



Big Data Problems In Monitoring At eBay

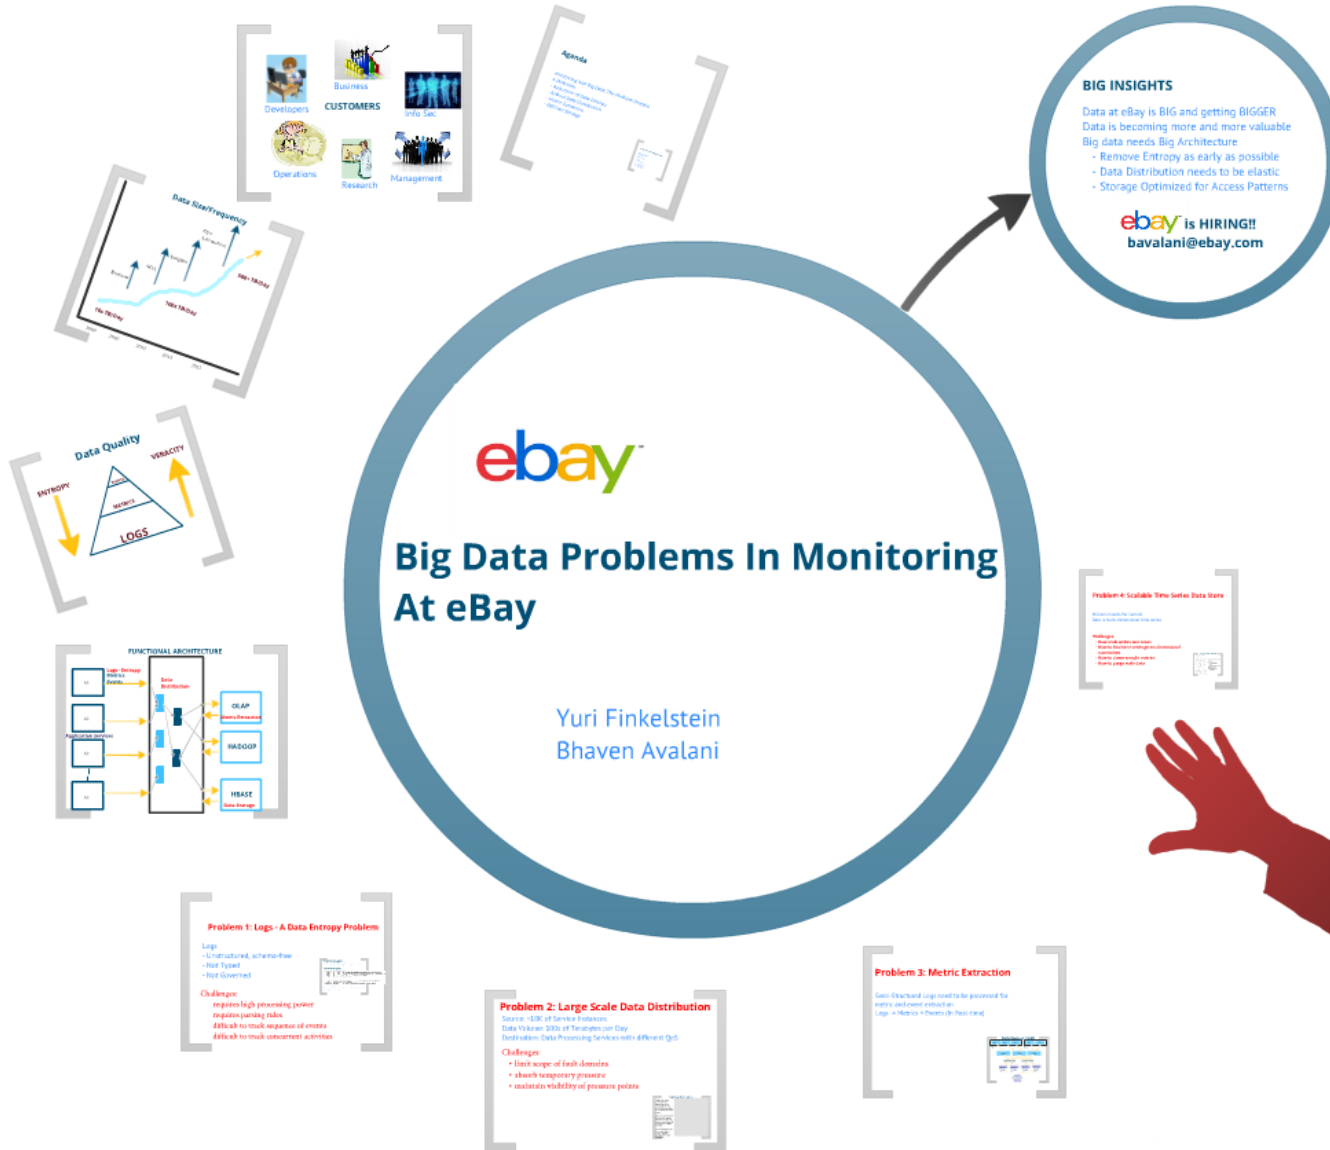
Yuri Finkelstein
Bhaven Avalani

BIG INSIGHTS

Data at eBay is BIG and getting BIGGER
Data is becoming more and more valuable
Big data needs Big Architecture

- Remove Entropy as early as possible
- Data Distribution needs to be elastic
- Storage Optimized for Access Patterns

ebay is HIRING!!
bavalani@ebay.com



Problem 1: Logs - A Data Entropy Problem

Logs

- Unstructured, unnormalized
- Not Typed
- Not Governed

Challenges

- requires high performing center
- requires parsing rules
- difficult to track separation of events
- difficult to track concurrent activities

Problem 2: Large Scale Data Distribution

Source - 10B of Service Instances
(Data Volume: 200 of Terabytes per Day)
Distribution: Data Processing Services with different QoS

Challenges

- limit scope of fault domains
- absorb temporary processes
- maintain visibility of pressure points

Problem 3: Metric Extraction

Source: Unstructured Logs need to be processed for metrics and need to be stored

Logs = Metrics + Pressure On Post-Processing

Problem 4: Scalable Time Series Data Store

Source: Unstructured Logs need to be processed for metrics and need to be stored

Challenges

- Real-time access to data
- Elasticity (scale up/down)
- Query (complex queries)
- Multi-tenant support





Big Data Problems In Monitoring At eBay


Yuri Finkelstein
Bhaven Avalani

BIG INSIGHTS

Data at eBay is BIG and getting BIGGER
Data is becoming more and more valuable
Big data needs Big Architecture

- Remove Entropy as early as possible
- Data Distribution needs to be elastic
- Storage Optimized for Access Patterns

ebay is HIRING!!
bavalani@ebay.com



Problem 1: Logs - A Data Entropy Problem

Log

- Unstructured, unnormalized
- Not Typed
- Not Governed

Challenges

- requires high performing center
- requires parsing rules
- difficult to track sequence of events
- difficult to track concurrent activities

Problem 2: Large Scale Data Distribution

Source - 10B of Service Instances
(Data Volume: 200 of Terabytes per Day)
Distribution - Data Processing Services with different QoS

Challenges

- limit scope of fault domains
- absorb temporary pressure
- maintain visibility of pressure points

Problem 3: Metric Extraction

Source - Unstructured Logs need to be processed for metrics and needs extracted as...

Log -> Metrics -> Present On Host -> Ring

Problem 4: Scalable Time Series Data Store

Source - Unstructured Logs need to be processed for metrics and needs extracted as...

Log -> Metrics -> Present On Host -> Ring



Agenda

Monitoring And Big Data: The Problem Domain

4 Problems

- Reduction of Data Entropy
- Robust Data Distribution
- Metric Extraction
- Efficient Storage

Monitoring: The Problem Domain

- Operational Management
 - Performance
 - Errors
 - Anomaly Detection
- Triaging
 - Root Cause Analysis
- Business Monitoring
 - Customer Behaviors
 - Click Stream Analytics

Monitoring: The Problem Domain

Operational Management

- Performance
- Errors
- Anomaly Detection

Triaging

- Root Cause Analysis

Business Monitoring

- Customer Behaviors
- Click Stream Analytics



Developers



Business



Info Sec

CUSTOMERS



Operations

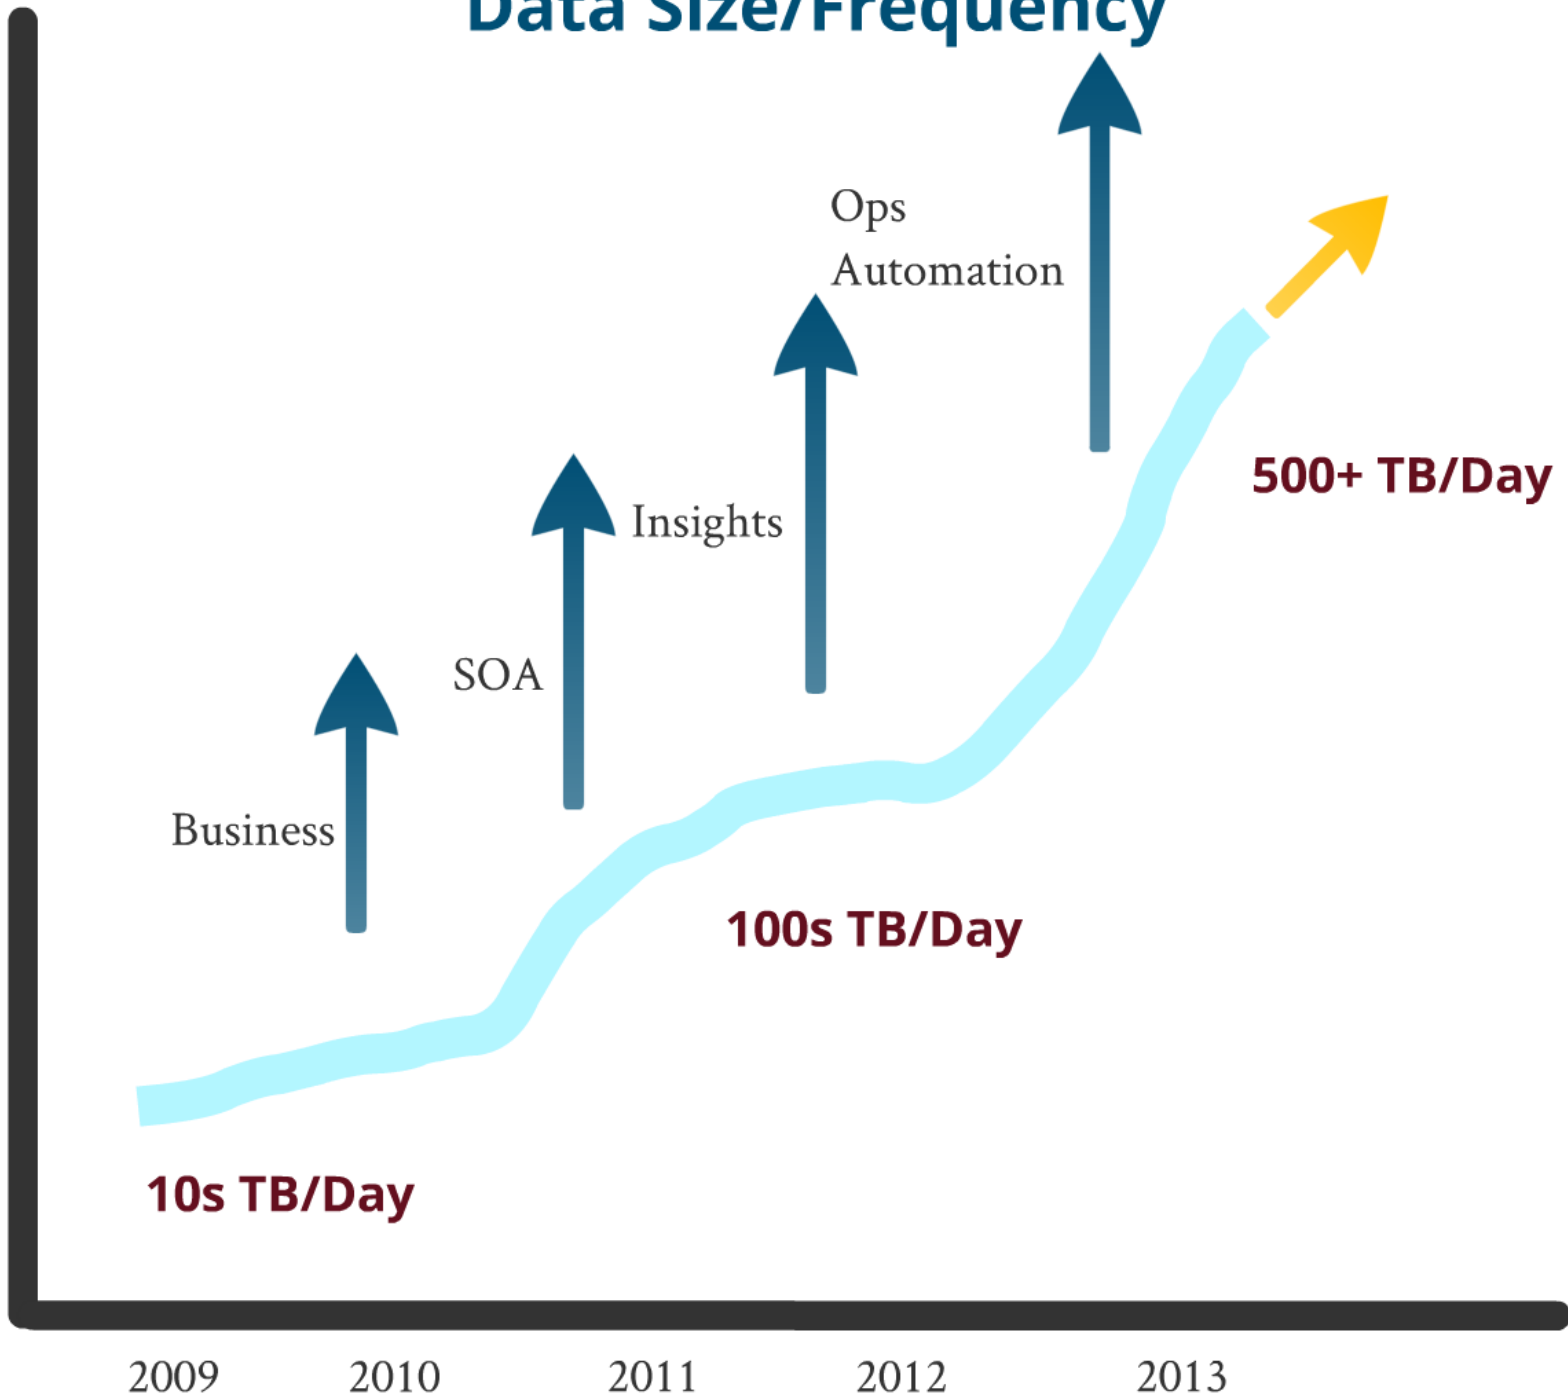


Research




Management

Data Size/Frequency



Data Quality

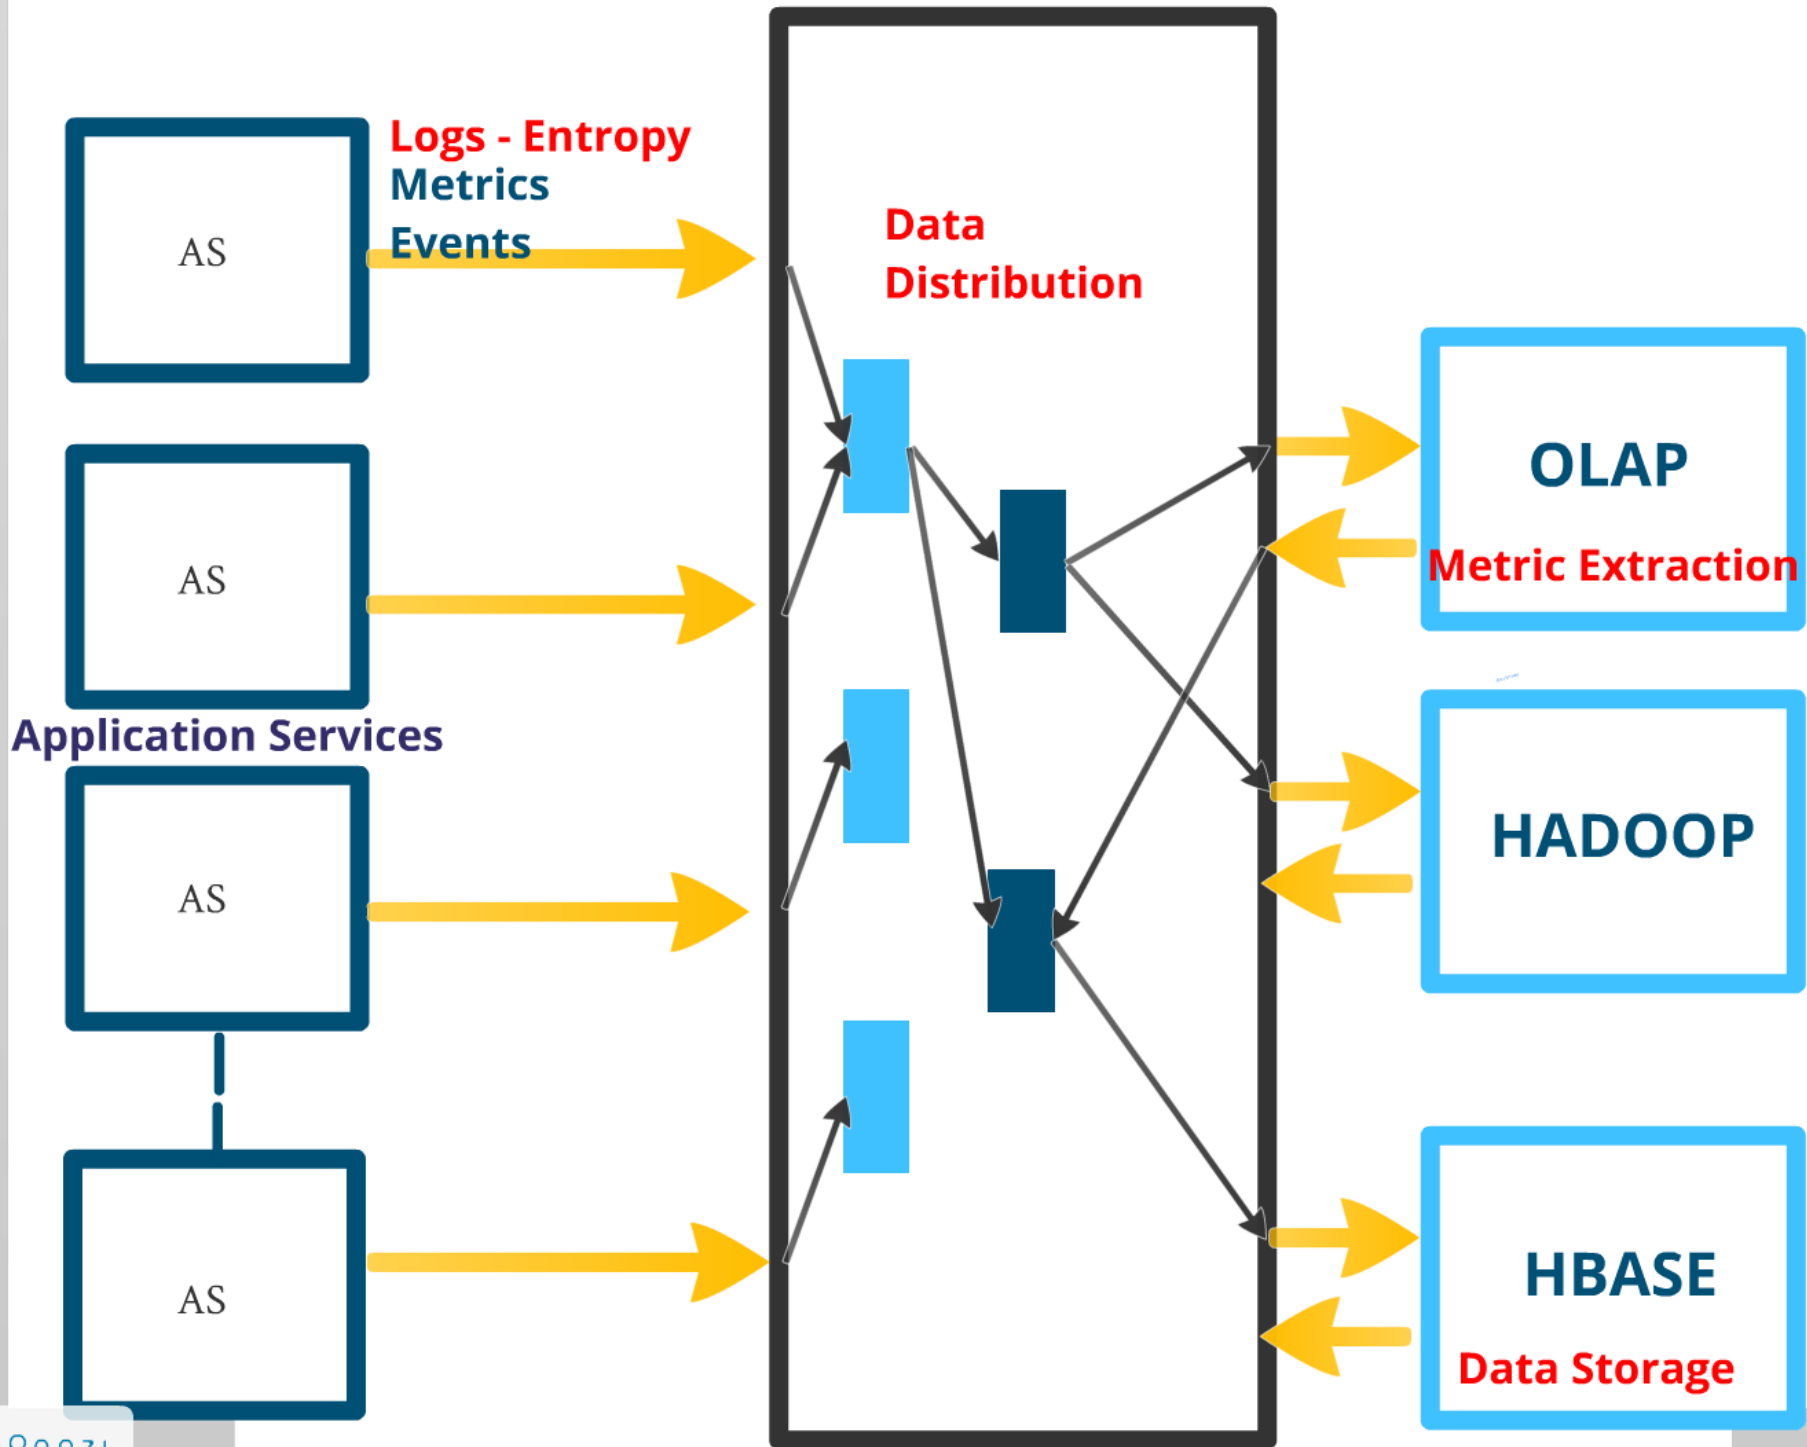
ENTROPY



VERACITY



FUNCTIONAL ARCHITECTURE



Traditional Logging

```
[Wed Oct 11 14:32:52 2000] [error] [client 127.0.0.1] client denied by server configuration:
/export/home/live/ap/htdocs/test
127.0.0.1 - frank [10/Oct/2000:13:55:36 -0700] "GET /apache_pb.gif HTTP/1.0" 200 2326
```

Structured Logging

```
-----
0 t21:48:14.01 URL Ginger.m.1.GET
1 E21:48:14.01 URL Request 0 /m/1/pictures/getObjectIdForMd5?md5=mGcMs8xD_j8SBgAq2Bx2AQ&widt
1 E21:48:14.01 URL ClientInfo 0 ClientIP=10.91.119.29&
1 E21:48:14.01 URL Headers 0 accept=application/json&cache-control=no-cache&pragma=no-cache&u
connection=keep-alive&x-ebay-client-ip=10.94.110.91&x-ebay-web-tier-ip=10.91.119.65&
1 A21:48:14.01 SEC Ginger.m.1.Security 0 0 SecurityEnabled=false&InternalIP=false
1 t21:48:14.03 Metadata GET Request com.ebay.zoom.eps.Image
2 t21:48:14.03 MetadataFind com.ebay.zoom.eps.dao.ImageDo
3 t21:48:14.03 PictureFailover Process GET Request
4 A21:48:14.03 EXEC 40305040 0 16 image05host
3 T21:48:14.04 PictureFailover Process GET Request 0 17 Got request for Oracle lookup&Su
-----
2 T21:48:14.04 MetadataFind com.ebay.zoom.eps.dao.ImageDo 0 17 key=mGcMs8xD_j8SBgAq2Bx
-----
1 T21:48:14.04 Metadata GET Request com.ebay.zoom.eps.Image 0 17
-----
0 T21:48:14.04 URL Ginger.m.1.GET 0 18 REQUEST-ID=13ac9bf6-2be0-a5ac-4594-42a4ffffd96d!m.1!10.9
```

t - open transaction (time, type, name)

T - close transaction (time, type, name, status_code, duration, data (name=value&name=value&...))

A - atomic transaction (like t and T in one line)

E - event (time, class, name, status_code, data)

H - heartbeat (time, class, name, data)

Transactions can be nested,

Problem 2: Large Scale Data Distribution

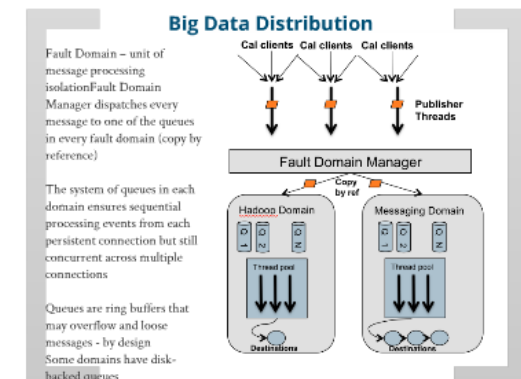
Source: ~10K of Service Instances

Data Volume: 100s of Terabytes per Day

Destination: Data Processing Services with different QoS

Challenges:

- limit scope of fault domains
- absorb temporary pressure
- maintain visibility of pressure points



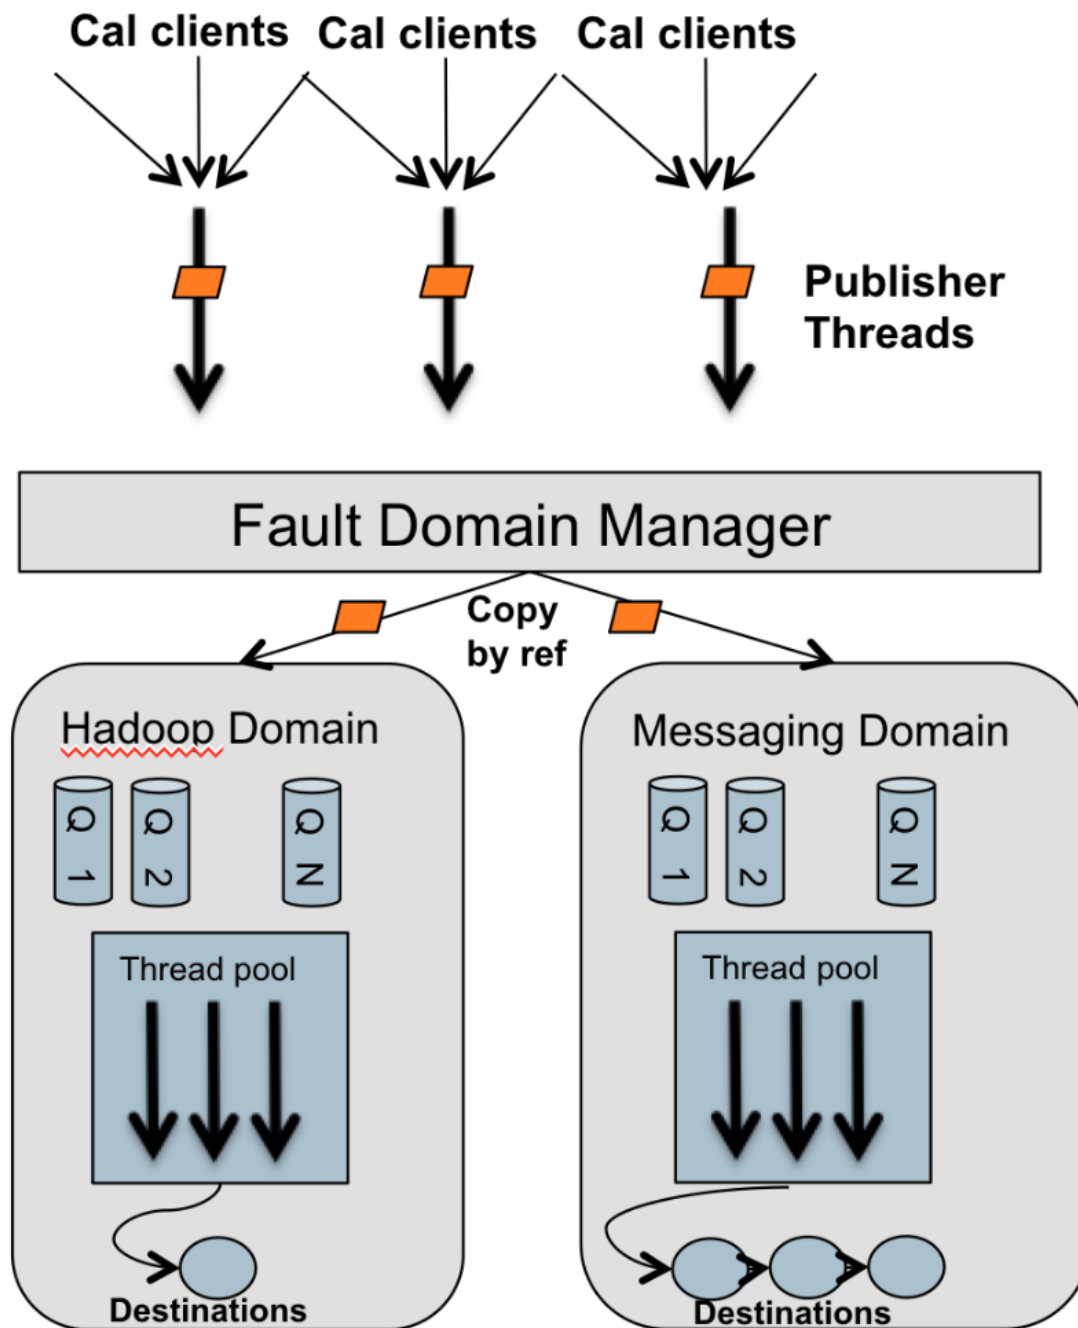
Big Data Distribution

Fault Domain – unit of message processing isolation
Fault Domain Manager dispatches every message to one of the queues in every fault domain (copy by reference)

The system of queues in each domain ensures sequential processing events from each persistent connection but still concurrent across multiple connections

Queues are ring buffers that may overflow and lose messages - by design

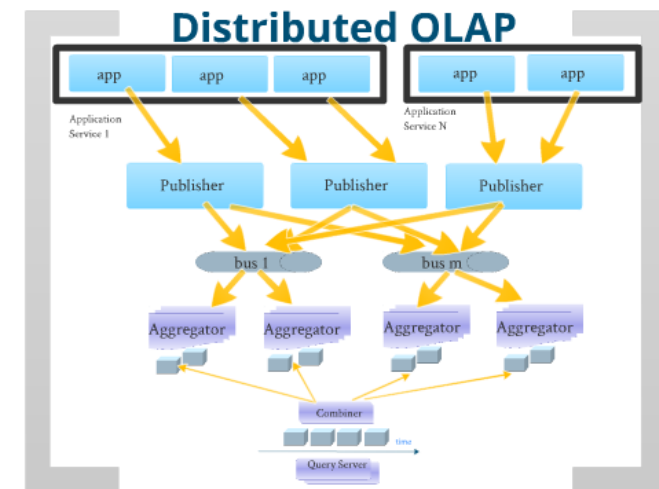
Some domains have disk-



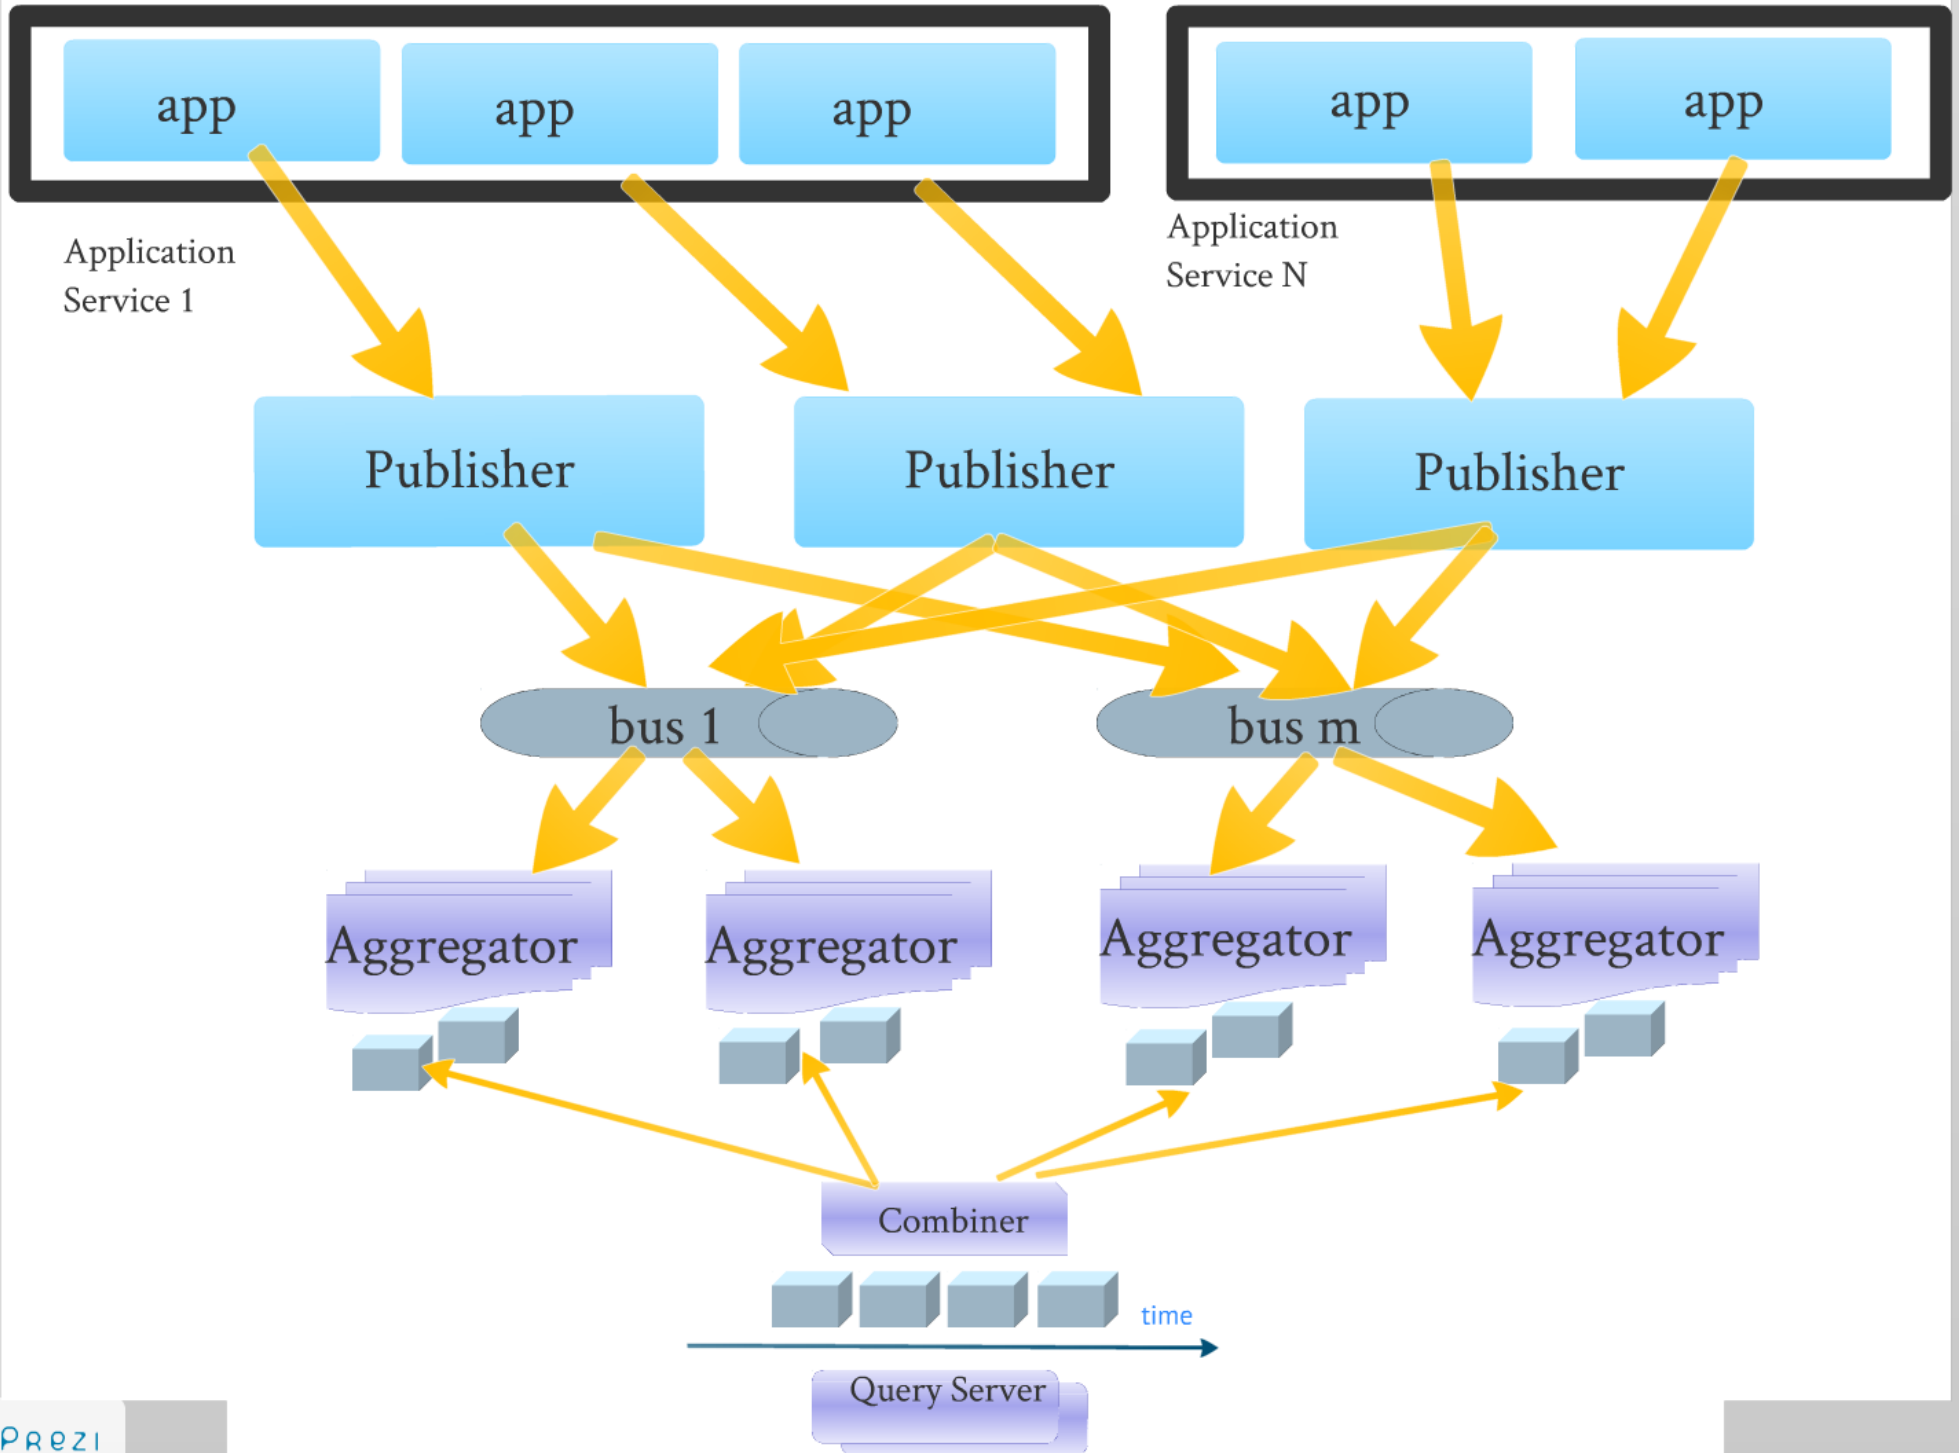
Problem 3: Metric Extraction

Semi-Structured Logs need to be processed for metric and event extraction

Logs -> Metrics + Events (In Real-time)



Distributed OLAP



Problem 4: Scalable Time Series Data Store

Millions Inserts Per Second
Data is multi-dimensional time series

Challenges:

- How scale writes and reads
- How to find time series given dimensional constraints
- How to down-sample metrics
- How to purge stale data

Time series storage consideration

Column-oriented models are good match for the needs of storing time series data. Consider:

However, what should be the row length?

- It can not grow forever
- Short rows defeat efficient writes
- Ideal row length has constant optimal resolution

OpenTSDB chose fixed row span - 1 hour

- 10 samples/takes every 10 sec, 10 - 10 every min
- Does not work well for temporally aggregated metrics (10 min, 1 hour, etc)
- Does not work for complex/quirky events

A better model is to have multiple row spans

- The lower the sampling frequency the longer the row span
- Use app's knowledge of what is the frequency of requested metric when querying for it across

Another improvement is to store complex values instead of just scalars

- Can pack more data points into a single cell to improve efficiency and reduce write TPS

Consider the need of a QA service to keep count of top-N operations

Non-obvious challenge: aggregation of complex values

Column	Row 1	Row 2	Row 3	Row 4
Value	100	200	300	400
Value	100	200	300	400
Value	100	200	300	400
Value	100	200	300	400

OpenTSDB model:
MPS: 10 = Metric ID, Time % 10, duration
Column name is often time 10 sec
Value is long float

Why not? It's a match
MPS: 10 = Metric ID, Metric Frequency, Time % 10, etc., 1, duration
Column name is often time 10 sec
Value is Complex/quirky

Complex? (Type is float, MaxValue, unit)
float, unit, etc., 1
float, Metric, context

Time series storage consideration

Column-oriented model is a good match for the needs of storing time series: Hbase, Cassandra

However, what should be the row length?

- It can not grow forever
- Short rows defeat efficiency wins
- Ideal row length has constant optimum

#columns

OpenTSDB chose fixed row span – 1 hour

- 360 sample if taken every 10 sec, 60 – if every min
- Does not work well for temporally-aggregated metrics 10 min, 1 hour, etc.
- Does not work for irregular/sparse events

A better model is to have multiple row spans


- The lower the sampling frequency the longer the row span
- Use apriory knowledge of what is the frequency of requested metric when querying for its rows

Another improvement is to store complex values instead of just scalars

- Can pack more data points into a single cell to improve efficiency and reduce write TPS

Consider the need of a SOA service to keep count of top N operations

New interesting challenge: aggregation of complex



OpenTSDB model:

ROW_ID ::= MetricID, Time/1h, dimensions

Column_name ::= offset from 1h time

Value ::= long | float

eBay AppWatch model:

ROW_ID ::= MetricID, MetricFrequency, Time/{1h,1d,1mo,...}, dimensions

Column_name ::= offset from 1h time

Value ::= ComplexType value

ComplexType ::= (count, Map<name, count> |
(val1, val2, val3..) |
(count, Map<name, prop>)

BIG INSIGHTS

Data at eBay is BIG and getting BIGGER
Data is becoming more and more valuable
Big data needs Big Architecture

- Remove Entropy as early as possible
- Data Distribution needs to be elastic
- Storage Optimized for Access Patterns

ebay™ is **HIRING!!**
bavalani@ebay.com

