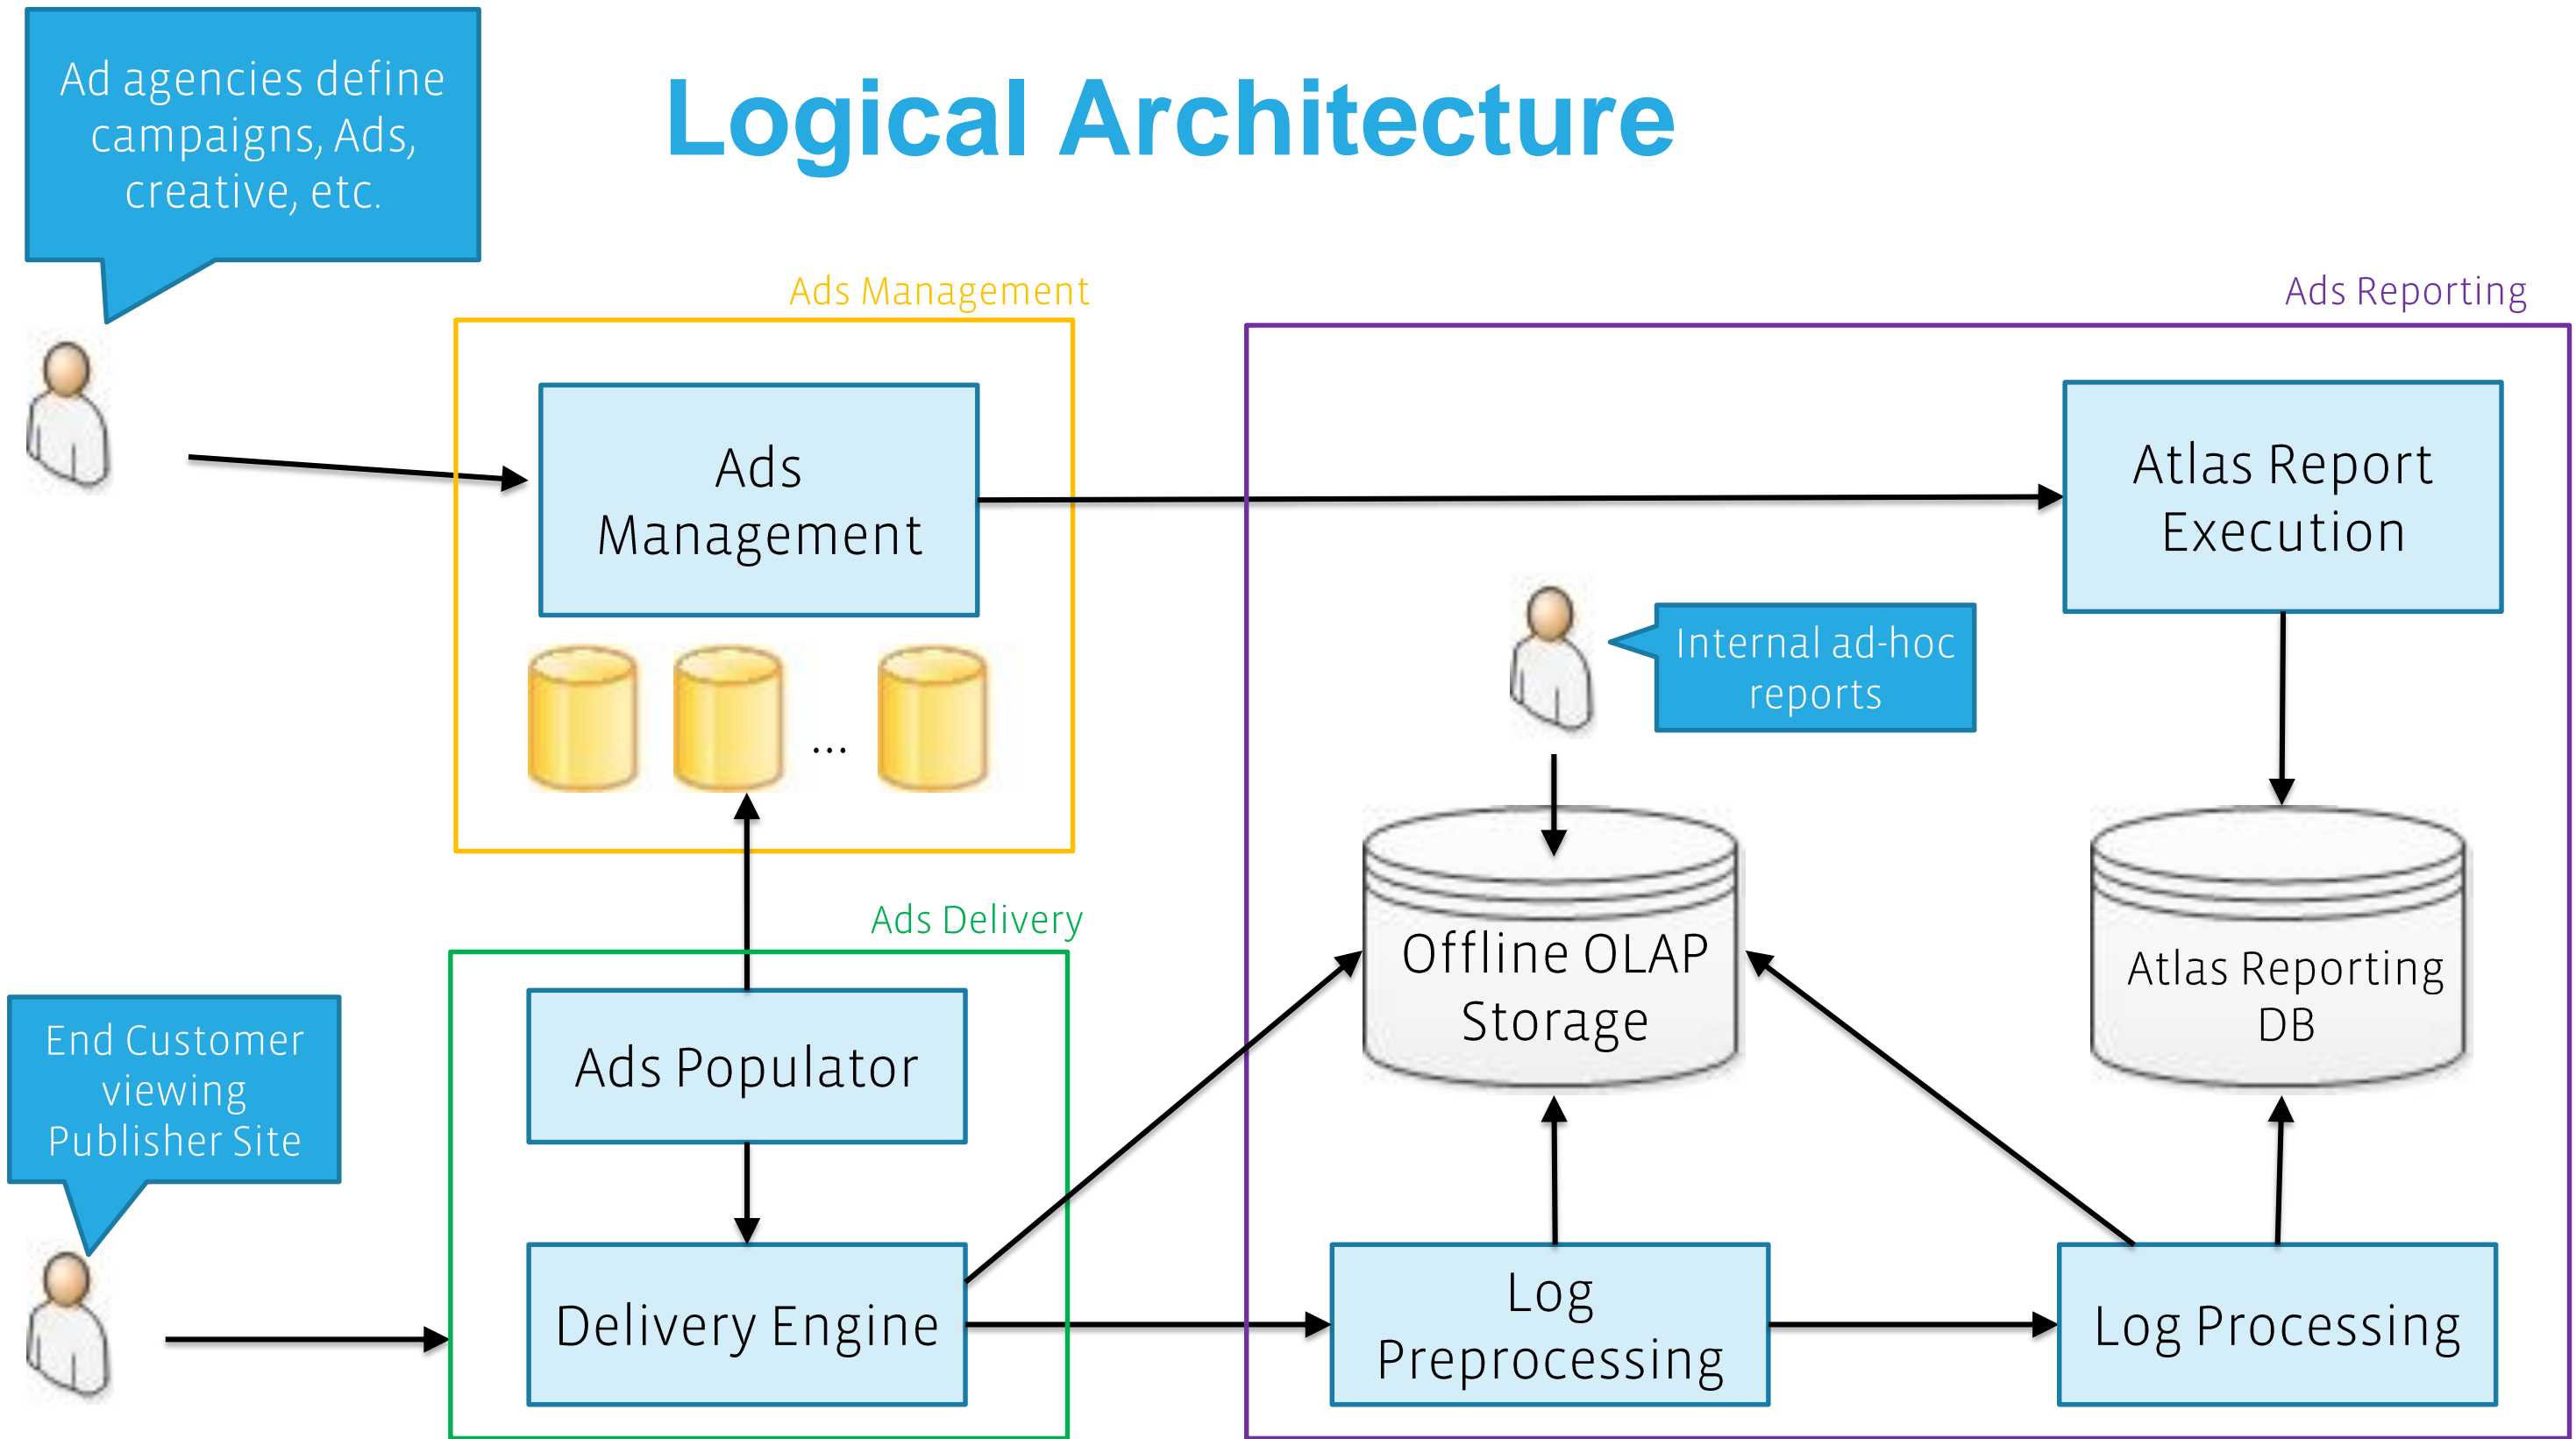
Ad agencies define campaigns, Ads, creative, etc.

Ads Management

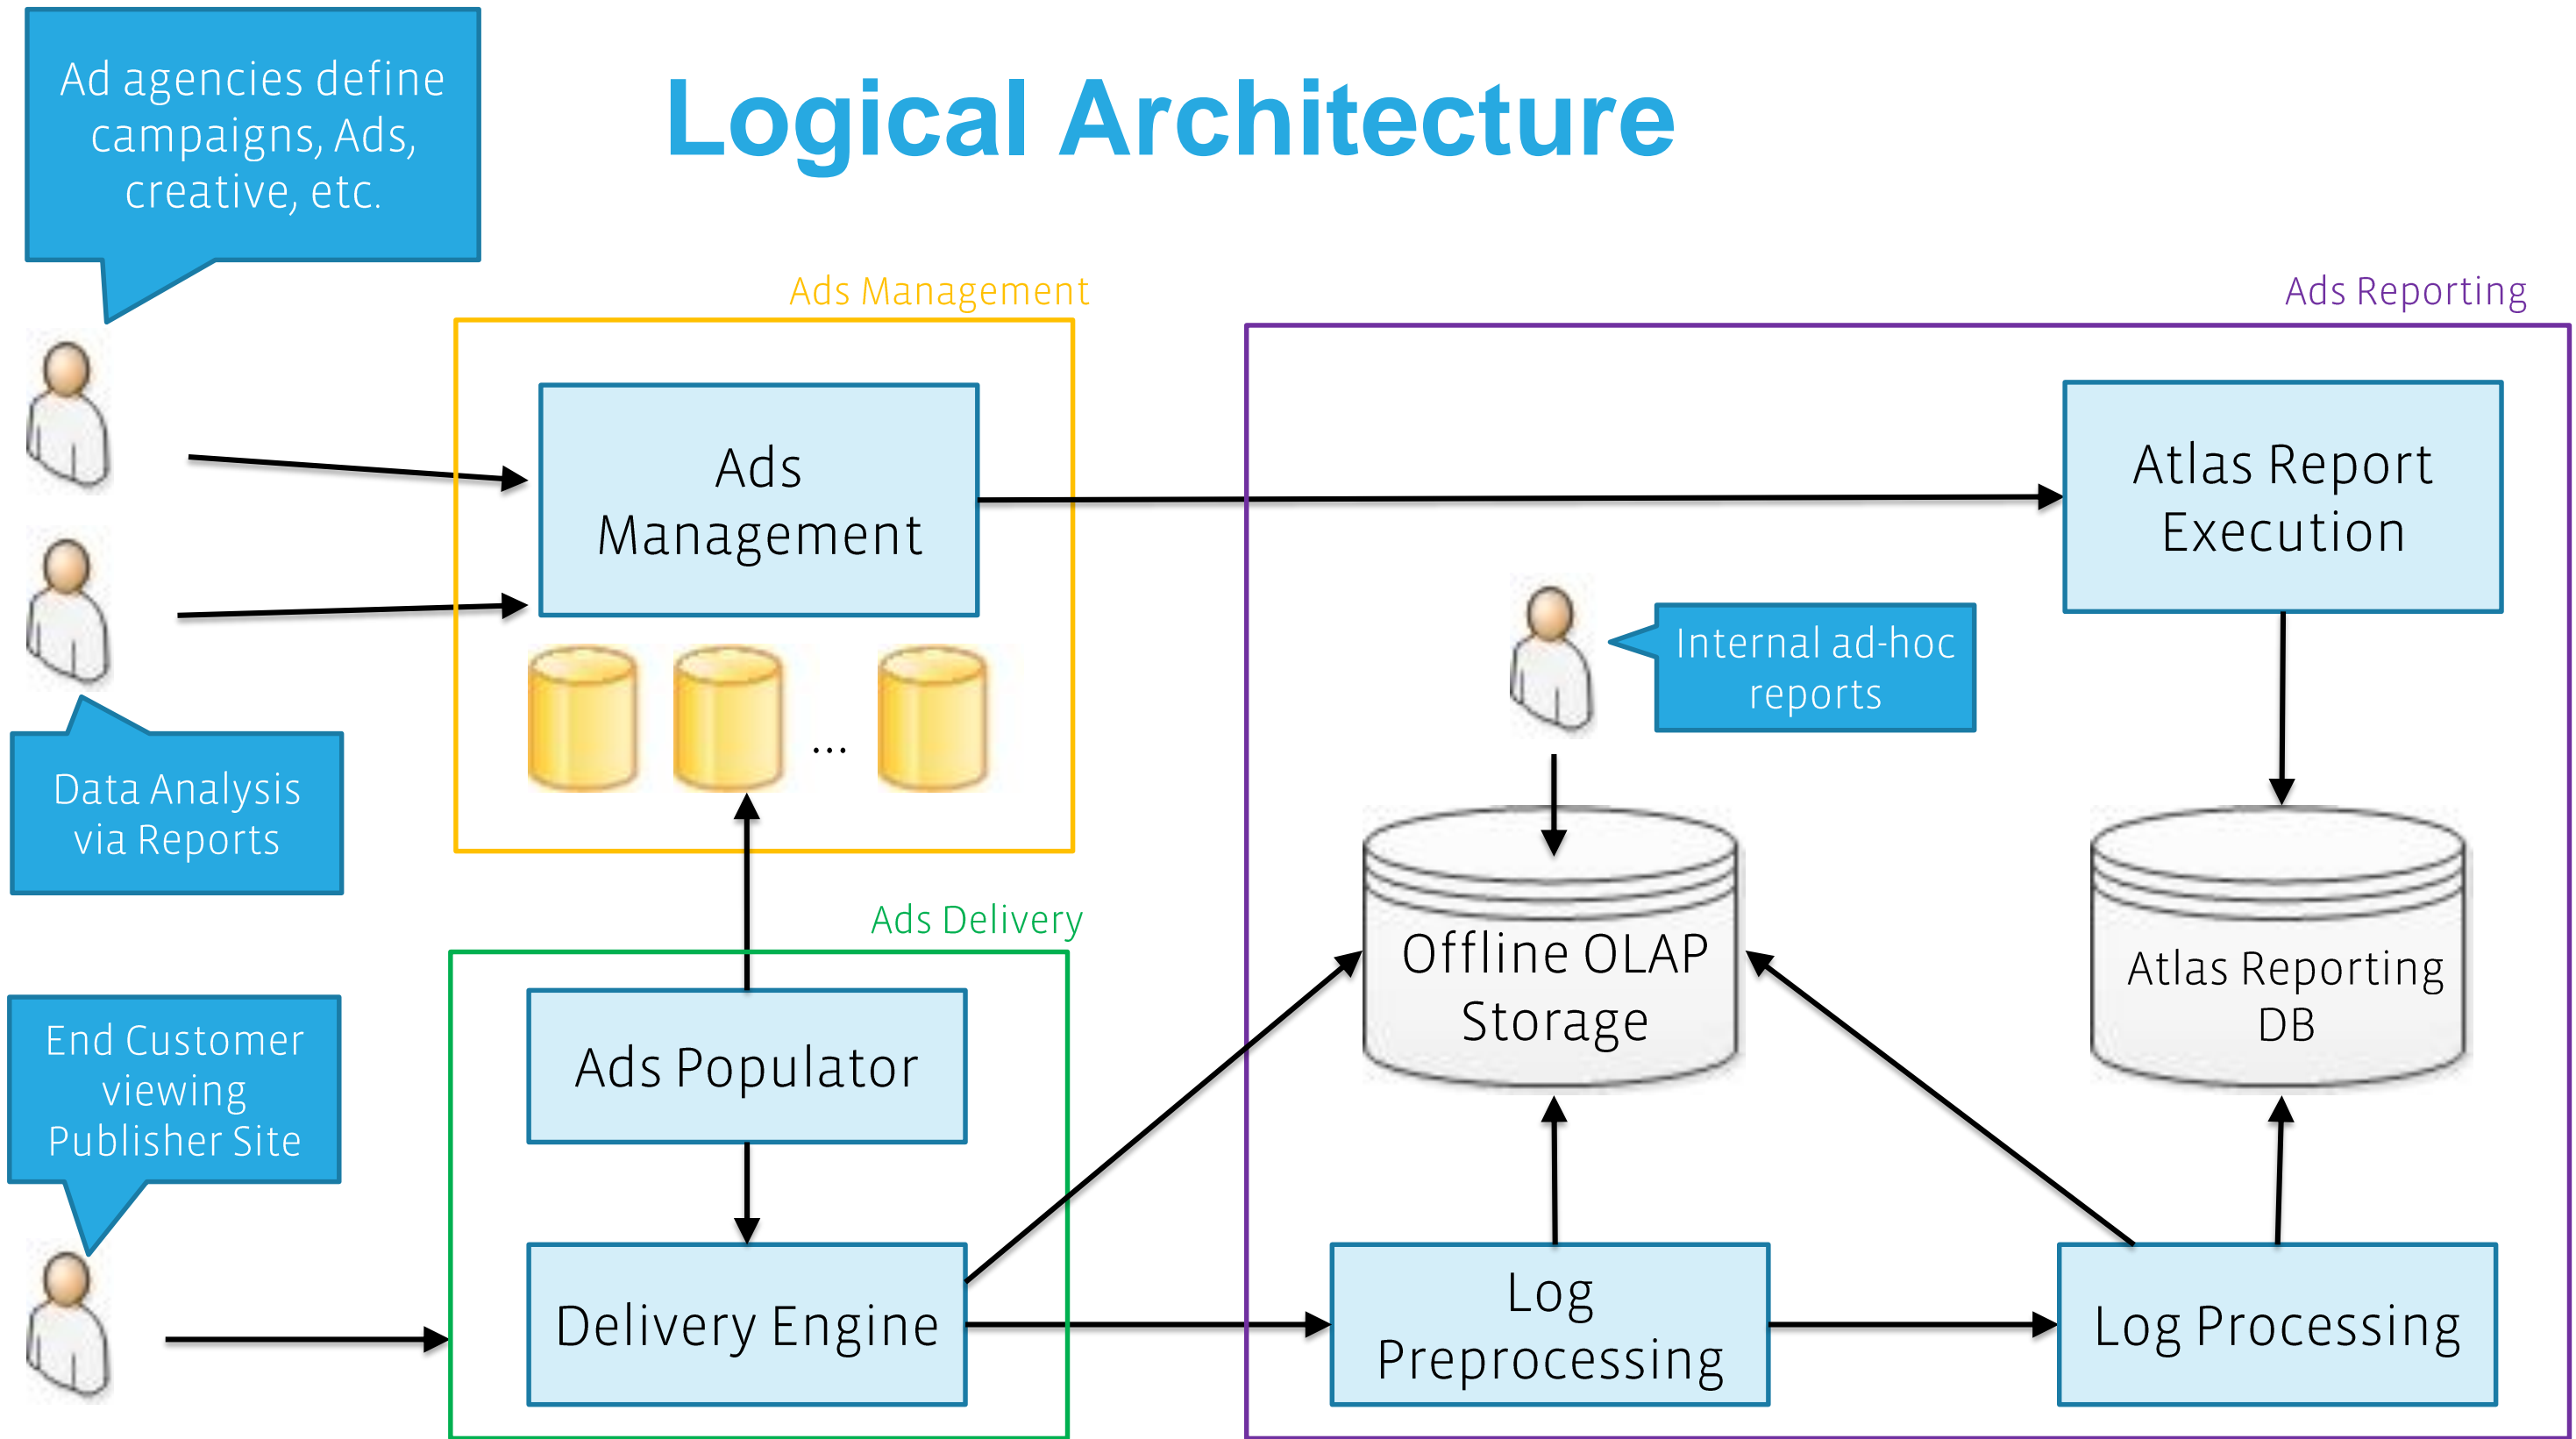
Ads Reporting



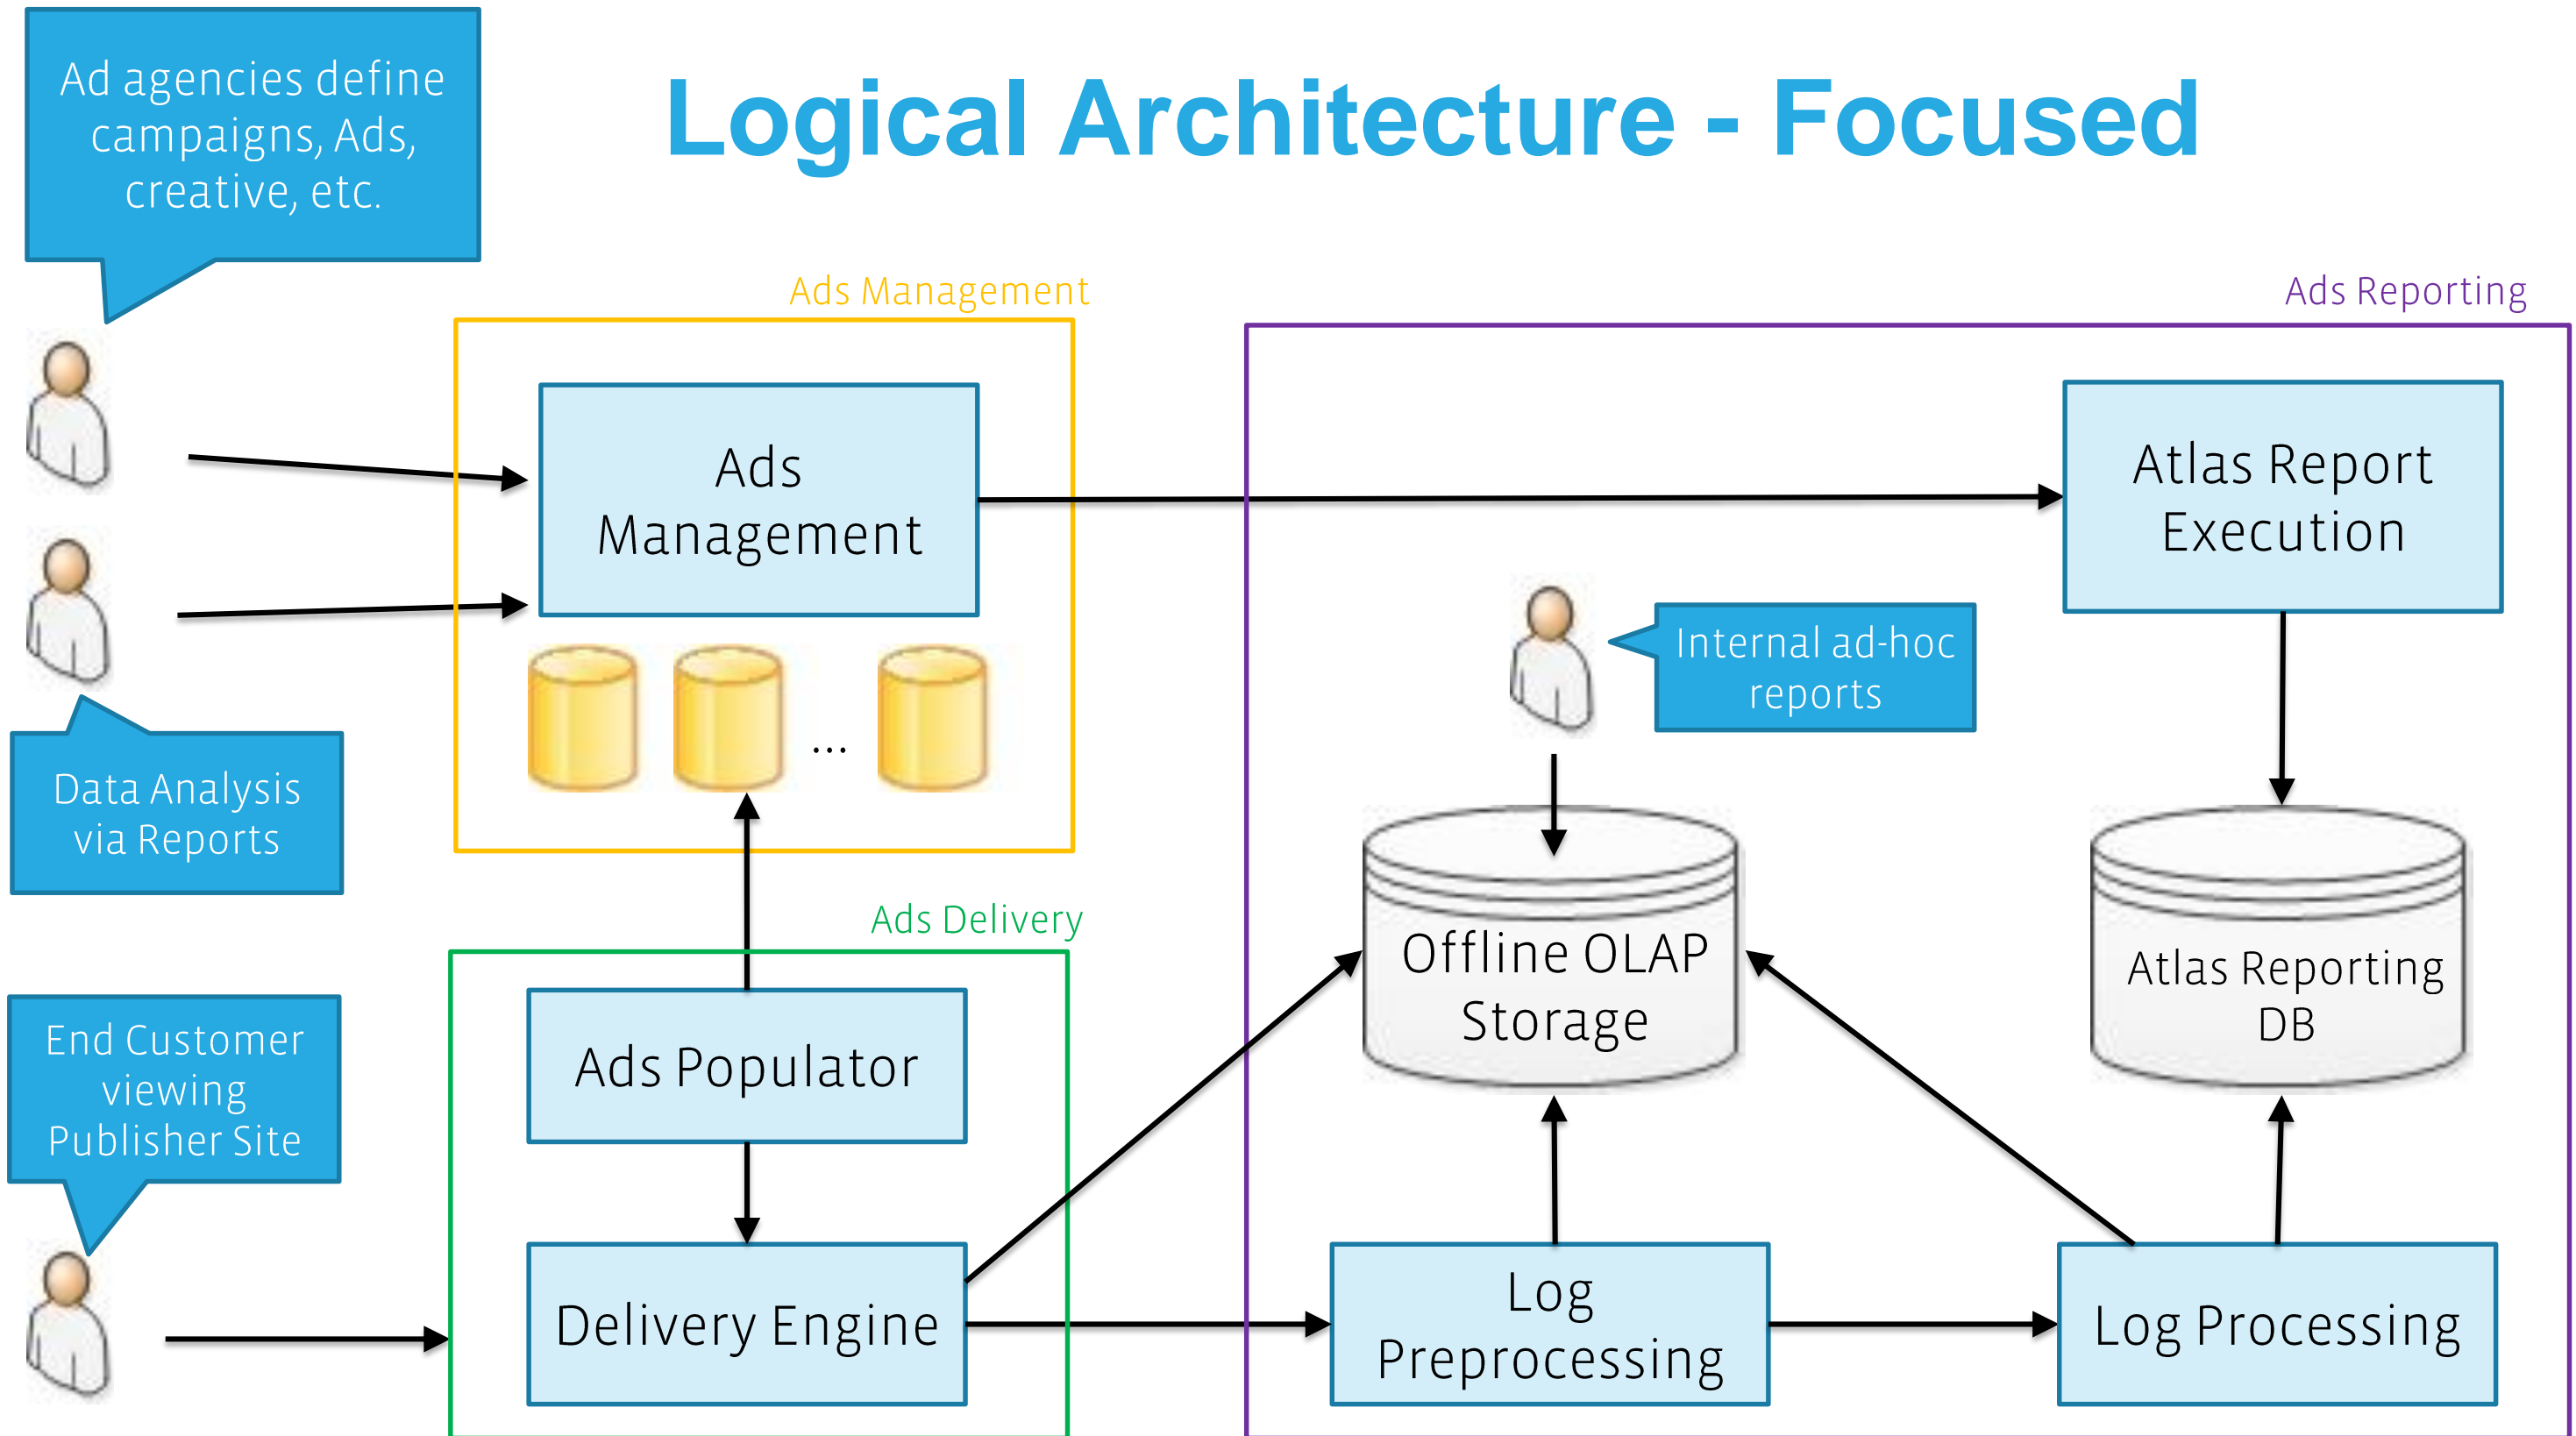
Logical Architecture



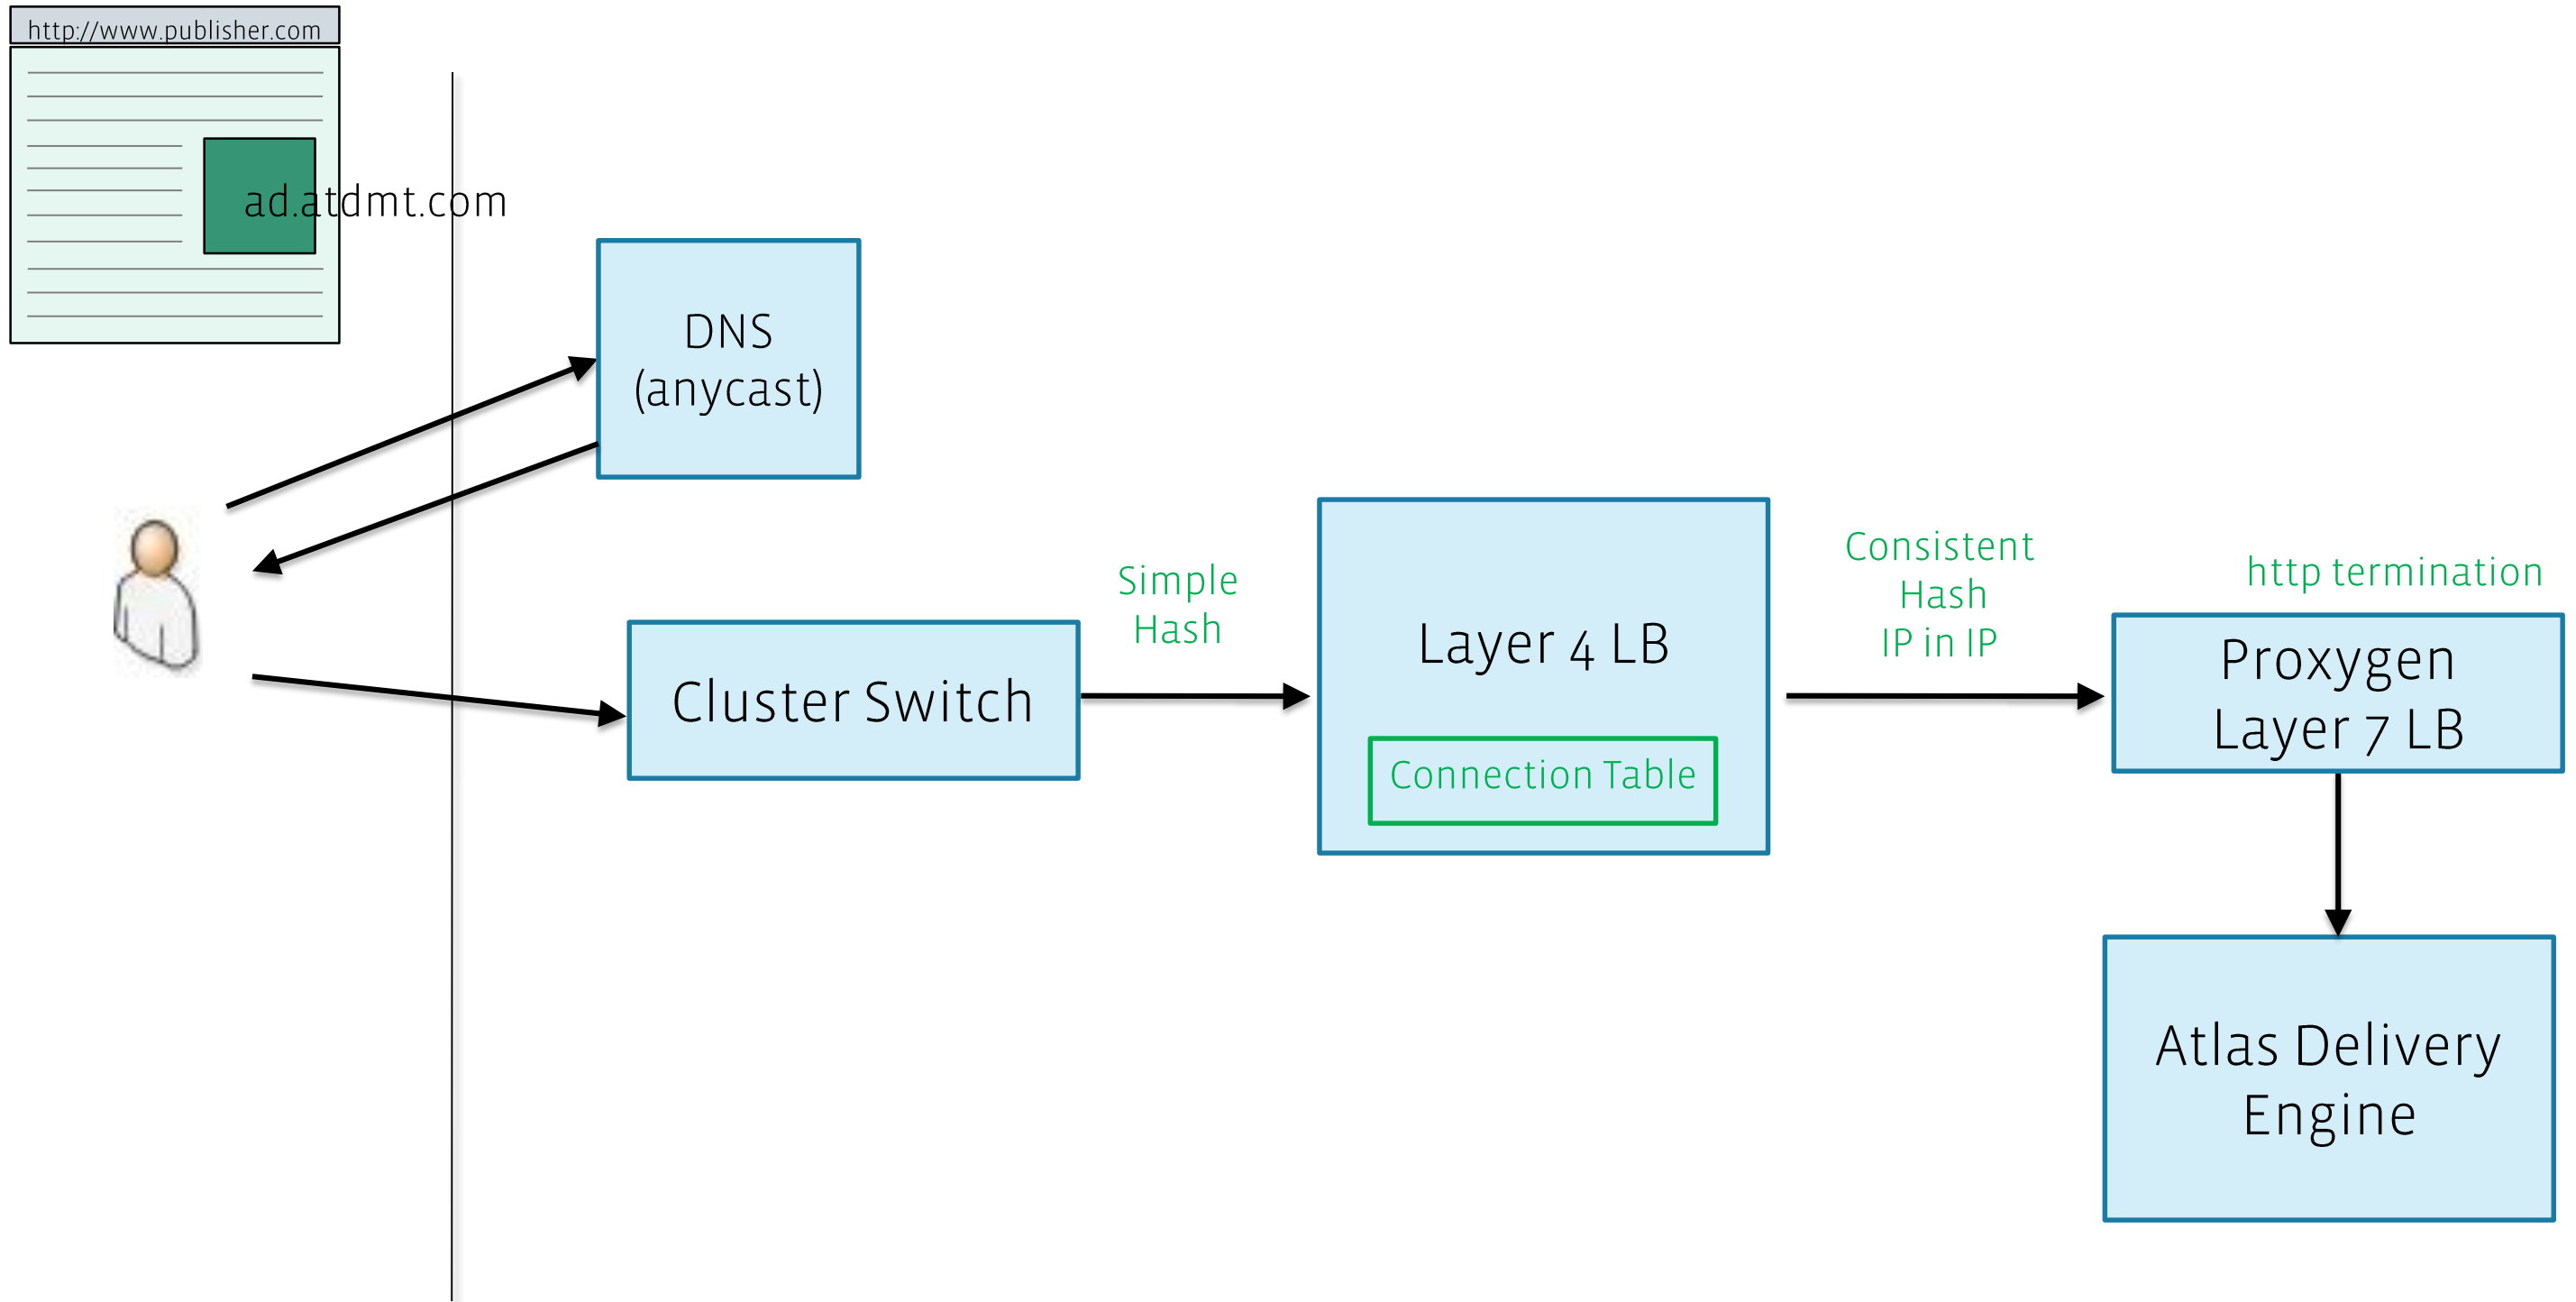
Logical Architecture



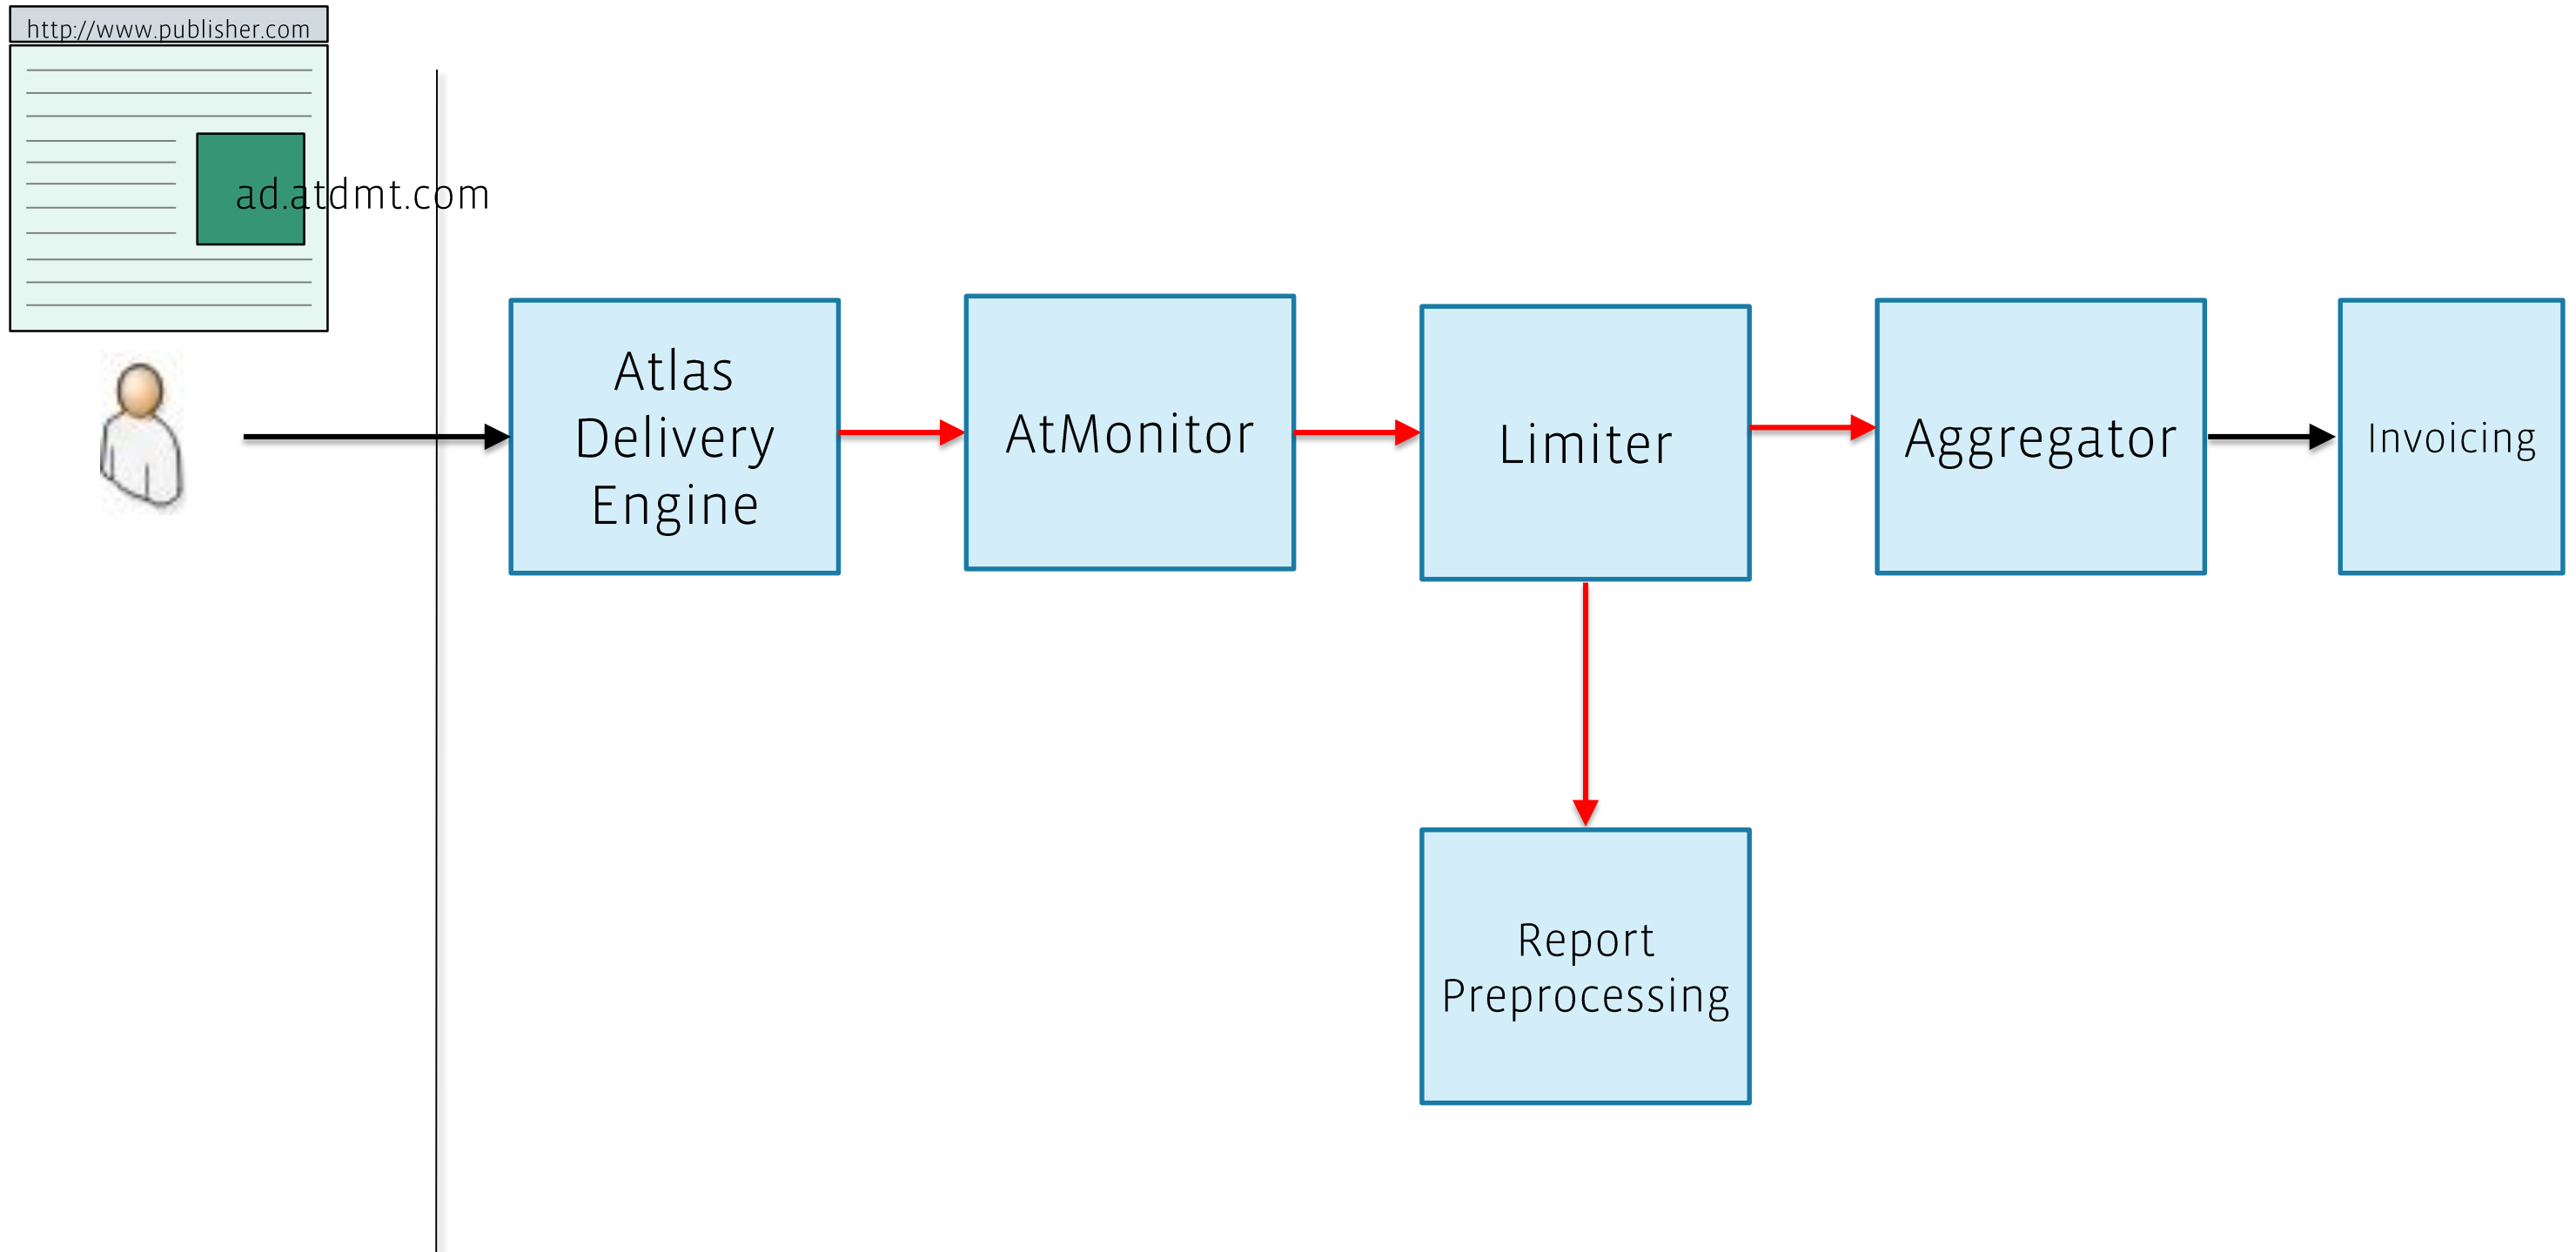
Logical Architecture - Focused



Physical Architecture – Ad Delivery



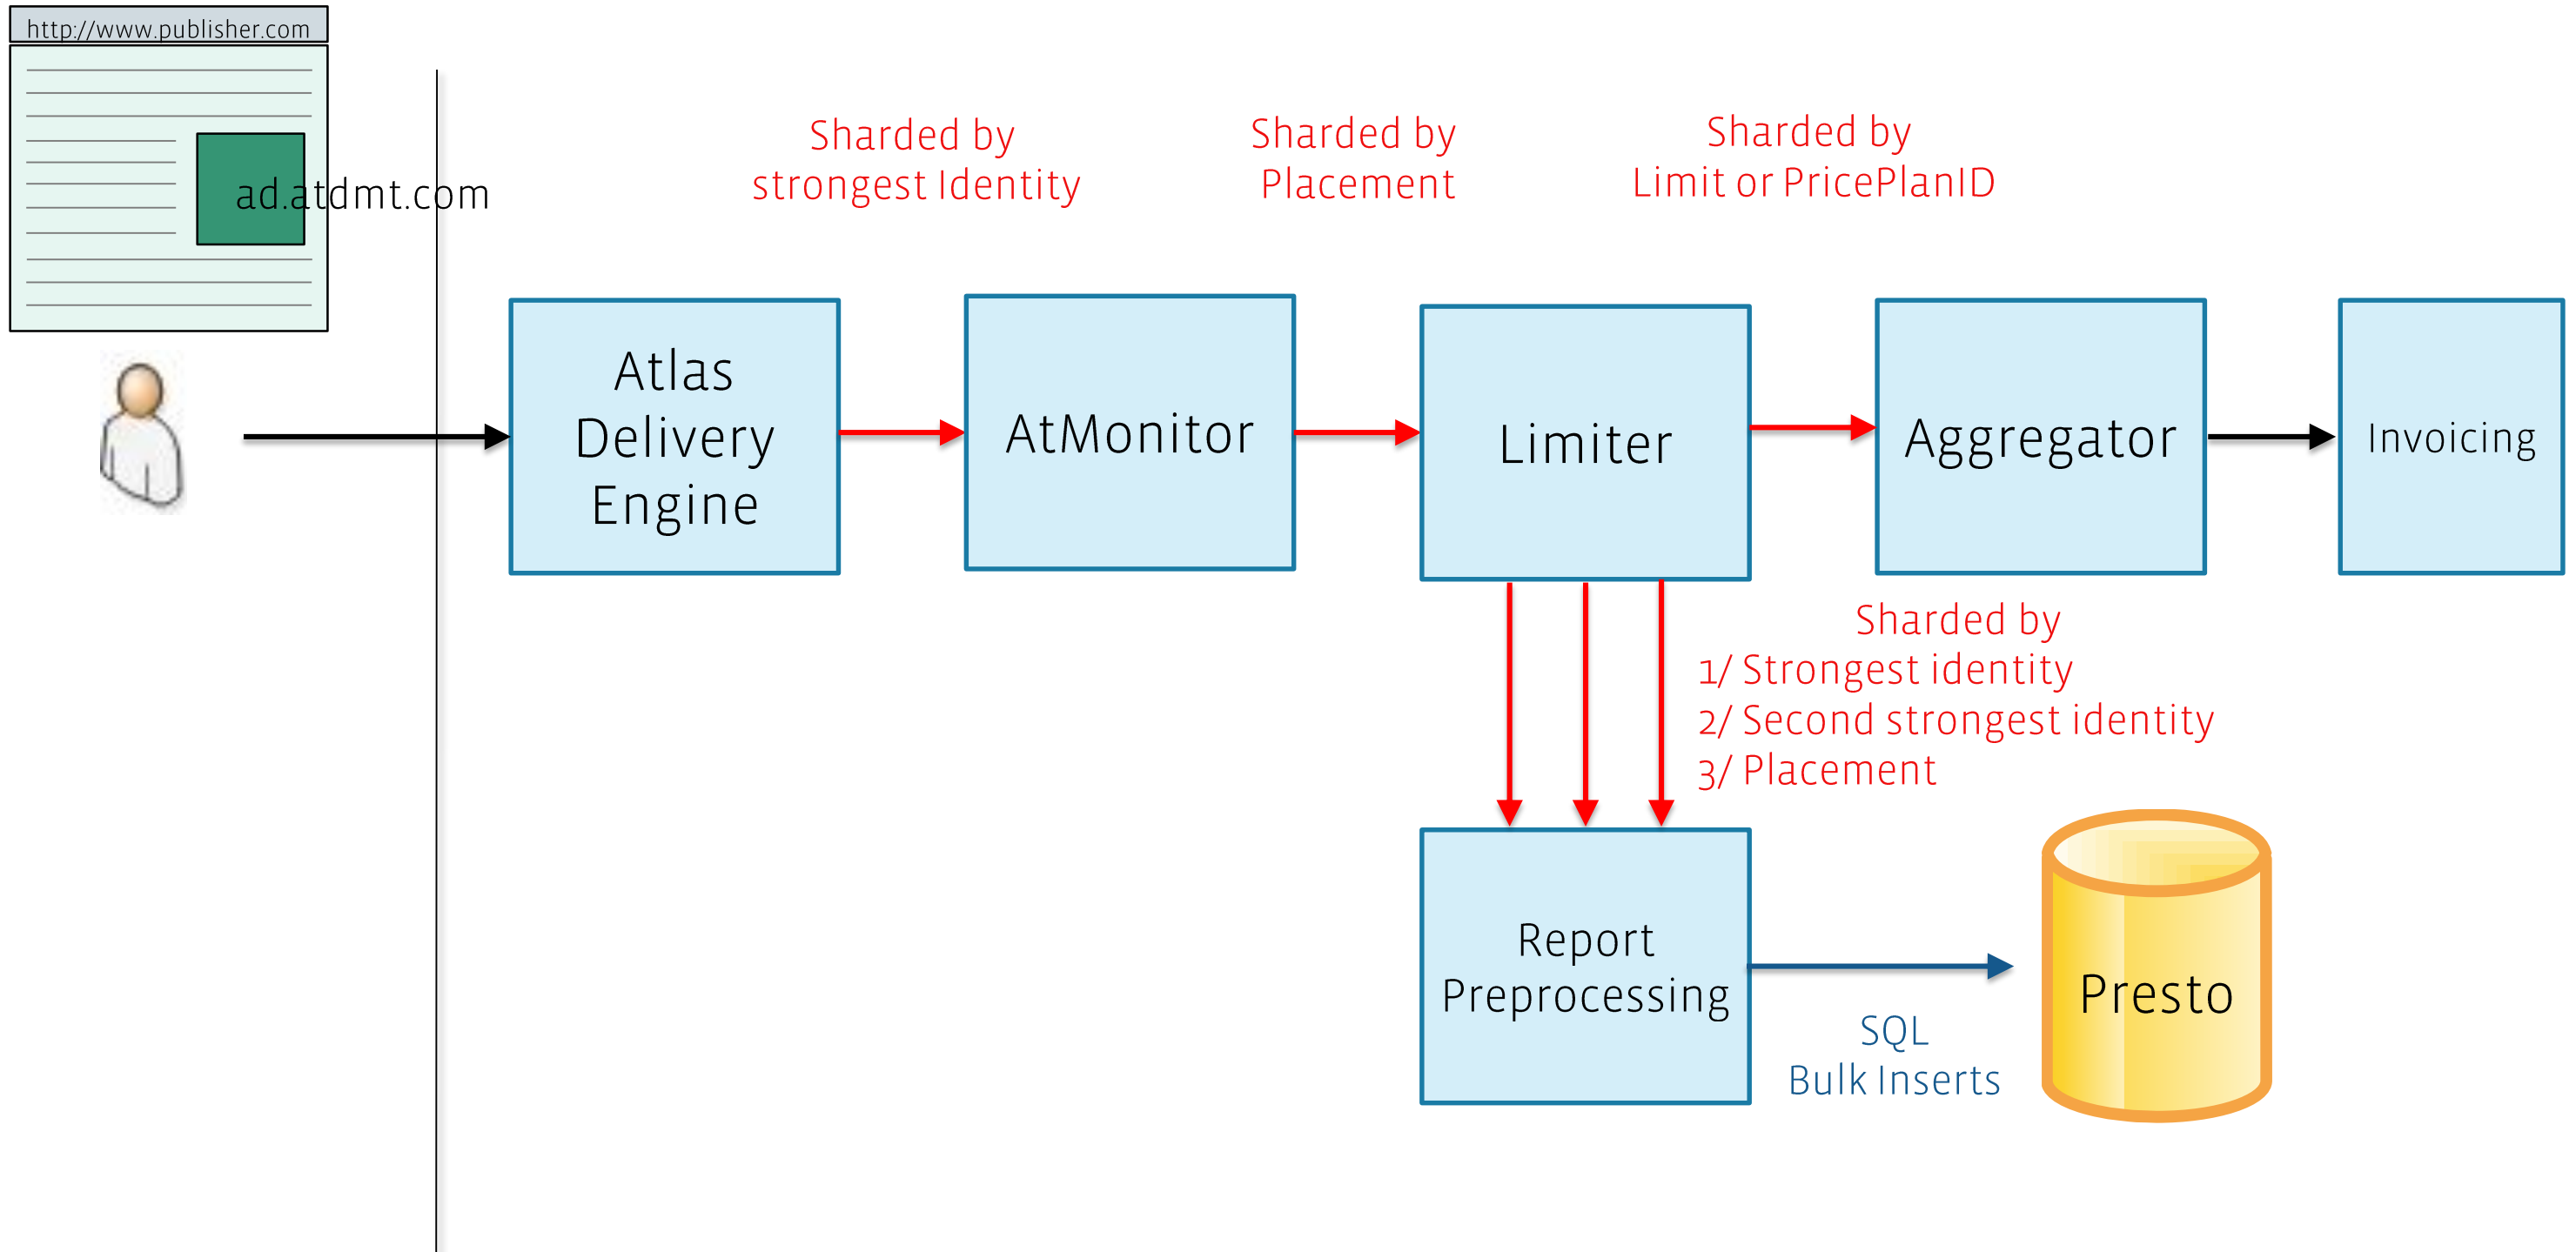
Physical Architecture – Data Processing Pipeline



Scribe

- Large-scale, high throughput message queue
- Not lossless but guarantees are excellent and perfect for us
- Decouples producers from consumers
- Persistent for n days
- Sharded consumption
- Checkpoint streams

Physical Architecture – Data Processing Pipeline



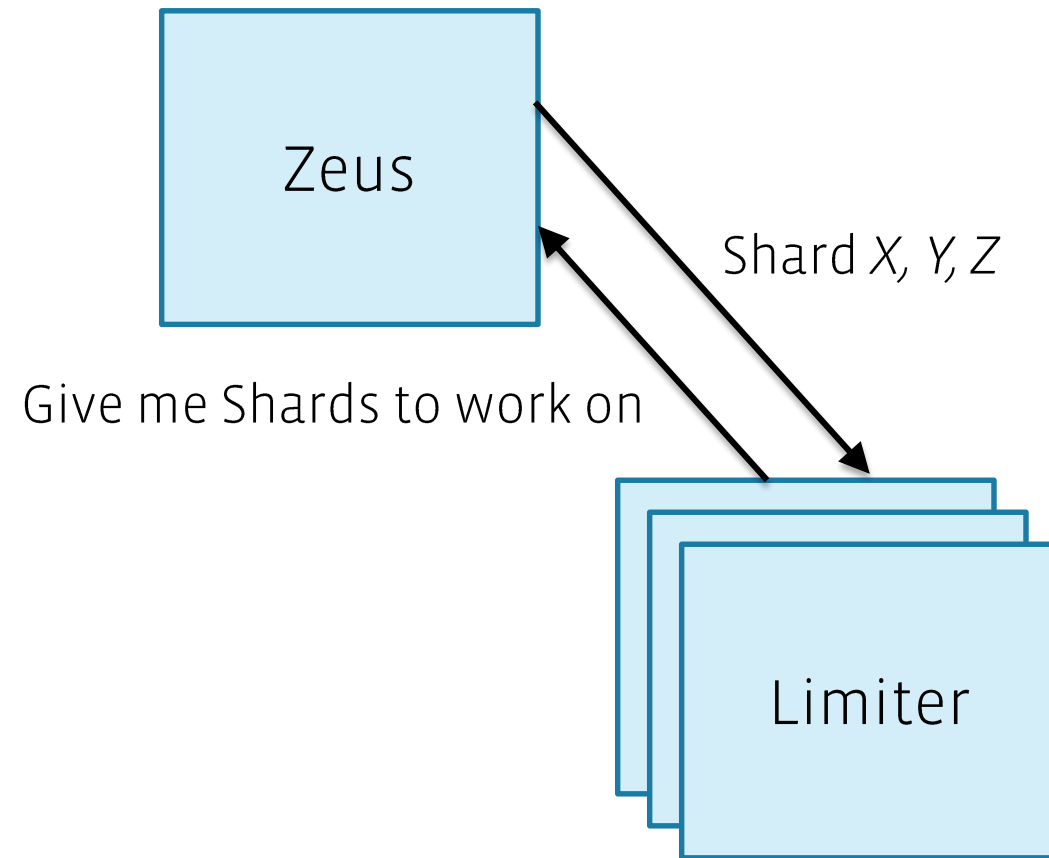
Detailed Look – Stream Processing

- Mini workflows
- Scalability is hard
 - Message queues can be costly
 - Repeatable re-execution

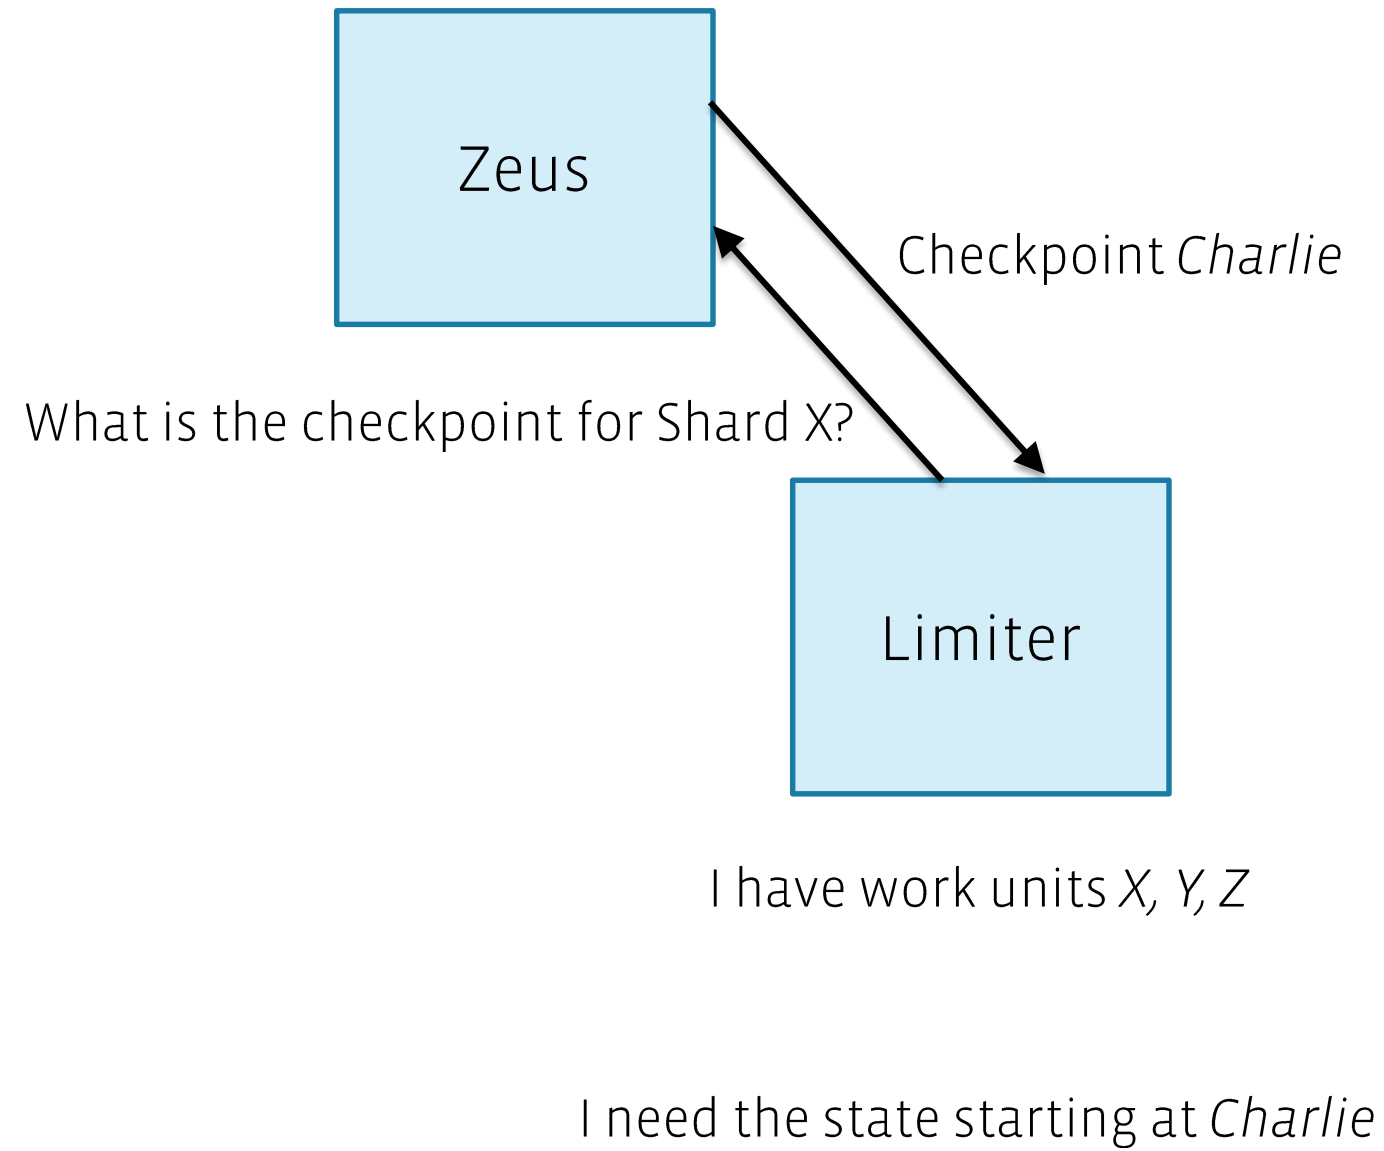
Physical Architecture



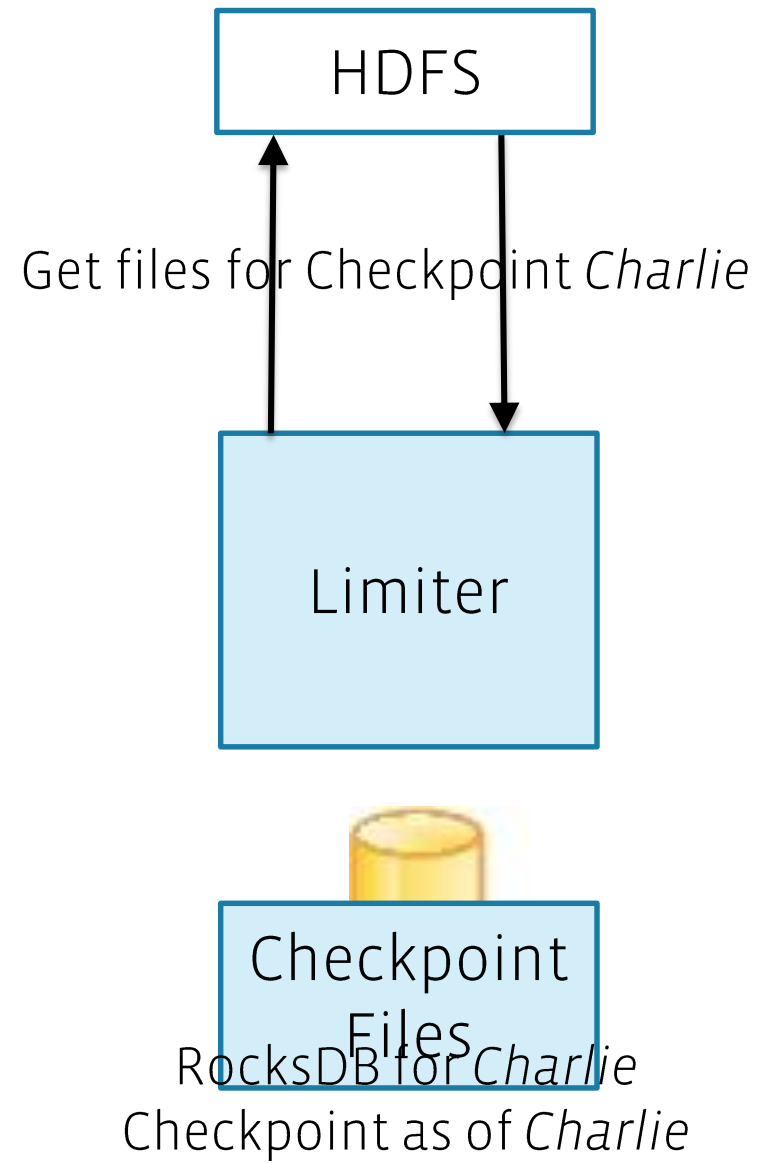
Physical Architecture



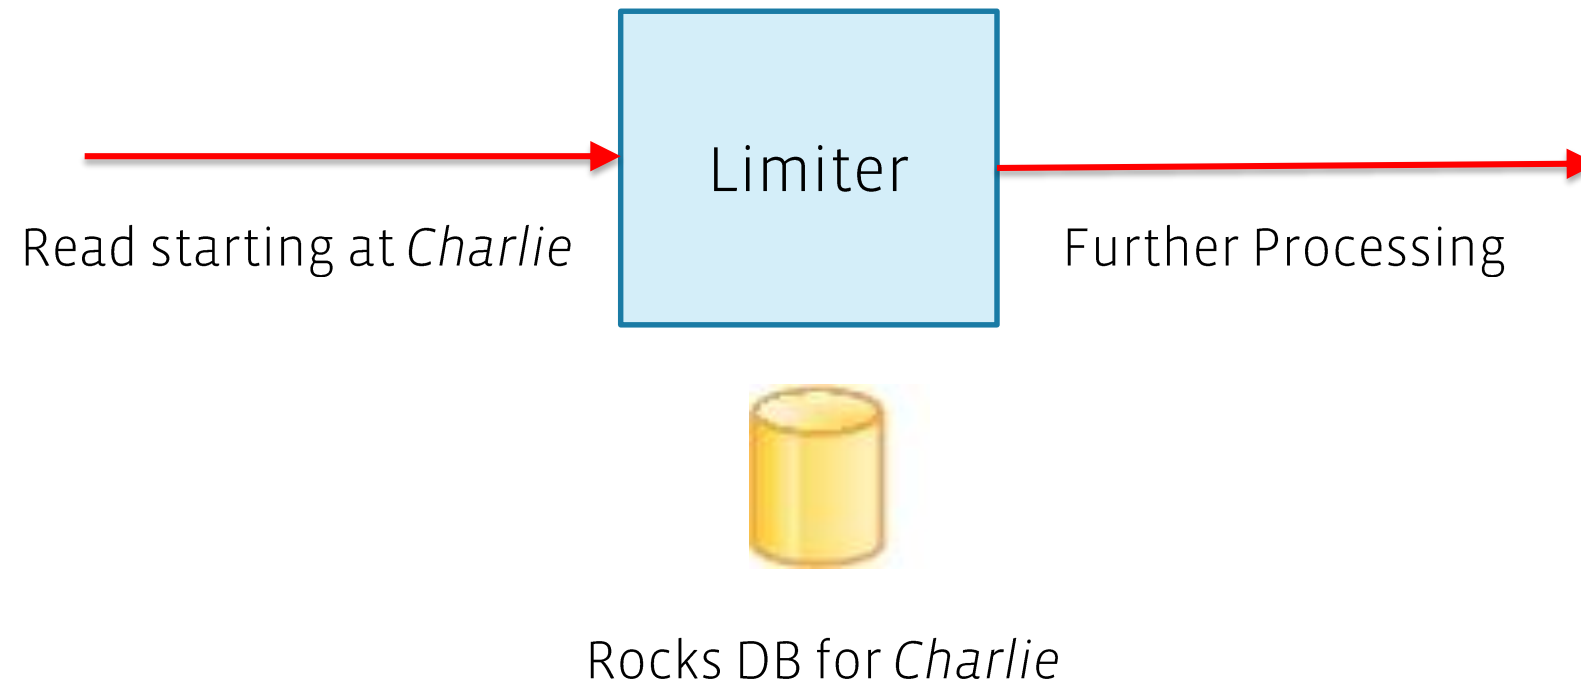
Physical Architecture



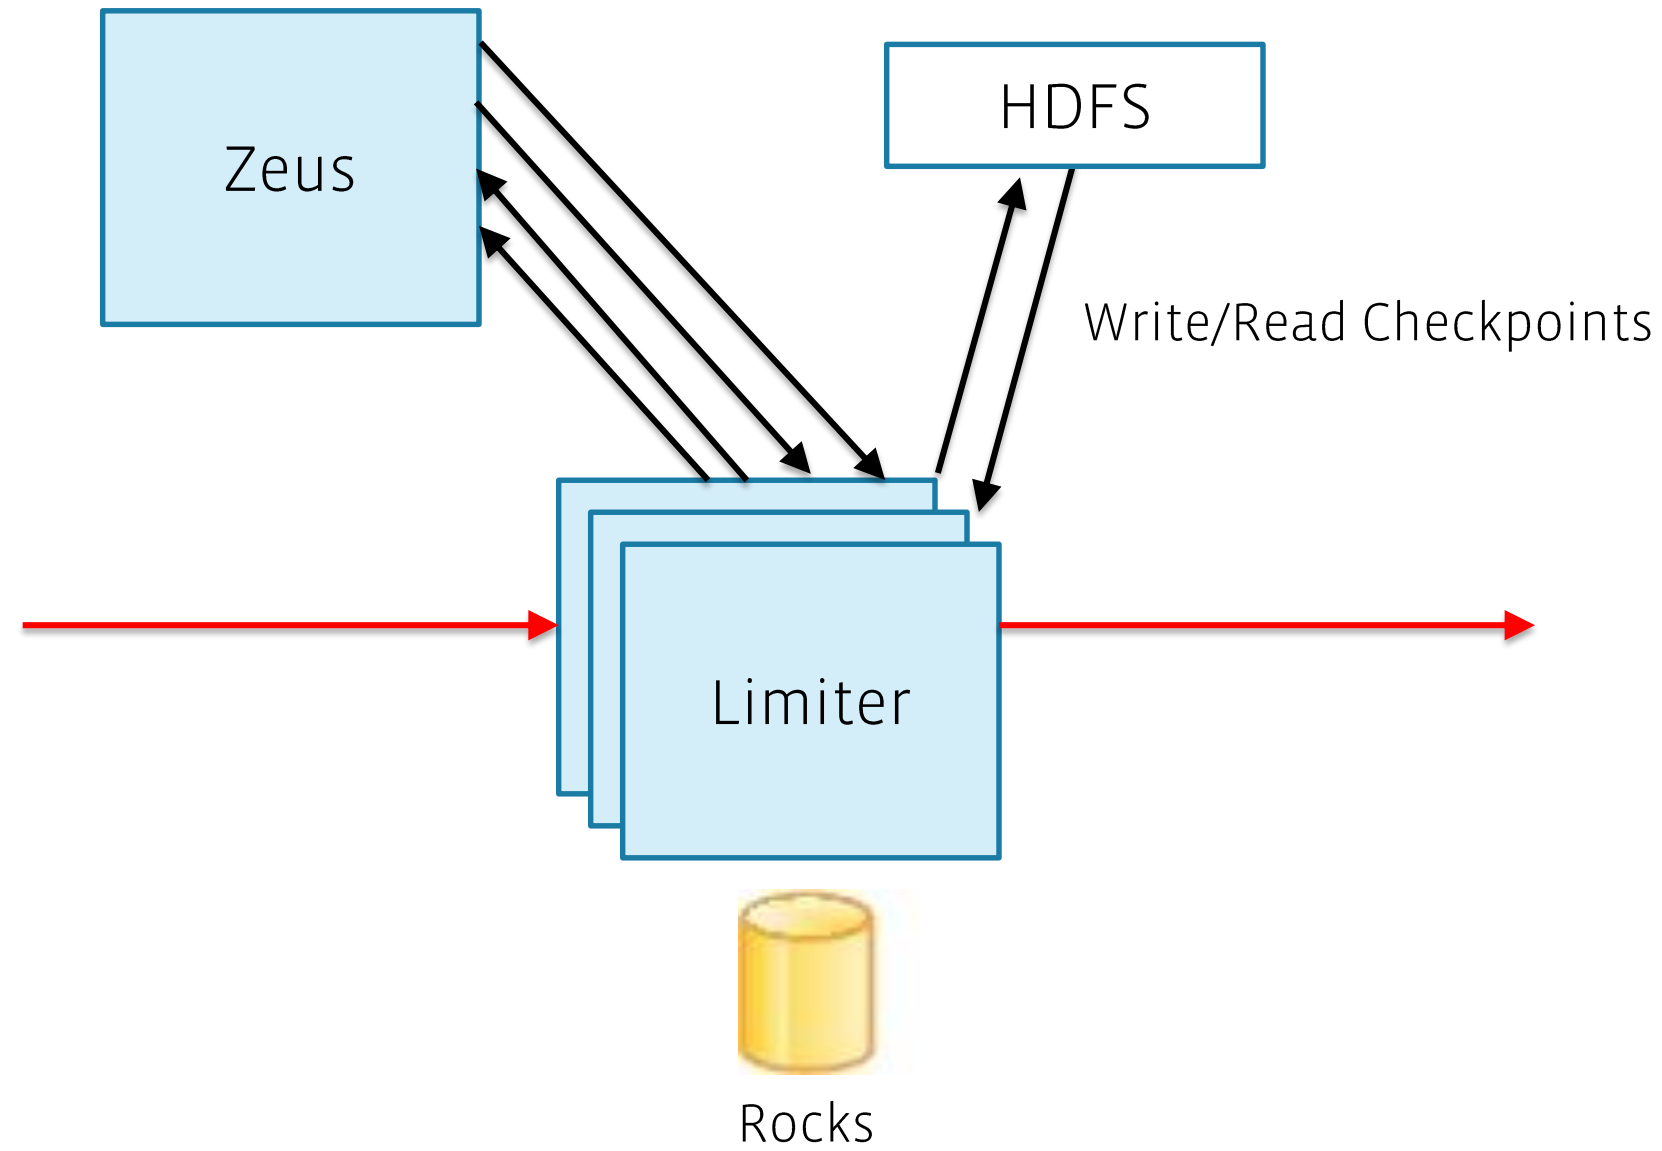
Physical Architecture



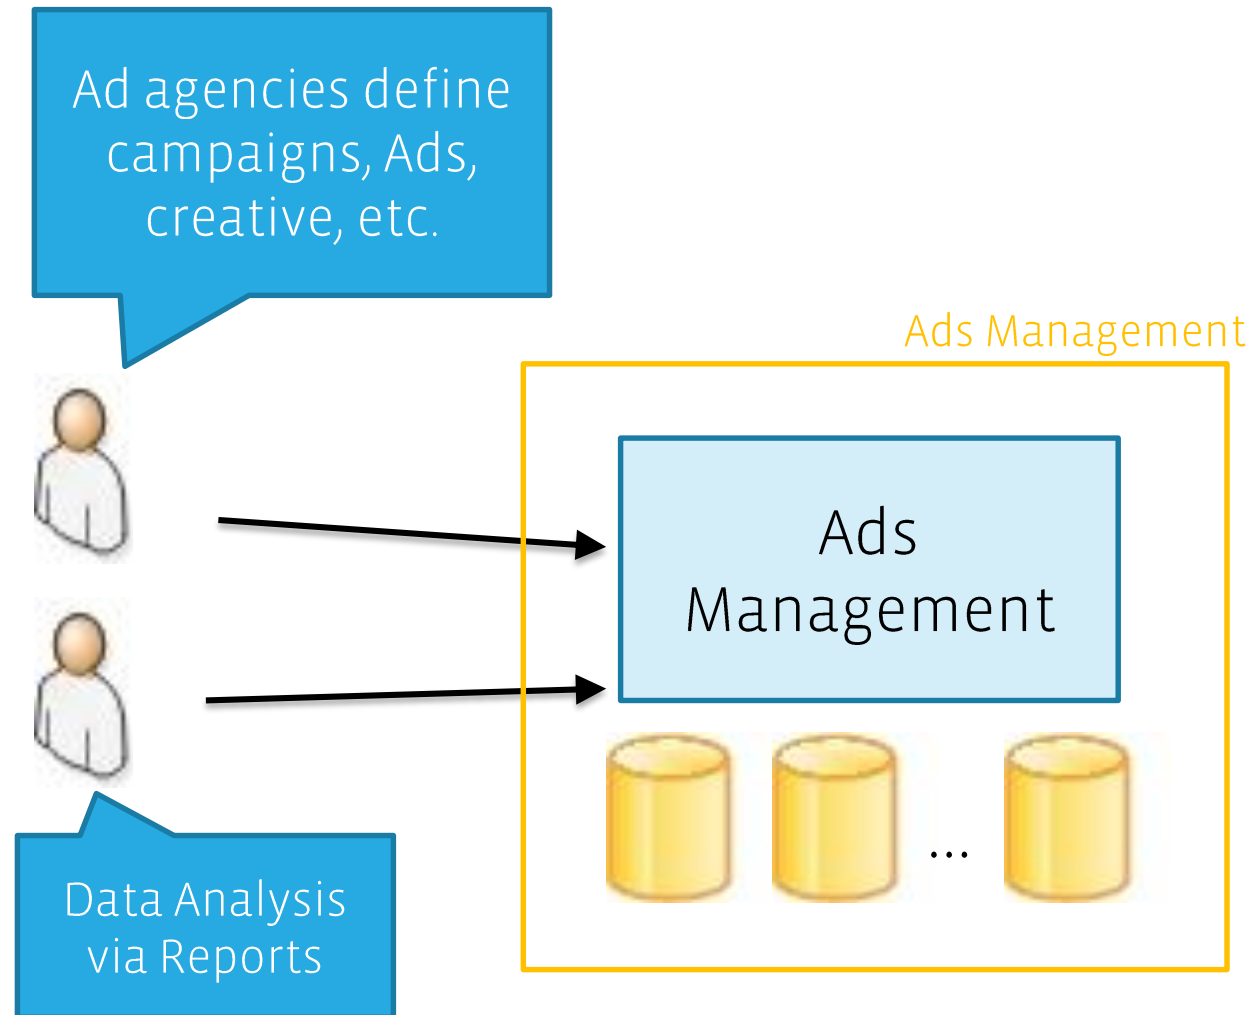
Physical Architecture



Physical Architecture

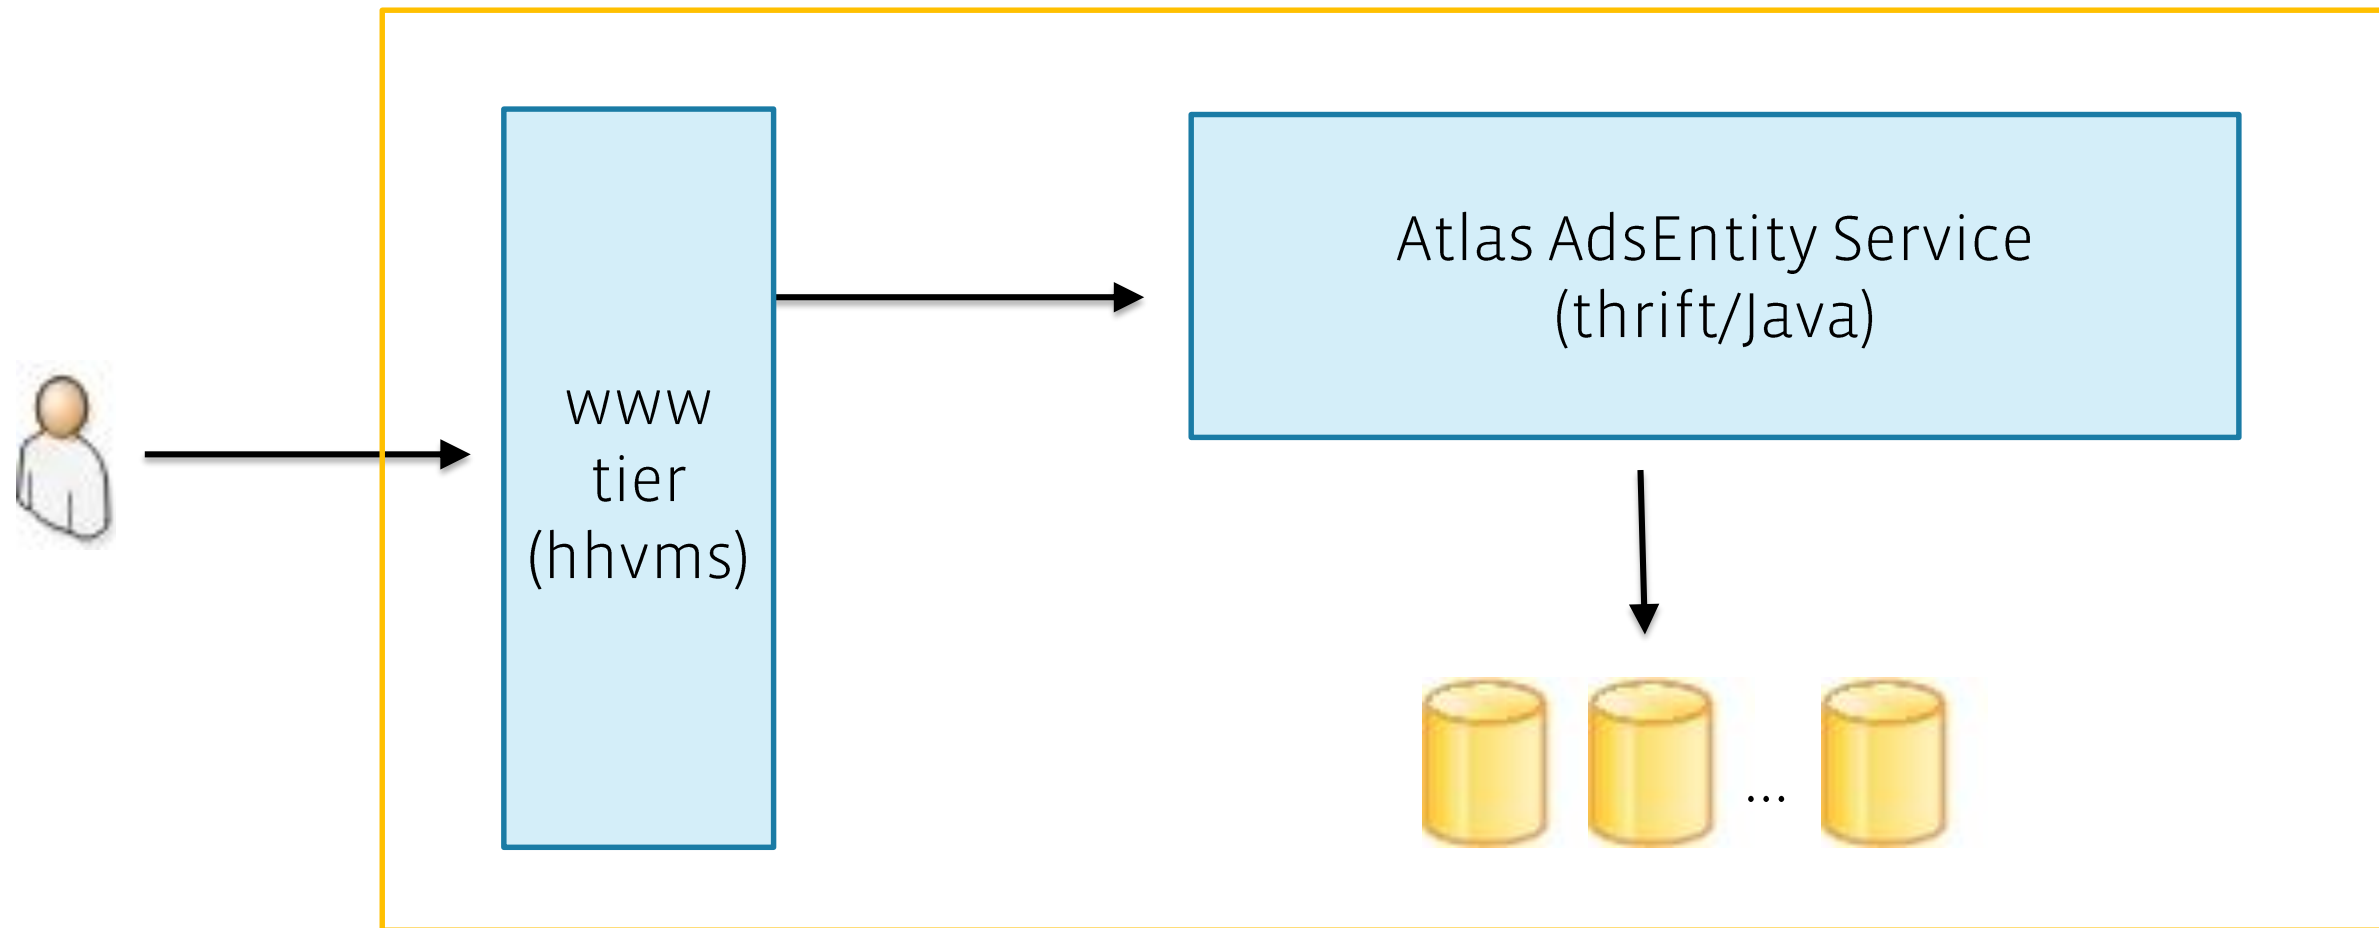


Lessons Learned



Mistake – Minimize the www/hack tier

Ads Management



Mistake – Minimize the www/hack tier

Goal

- Minimize the code in www tier to ensure a higher level of availability

Why this was a mistake

- Hack and the www tier have come a long way in 2 years
- Huge improvements in availability

The Impact

- API changes hit two separate systems
- Couldn't leverage all the improvements and investments in hack

Mistake – Looking ahead

- Lesson
 - Look ahead to where the industry is going
 - Or look to where an organization (or group) will invest efforts
 - Plan to meet them there

Questions?