Big-Data Innovation:
Hadoop, Real-time, and Machine Learning

Andy Feng
Yahoo

Hadoop at Yahoo

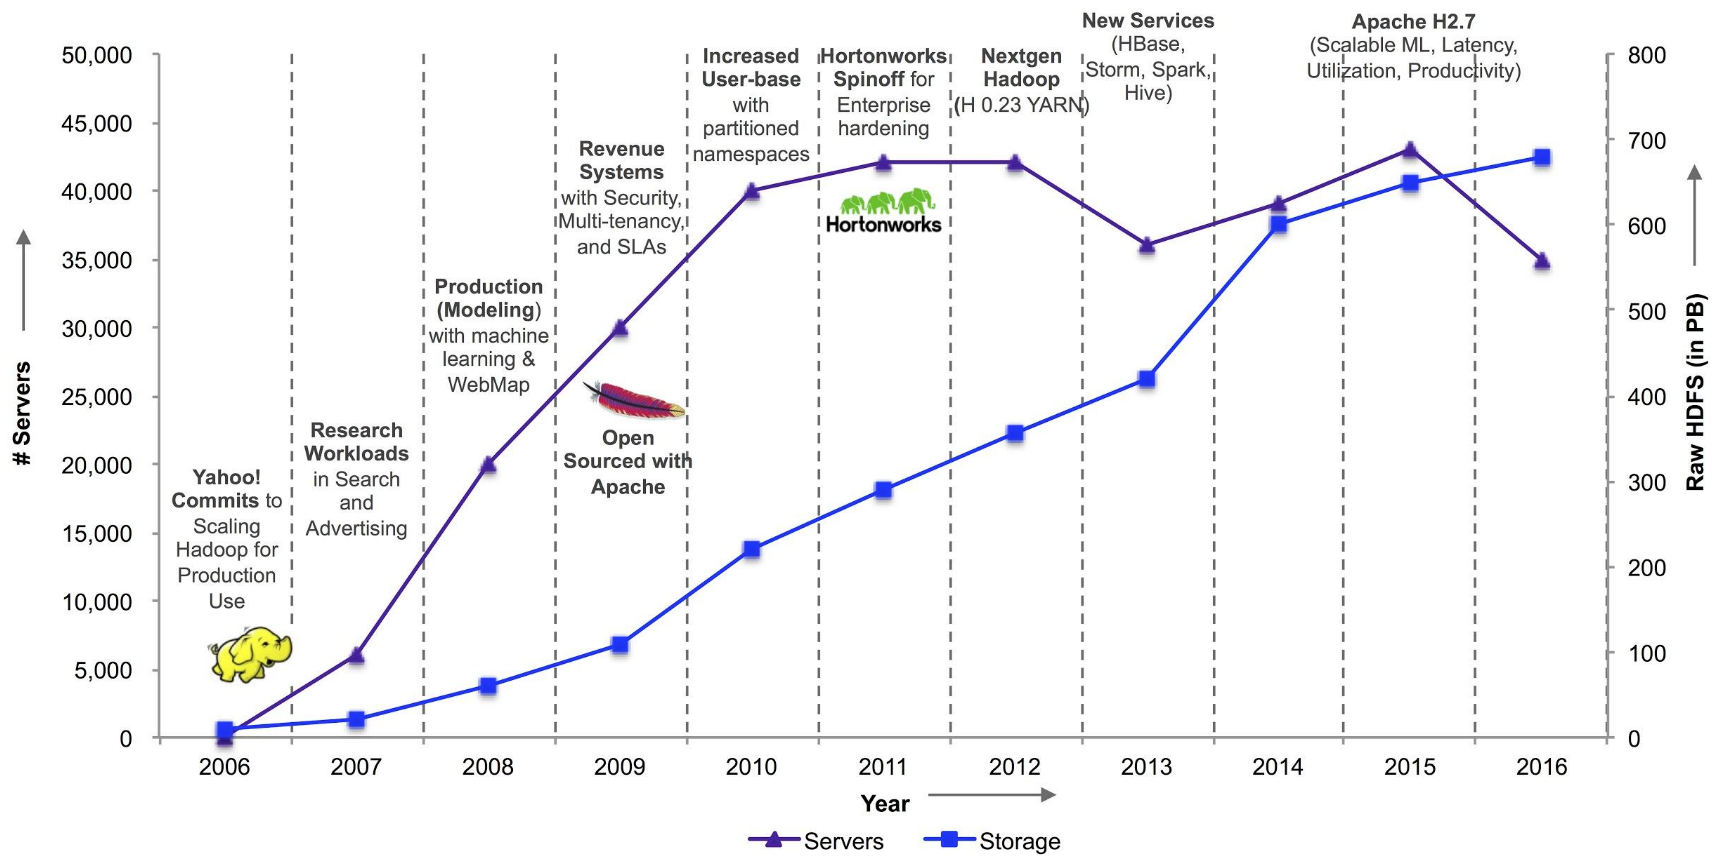

YAHOO!

Real-time Processing at Yahoo

Big-Data Applications

Weather App

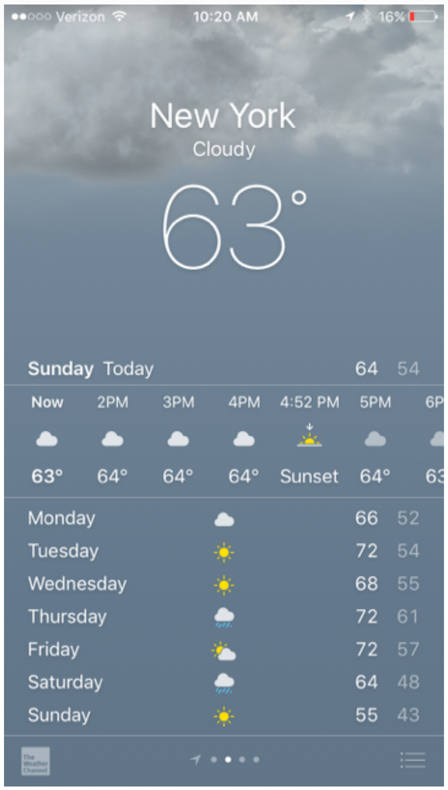

Weather App

- Beauty
  - Computational assessed
- Relevant
  - Location
  - Time
  - Cloudy
  - Shower
  - …

Yahoo Weather App

YAHOO!

Flickr flickr.com/cameraroll

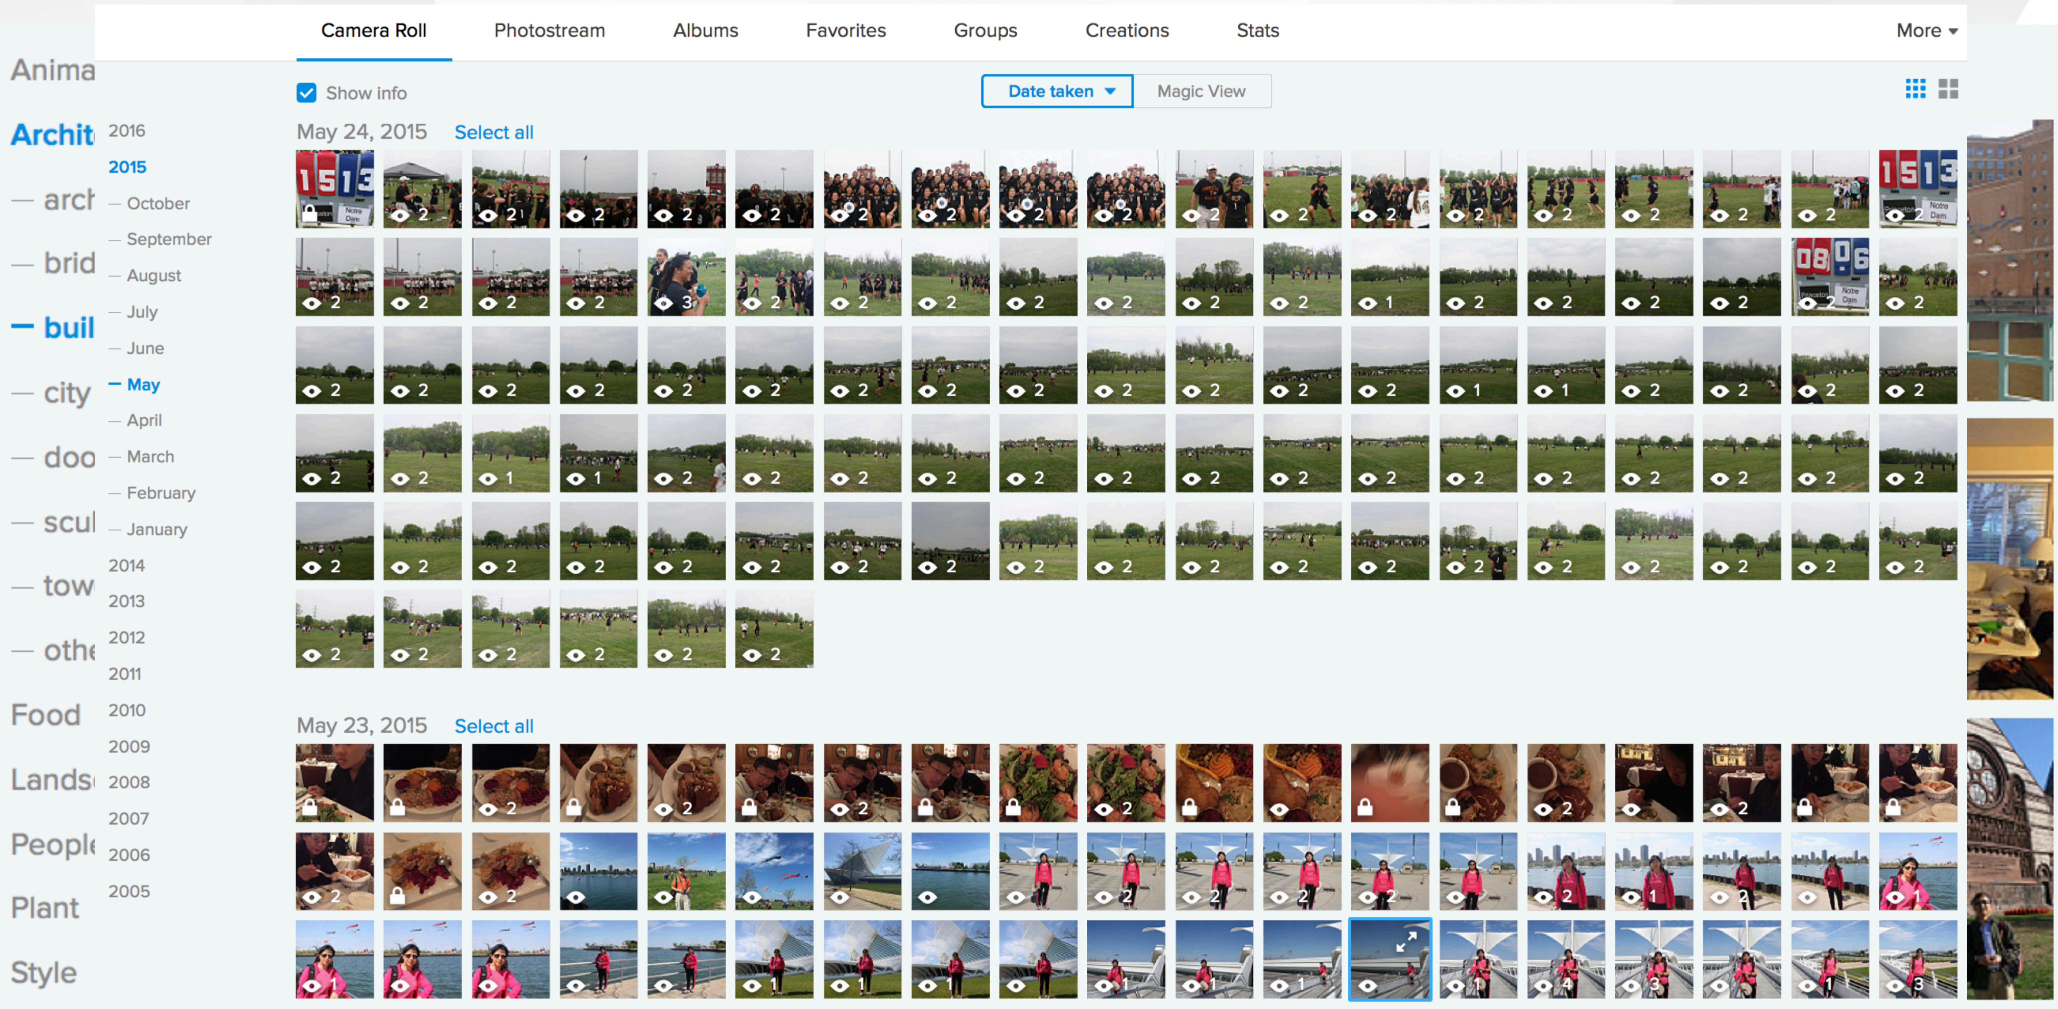

Yahoo Vision Kit

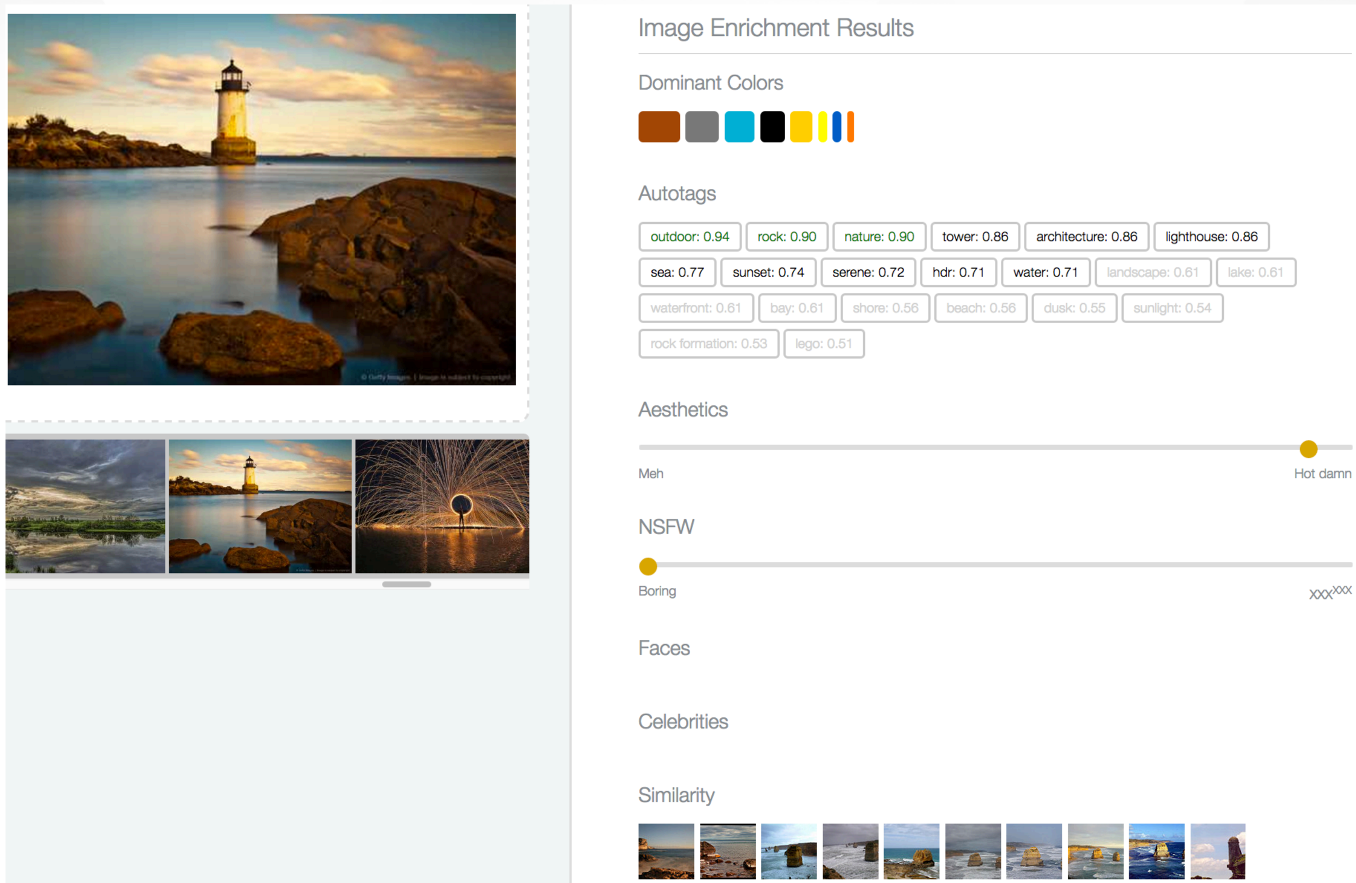

Personalized Content

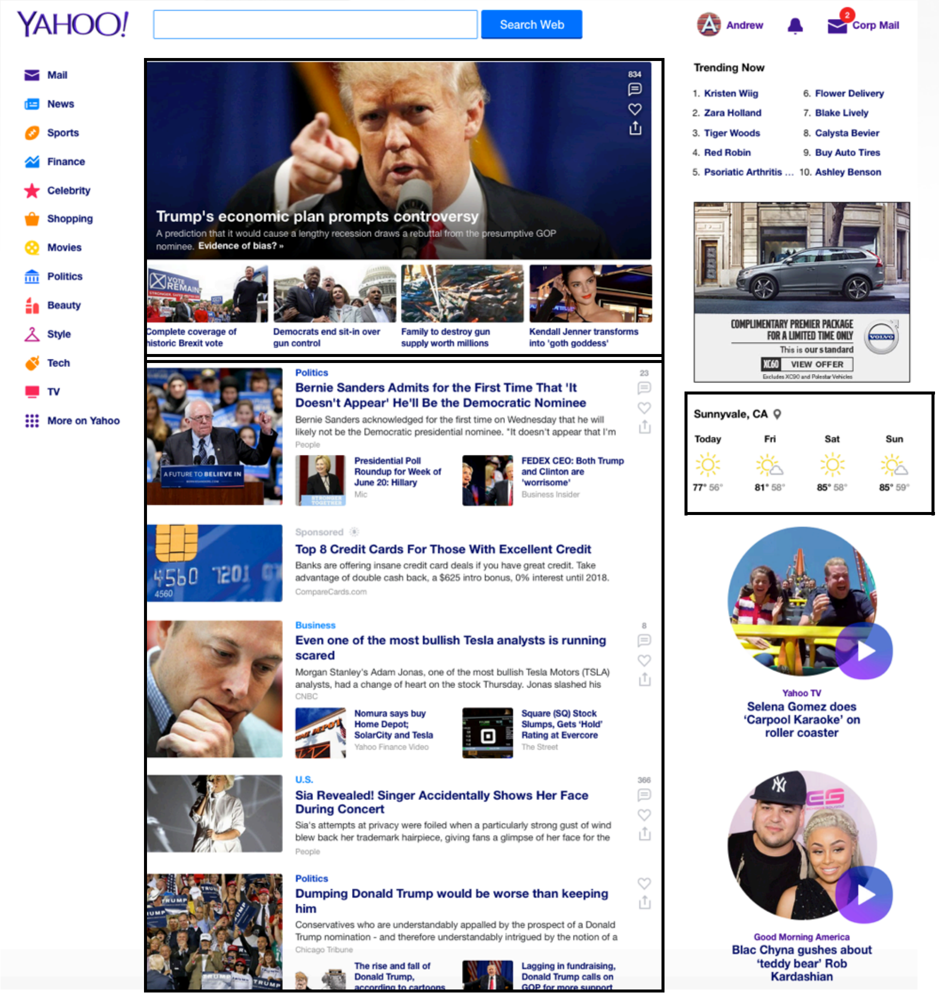

Content augmentation

User profiling

Recommendation

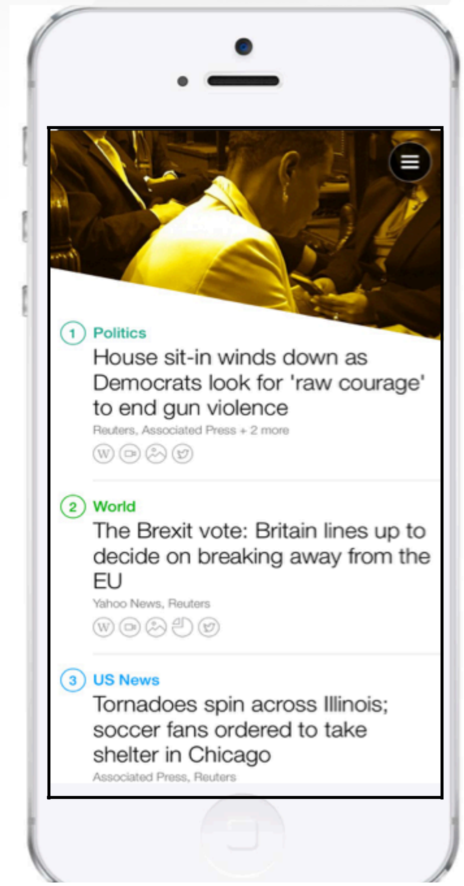

Mail Smart Views

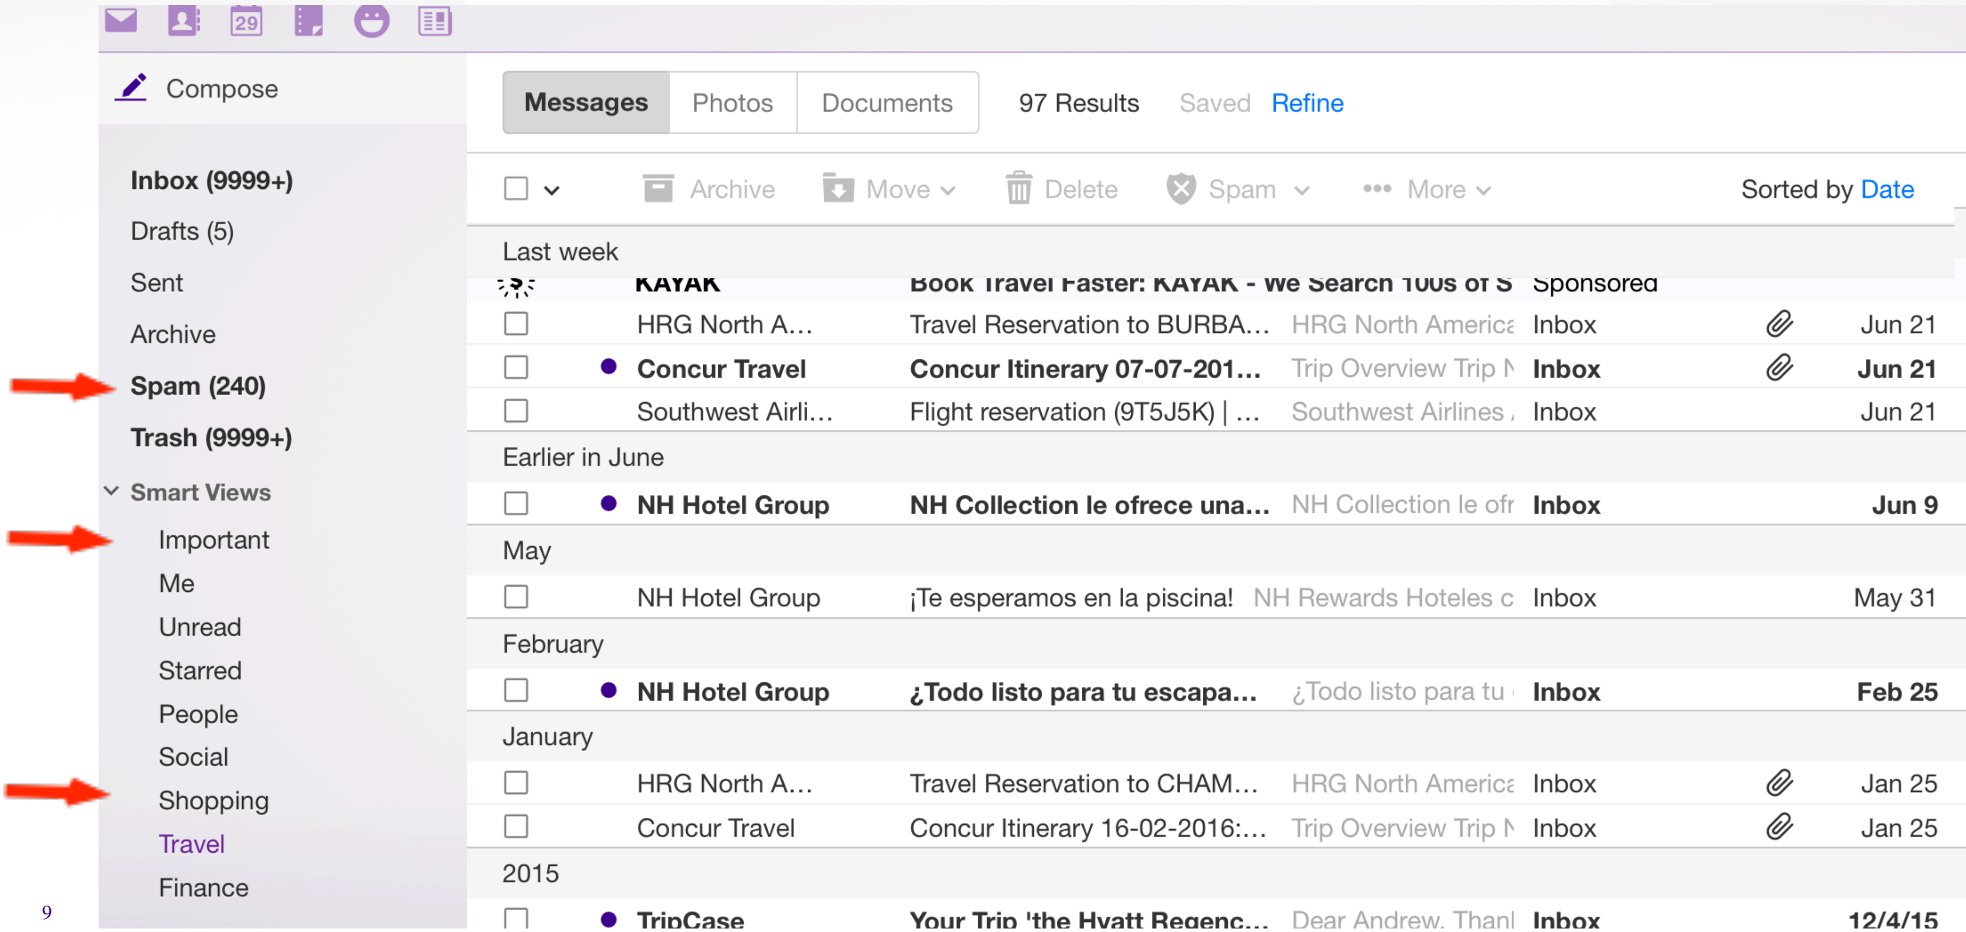

Search

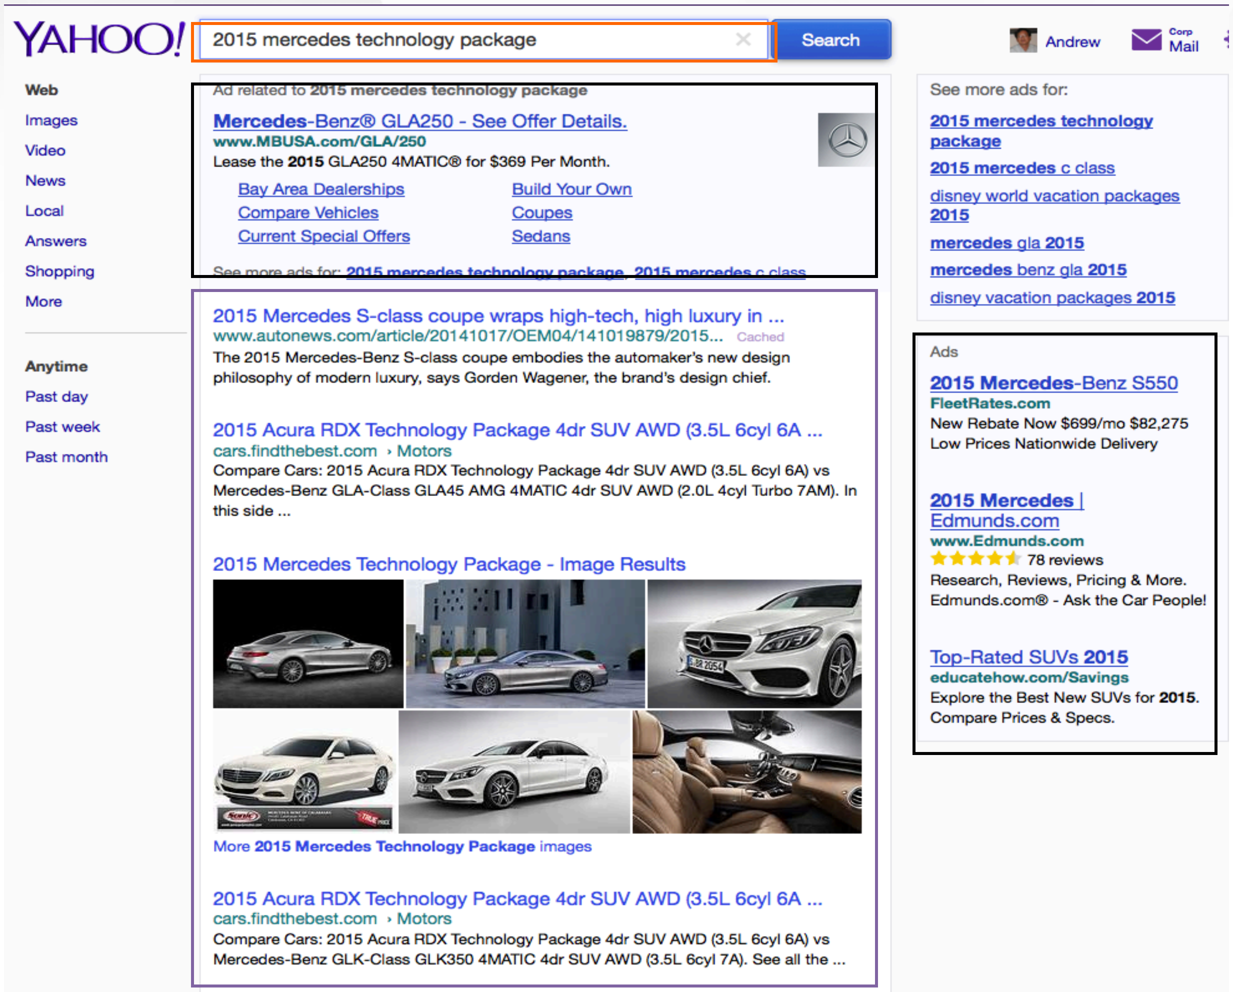

Query intention

Page ranking

Ads matching

Yahoo Search2Vec: Big Data & Machine Learning http://bit.ly/28SLdjU

Bucket Tests	Query Coverag e	Auction Depth	Revenue per Search
Simple model vs. Baseline	+1.14%	+2.13%	+7.07%
Large model vs. Baseline	+3.61%	+4.57%	+17.12%

Big-Data Innovations

Technology Stack

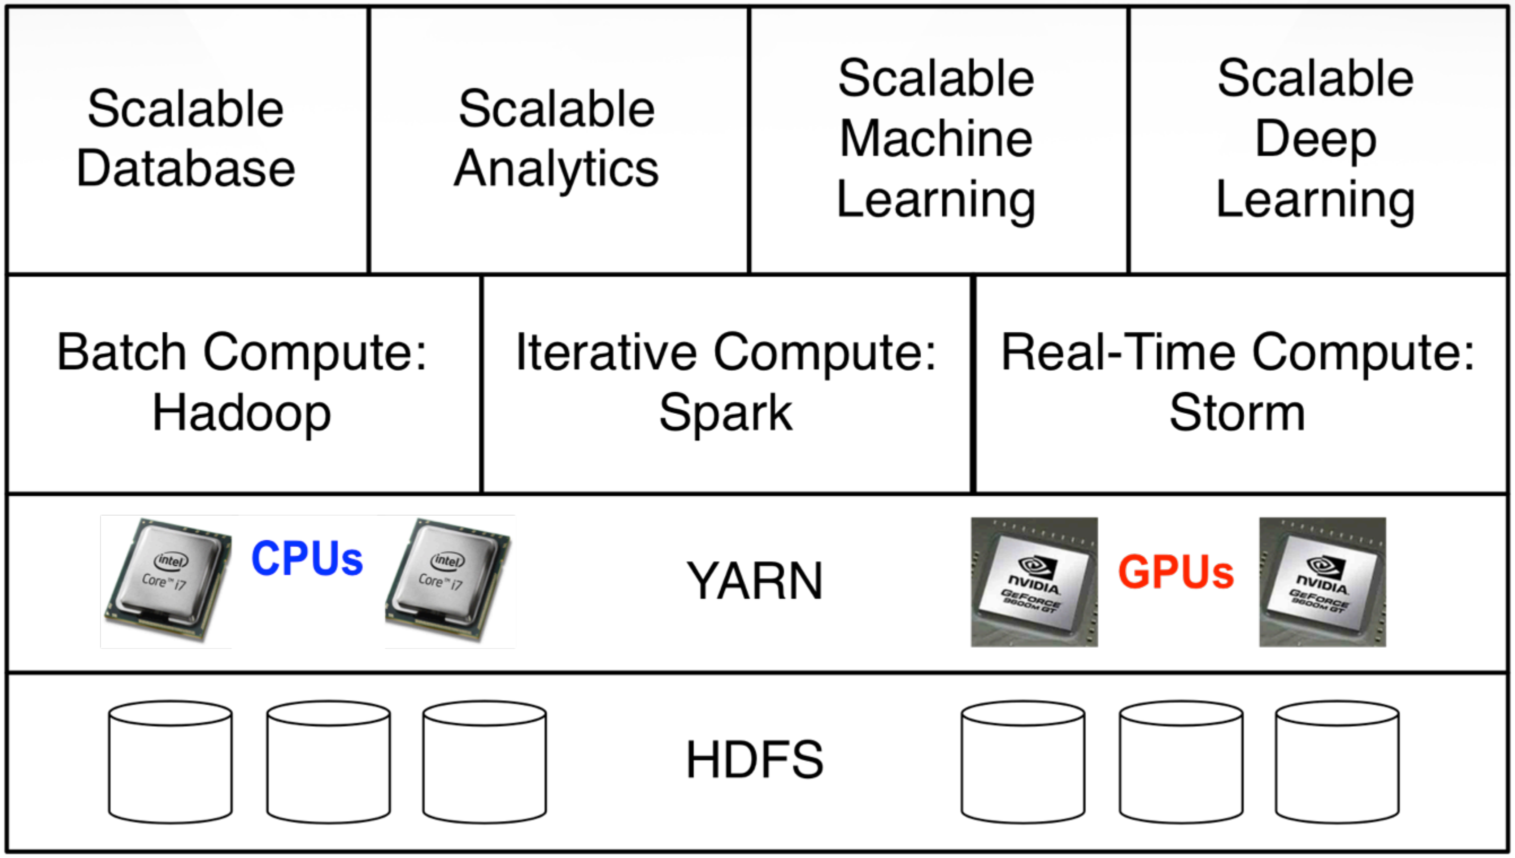

Flickr Pipeline: 10B photos, 8M per day

(1) Feature Engineering @ Scale → (2) Deep Learning @ Scale → (3) Non-deep Learning @ Scale → (4) Apply Model @ Scale → (5) Product Integration @ Scale

* http://bit.ly/1KIDfof by Pierre Garrigues, Deep Learning Summit 2015

System Configuration: Homogenous Clusters

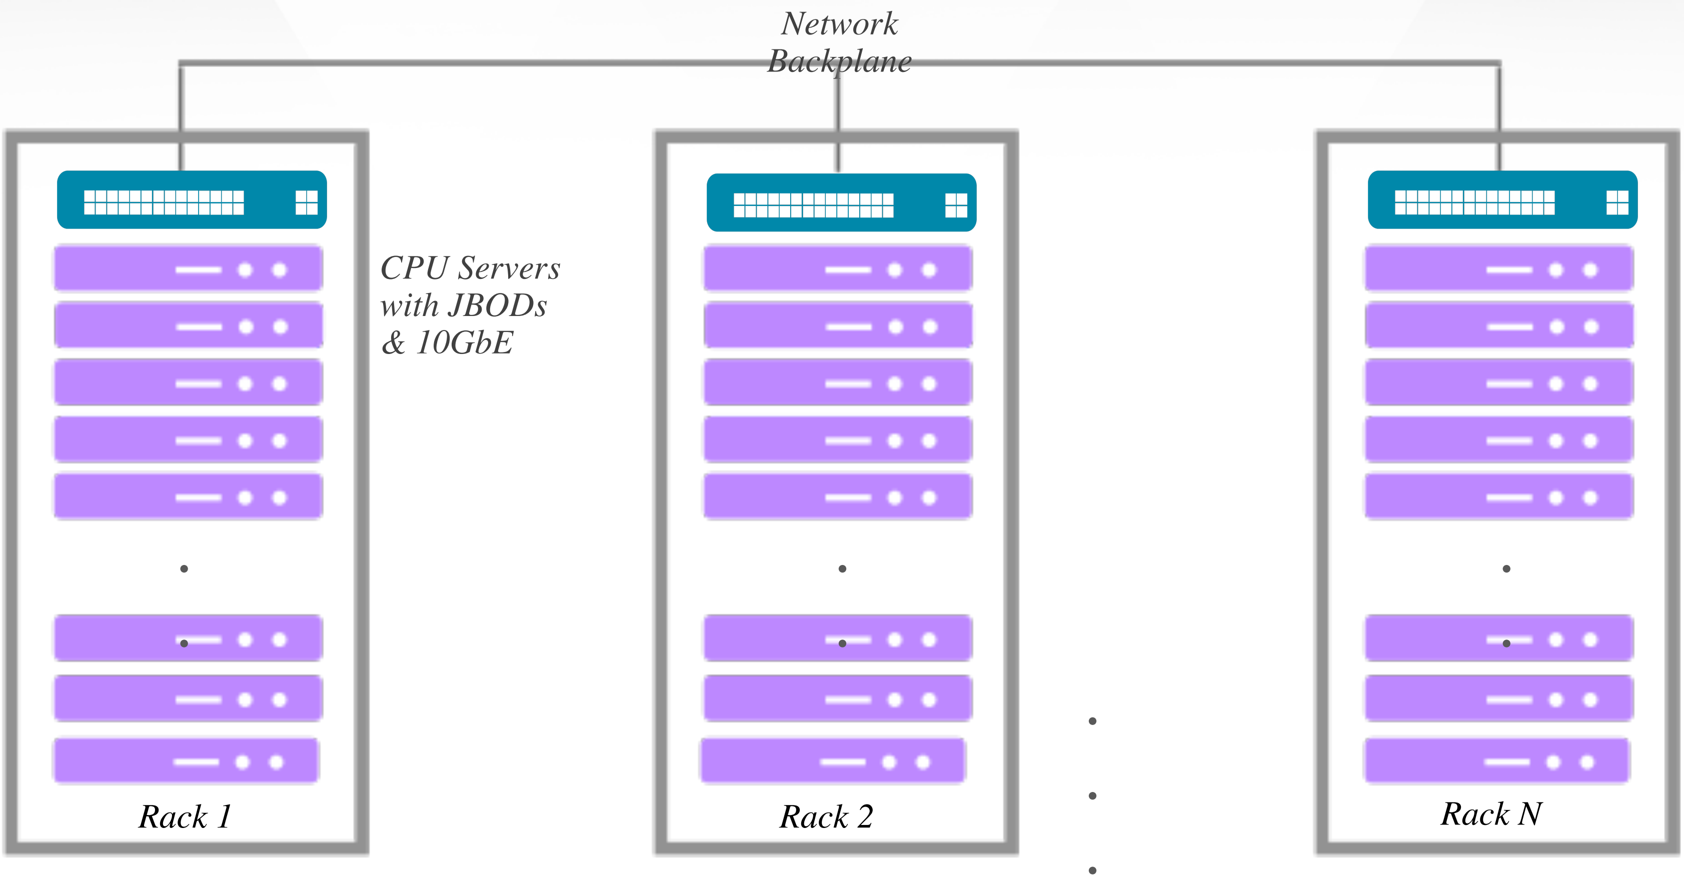

YAHOO!

System Configuration: Heterogeneous Clusters
http://bit.ly/295RNbU

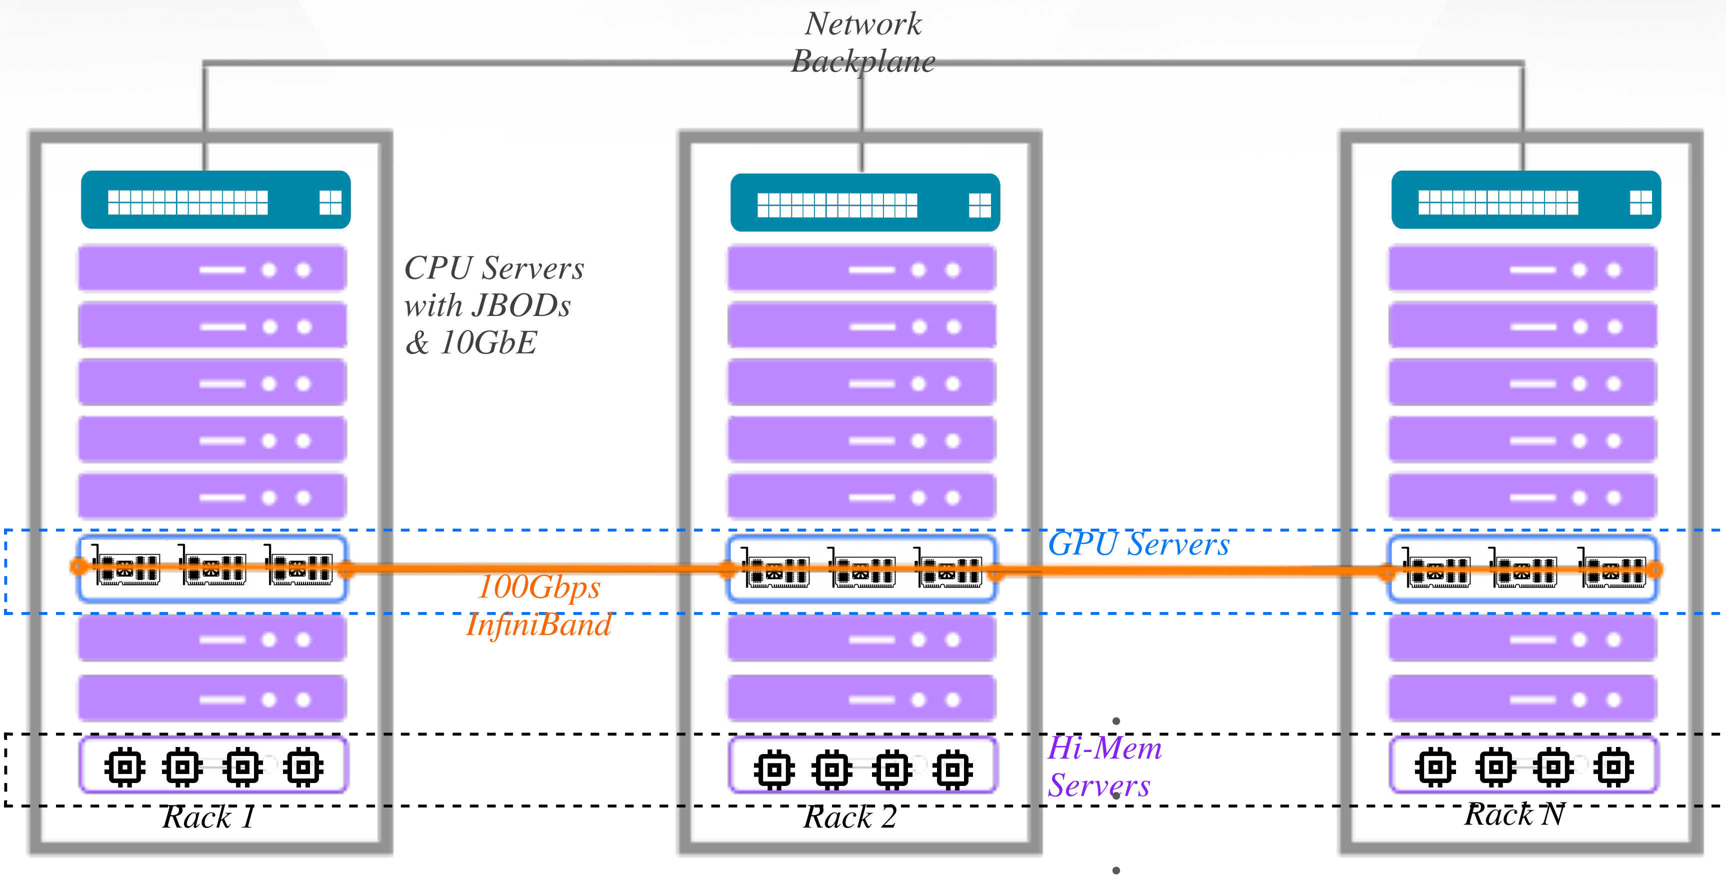

YAHOO!

Resource Allocation: Early Days

Compute Total and Used (TB)

One Month Sample (2015)

Cluster 1 (2,000 servers)
HDFS 12 PB
Memory 23 TB
Avg. Util: 26%

Cluster 2 (3,100 servers)
HDFS 21 PB
Memory 52 TB
Avg. Util: 40%

Cluster 3 (5,400 servers)
HDFS 36 PB
Memory 70 TB
Avg. Util: 59%

Resource Allocation: 40% decrease in TCO http://bit.ly/295RNbU

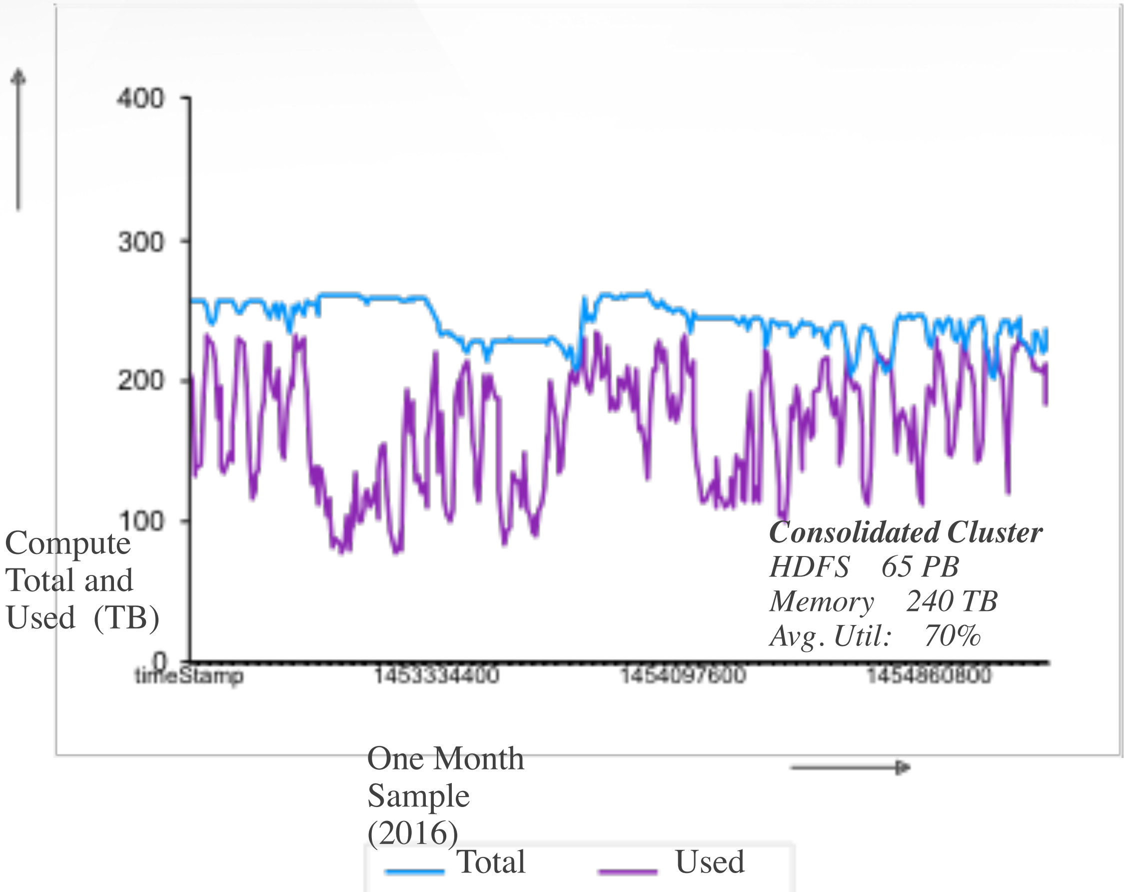

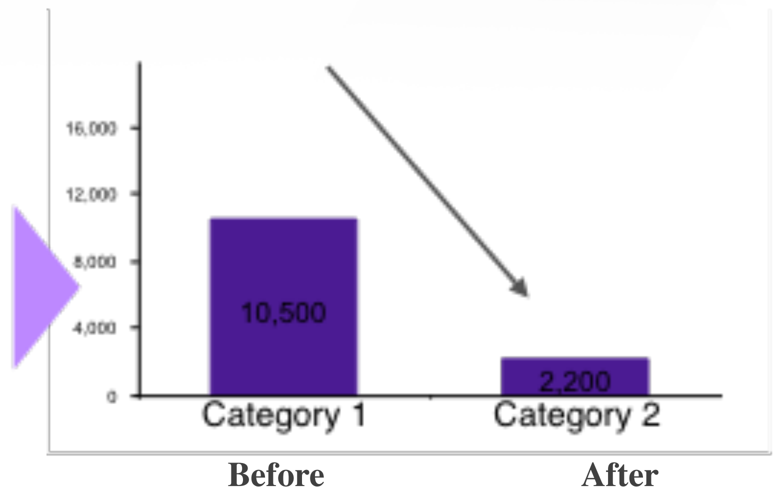

65% increase in compute capacity

50% increase in avg. utilization

Apache Storm: Real-time Processing
https://storm.apache.org

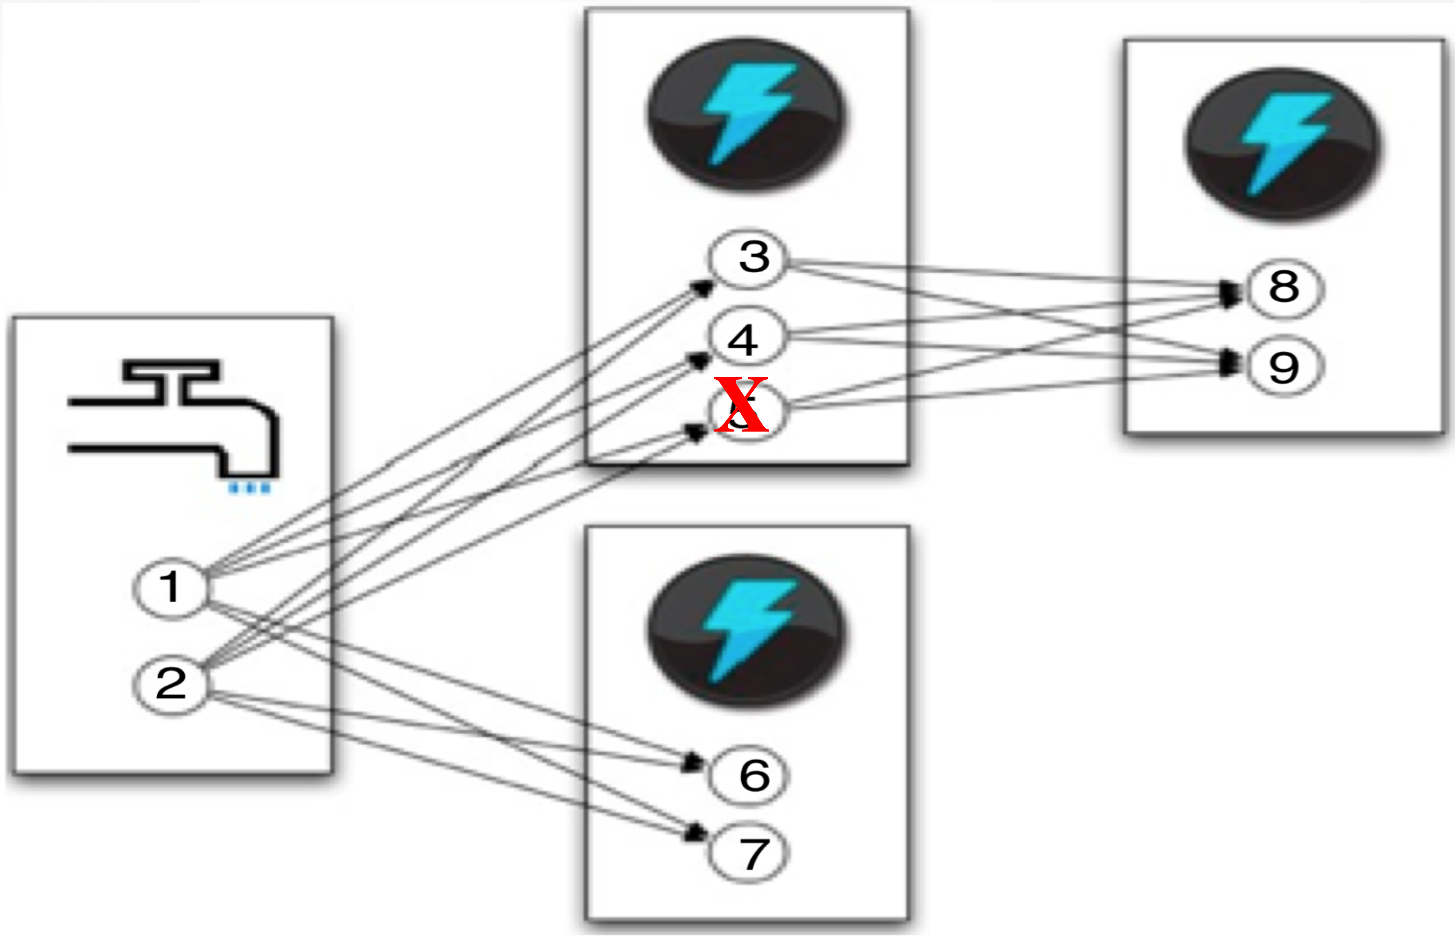

Resource Aware Scheduling: 100% increase utilization

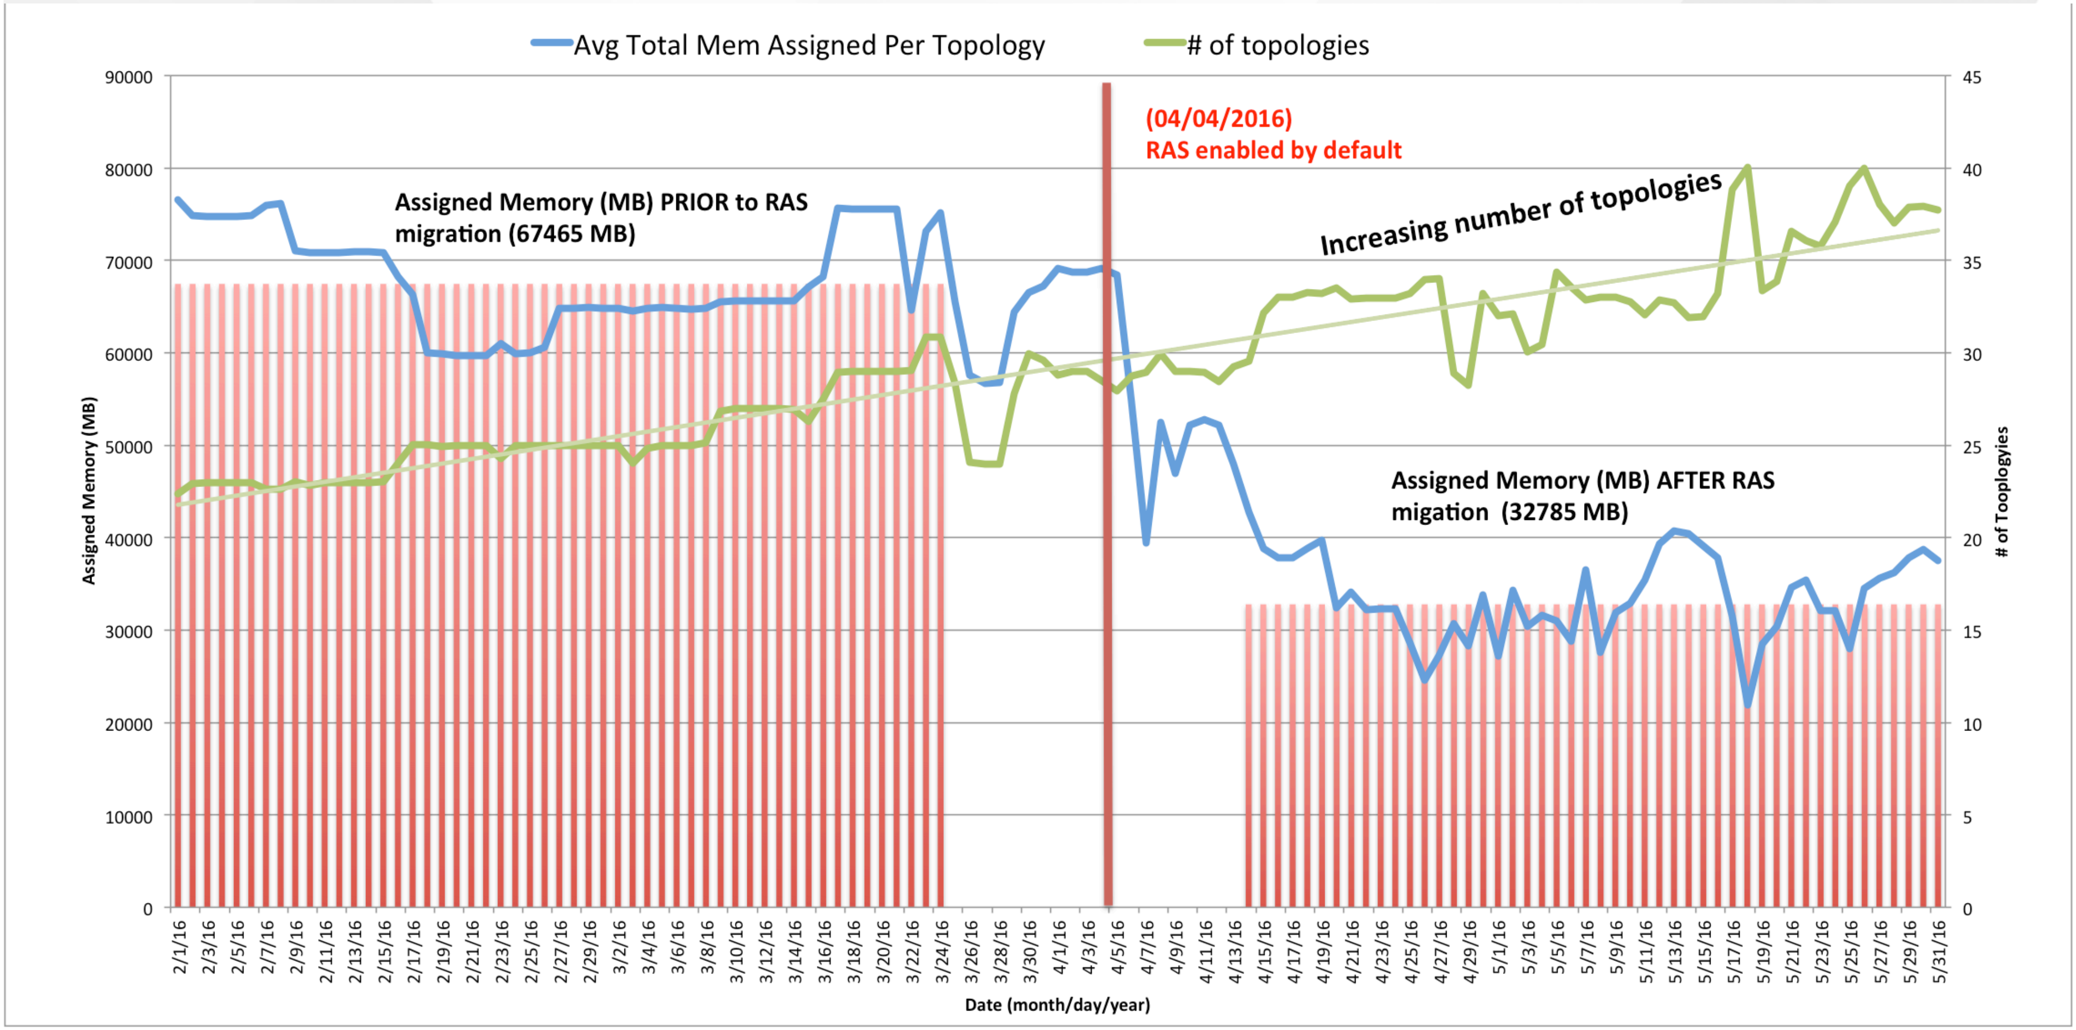

Apache HBase: NoSQL Database for Hadoop
https://hbase.apache.org

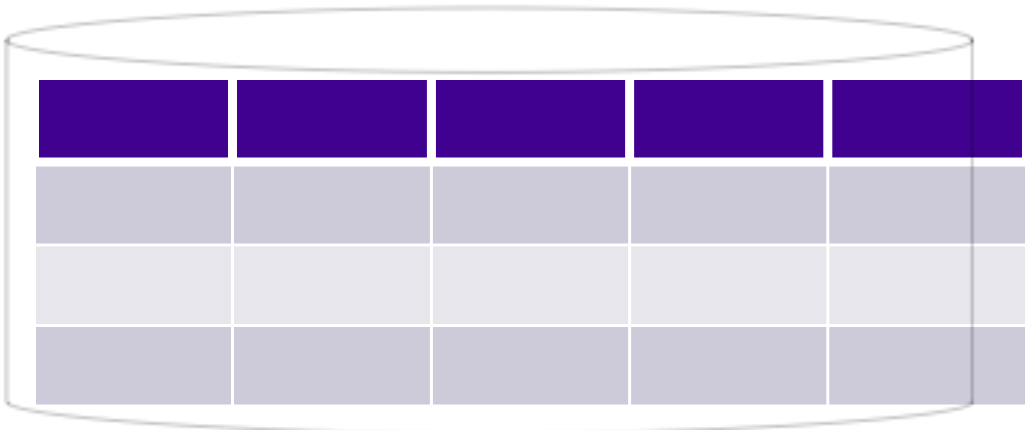

- Several million regions (partitions of tables)
- 800 region servers

YAHOO!

Support Millions of Regions
http://bit.ly/29hlLFW

3 Million Regions

	Scan Meta	**Assignment**
1 Meta / 1 RS	56min	19.79min
32 Meta / 3 RS	2.91min	12.56min

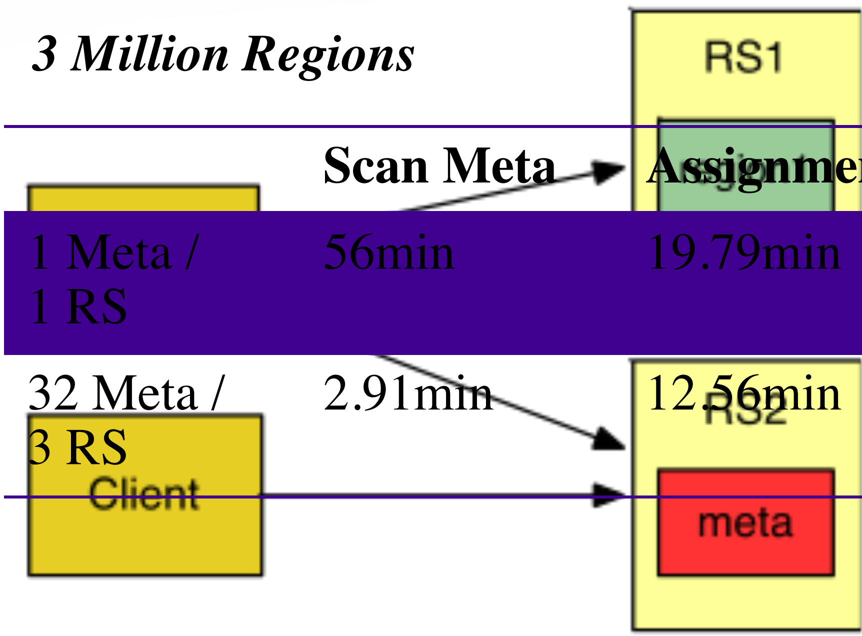

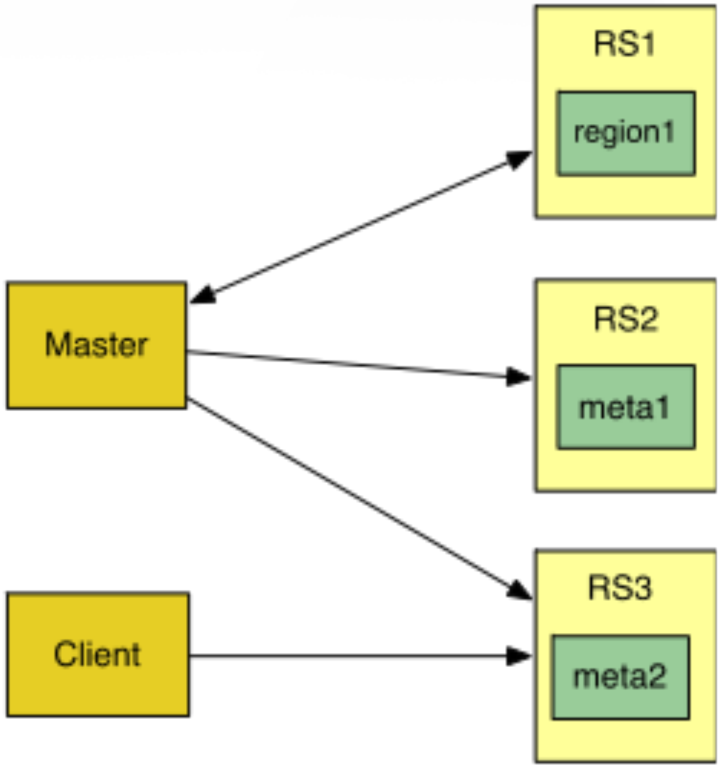

HBase Multi-tenancy & Isolation
http://bit.ly/28YK4a7

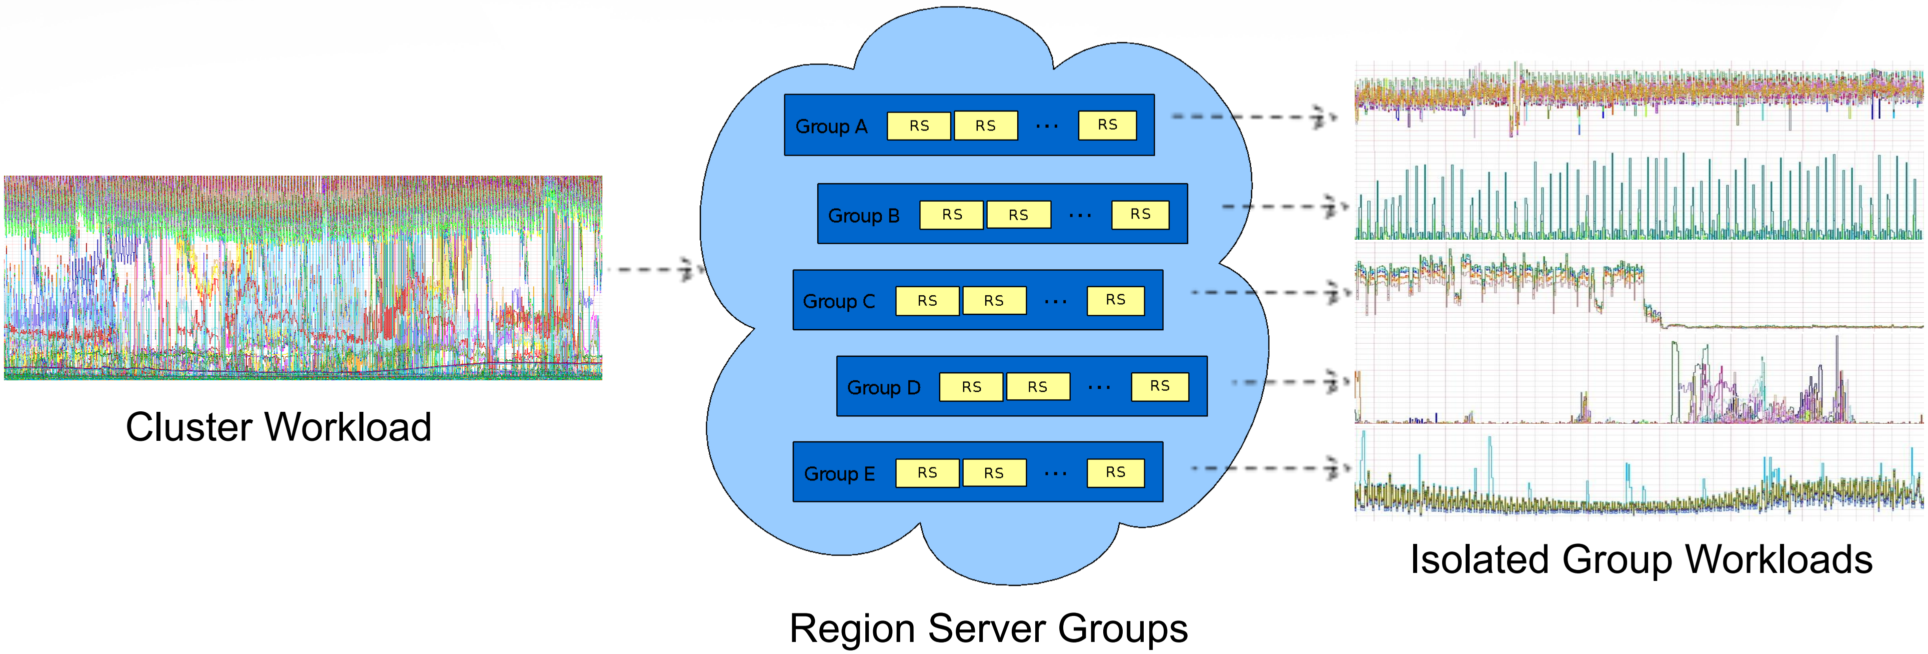

YAHOO!

Apache Omid: Transactions for NoSQL DB

http://omid.incubator.apache.org

ACID Transactions

- *Multi-row/multi-table transactions*
- *Snapshot isolation*
- *Lock-free*

Omid: Transactional HBase App

```
TransactionManager tm = HBaseTransactionManager.newBuilder().build();
Transaction tx = tm.begin();
Put row1 = new Put(Bytes.toBytes("EXAMPLE_ROW1"));
row1.add(family, qualifier, Bytes.toBytes("VALUE_1"));
tt.put(tx, row1);
Put row2 = new Put(Bytes.toBytes("EXAMPLE_ROW2"));
row2.add(family, qualifier, Bytes.toBytes("VALUE_2"));
tt.put(tx, row2);
tm.commit(tx);
```

Omid: Snapshot Isolation
http://bit.ly/28YZvQF

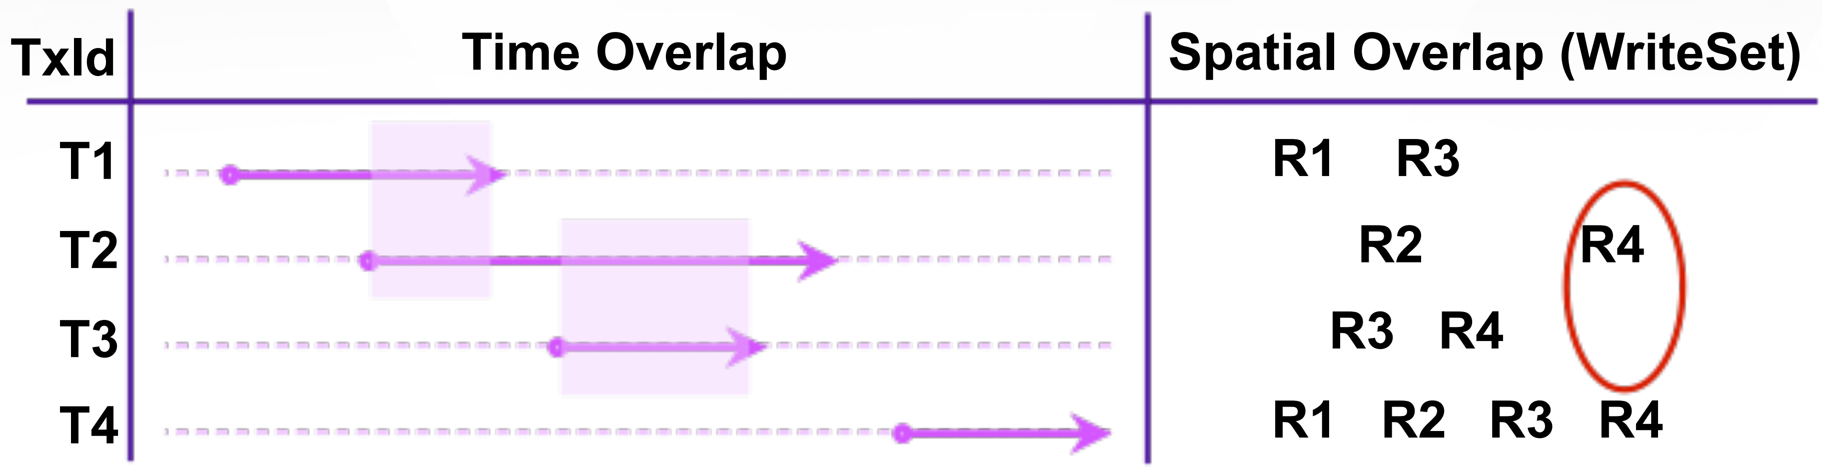

- T2 overlaps in time with T1 & T3, but
  - Spatially T1 ∩ T2 = ∅ ➔ T1 and T2 can both commit.
  - Spatially T2 ∩ T3 = { R4 } ➔ T2 or T3 will abort.

Machine Learning: Search2Vec
http://bit.ly/28SLdjU

Input: Training Data

S1: gas_caps gas_cap_replacement_for_car slc_679f037d gas_door_replacement_for_car slc_466145af1 fuel_door_covers adid_28540536 slc_348709d7 autozone_auto_parts adid_33183157 auoto_zone slc_8dcdab5d slc_58f979b6

S2: hoka_running shoe_reviews adid_22830711 hoka_shoes_for_bad_feet hoka_shoes_amazon slc_231gfaw zappos_shoes slc_7c126f71 hoka_walking_shoes

S3: king_tut king_tut_exhibit king_tut_exhibit_seattle_2015 slc_726y6j51 charlies_seattle adid_55774014

Output:
Numeric Vectors
Of queries & ads

Vector("san jose weather") ≈
Vector("weather 95113") ≈
Vector(ad123)

YAHOO!

Search2Vec: Cosine similarity b/w ads and queries

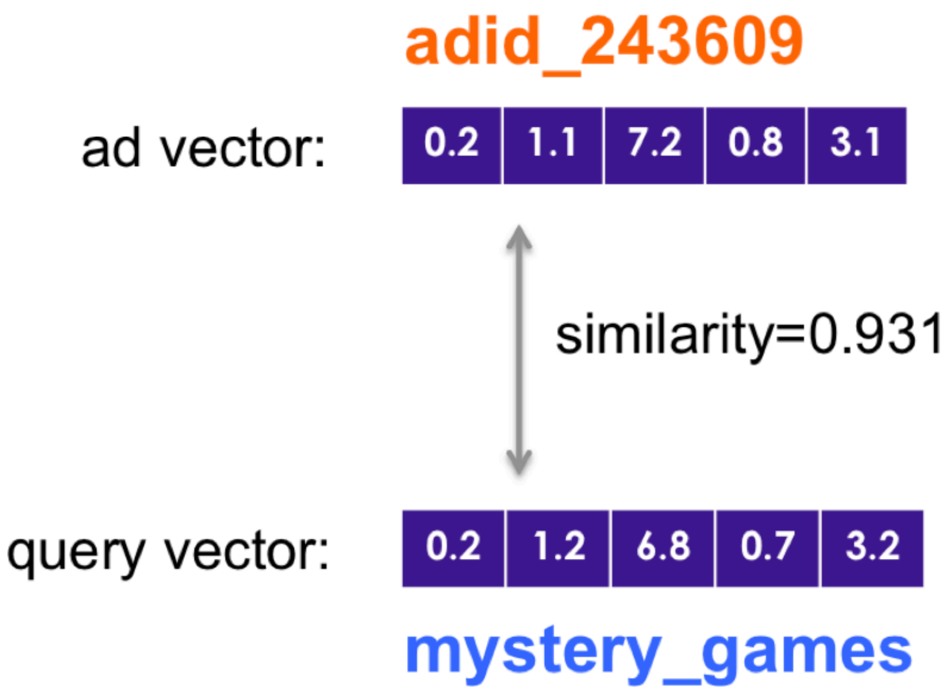

ad metadata	**ad id**: 243609 **bidterm id**: 341454 **ad title**: Host a Fun Murder Mystery Party **ad description**: Huge selection of fun murder mystery games for all ages, groups. #1 site for instant downloads and boxed sets of exciting murder mystery party games
generated keywords	murder mystery murder mystery games mystery games murder mystery party mystery party free murder mystery games for parties detective games murder mystery game how to host a mystery party for kids murder mystery dinner friends game night murder mystery parties at home murder mystery dinner party ...

YAHOO!

Search2vec: 300M Unique Queries/Ads/Links

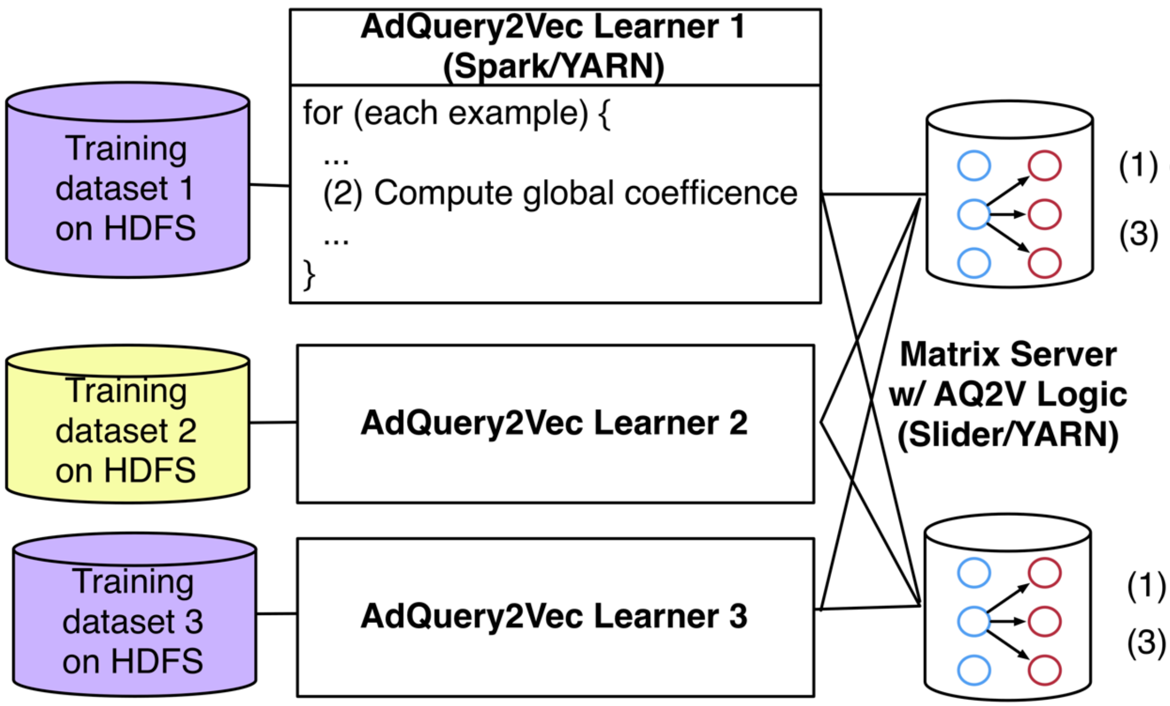

- Computation on servers
  - (1) Negative sampling
  - (1) Compute gradient: X*Y
  - (3) Adjust vectors: Y=aX+Y
- Daily training enabled
  - weeks ➔ hours

YAHOO!

Deep Learning: Previous Practice

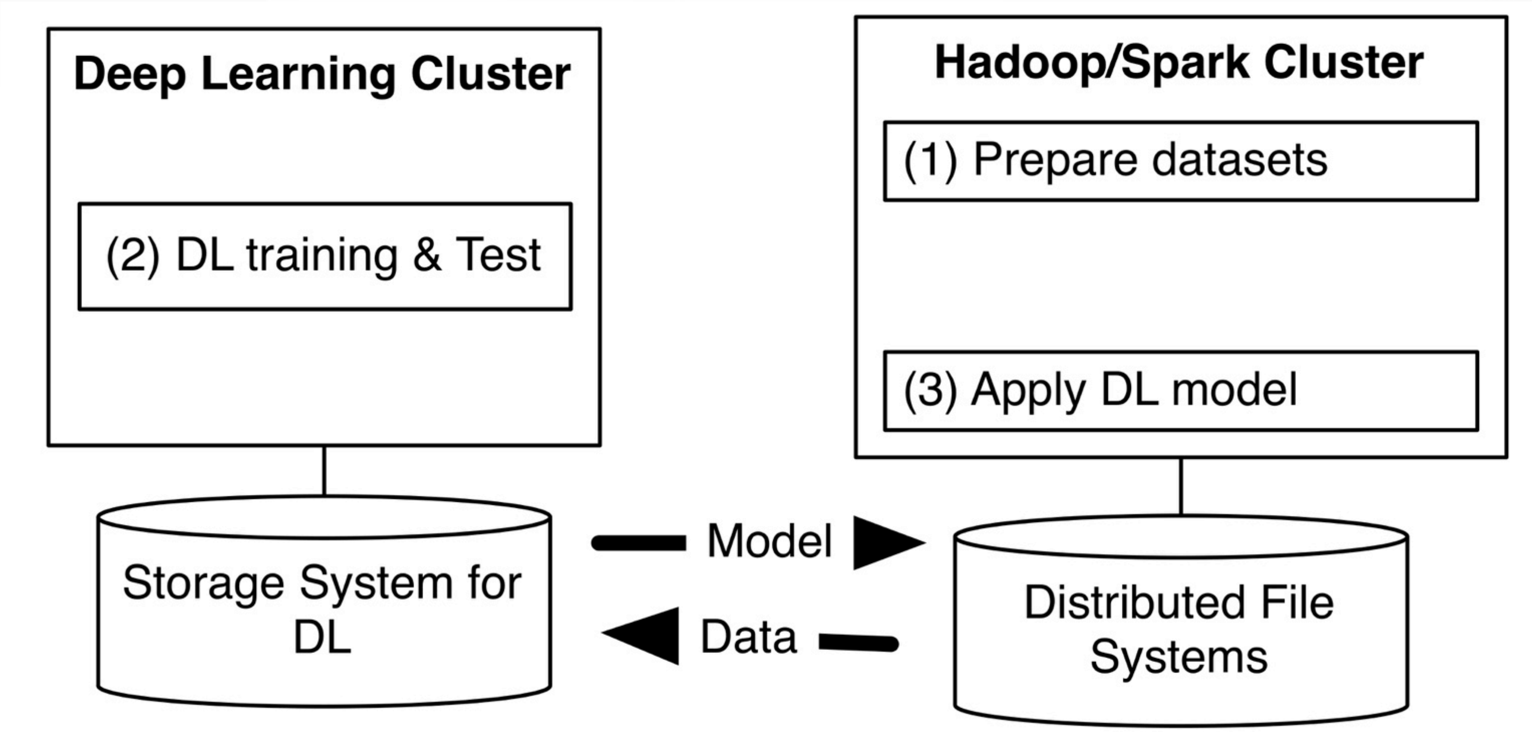

Scalable Deep Learning for Big Data

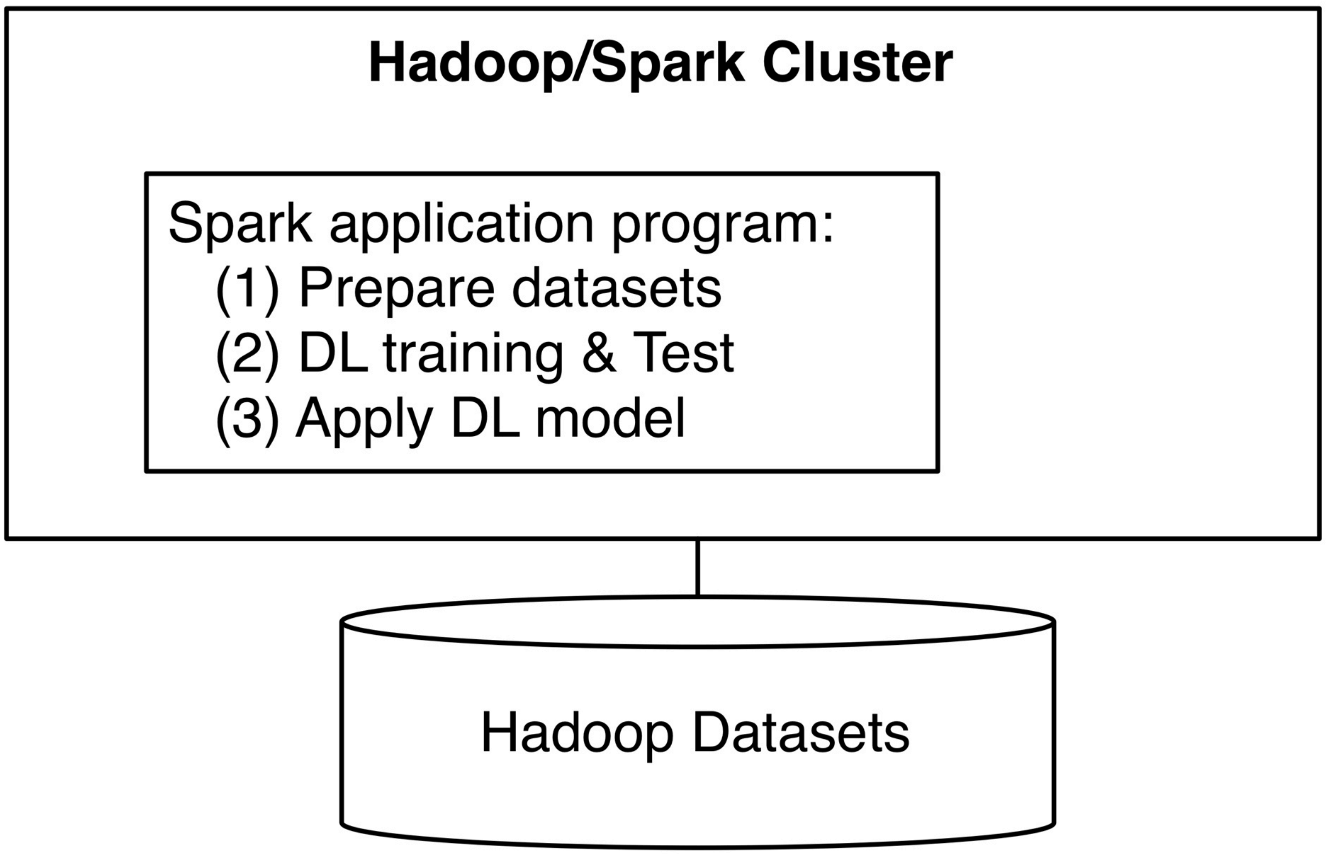

CaffeOnSpark Open Sourced

github.com/yahoo/CaffeOnSpark

- Released in Feb. 2016
  - Apache 2.0 license
- Distributed deep learning
  - GPU or CPU
  - Ethernet or InfiniBand
- Easily deployed on public cloud or private cloud

CaffeOnSpark: One Notebook http://bit.ly/1REZ0cN

File Edit View Insert Cell Kernel Help | Python 2

Logistic Regression using MLlib

```
In [57]: from pyspark.mllib.linalg import Vectors
         from pyspark.mllib.regression import LabeledPoint
         from pyspark.mllib.classification import LogisticRegressionWithLBFGS

In [58]: data = f.map(lambda row: LabeledPoint(row.label[0], Vectors.dense(row.ip1)))

In [59]: lr = LogisticRegressionWithLBFGS.train(data, numClasses=10, iterations=10)

In [60]: predictions = lr.predict(data.map(lambda pt : pt.features))

In [61]: predictions.take(5)

Out[61]: [7, 2, 1, 0, 4]
```

```
|00000001|    [1.0]|[1.3683326, -0.0,...|[2.0906663, 1.048...|[2.0]|
|00000002|    [1.0]|[1.5641443, -0.0,...|[-0.773368, 10.61...|[1.0]|
|00000003|    [1.0]|[-0.0, 1.9505613,...|[16.46351, -6.917...|[0.0]|
|00000004|    [1.0]|[0.5979191, 0.075...|[-0.48371825, -2....|[4.0]|
+--------+--------+--------------------+--------------------+-----+
```

```
In [45]: dl_train_source = DataSource(sc).getSource(cfg,True)

In [46]: cos.train(dl_train_source)
```

Datasets for YOU

Data, Data, Data.

- Machine learning is all about data
  - More data → Better model
- Industrial researchers work on large-scale data
  - Academic researchers need data, too ☹

YAHOO!

Yahoo Webscope Datasets

Categories	datatype	# of datasets
L: Language Data	l	25
G: Graph and Social Data	g	8
R: Ratings and Classification Data	r	10
A: Advertising and Market Data	a	4
C: Competition Data	c	3
S: Computing Systems Data	s	4
I: Image Data	i	5

- *Ex. https://webscope.sandbox.yahoo.com/catalog.php?datatype=r*
- Non-commercial use by academics and scientists: 8,000+

YAHOO!

L5 - Query Vectors

- **8M** query vectors trained using our search2vec system
- May serve as a testbed for query rewriting task
  - IR research
  - Word and sentence similarity task in NLP research

Method	oAUC	Macro NDCG@5
word2vec	0.817	0.929
search2vec	**0.880**	**0.959**

I3 - Flickr Creative Commons

- 100M Flickr photos & videos
  - › Largest public multimedia dataset ever released
    - ImageNet … 14M images
- Metadata
  - › Tags, title, des & geo
    - 68M photos have at least one tag
  - › Comments, favorites, social network data queried via Flickr API

longing by Robert-Crouse Baker

Tags: Ocean Beach, San Francisco, Pacific, Surf, surfing, waves, green, water, man, surfer, January, océan, le surf, vagues, mer, sea
Description: “My soul is full of longing for the secret of the sea,
and the heart of the great ocean sends a thrilling pulse through me.” — Henry Wadsworth

Summary

- Big-data technologies and continued innovation are critical for Yahoo business.
- Yahoo continues open source contribution ever since our donation of Hadoop to Apache
  - Apache Spark/Storm/Hbase/Omid, CaffeOnSpark
  - Webscope datasets

Thanks!

yahoohadoop.tumblr.com

bigdata@yahoo-inc.com