Data Applications: Past, Present & Future

Milind Bhandarkar
Chief Scientist, Pivotal
@techmilind

About Me

- http://www.linkedin.com/in/milindb
- Founding member of Hadoop team at Yahoo! [2005-2010]
- Contributor to Apache Hadoop since v0.1
- Built and led Grid Solutions Team at Yahoo! [2007-2010]
- Parallel Programming Paradigms [1989-today] (PhD cs.illinois.edu)
- Center for Development of Advanced Computing (C-DAC), National Center for Supercomputing Applications (NCSA), Center for Simulation of Advanced Rockets, Siebel Systems (acquired by Oracle), Pathscale Inc. (acquired by QLogic), Yahoo!, LinkedIn, and Pivotal (formerly Greenplum)

Pivotal

Kryptonite: First Hadoop Cluster At Yahoo!

M45

HP Labs

Intel Research

Yahoo!

University of Illinois at Urbana Champaign

Karlsruhe Institute of Technology, Germany

Infocomm Development Authority, Singapore

Electronics and Telecommunications Research Institute

Malaysian Institute for Microelectronic Systems

Russian Academy of Sciences

Carnegie Mellon University

ChinaMobile

CESGA

ChinaTelecom

GaTech

Computer Network information Center
Chines Academy of Sciences

OpenCirrus

Analytics Workbench

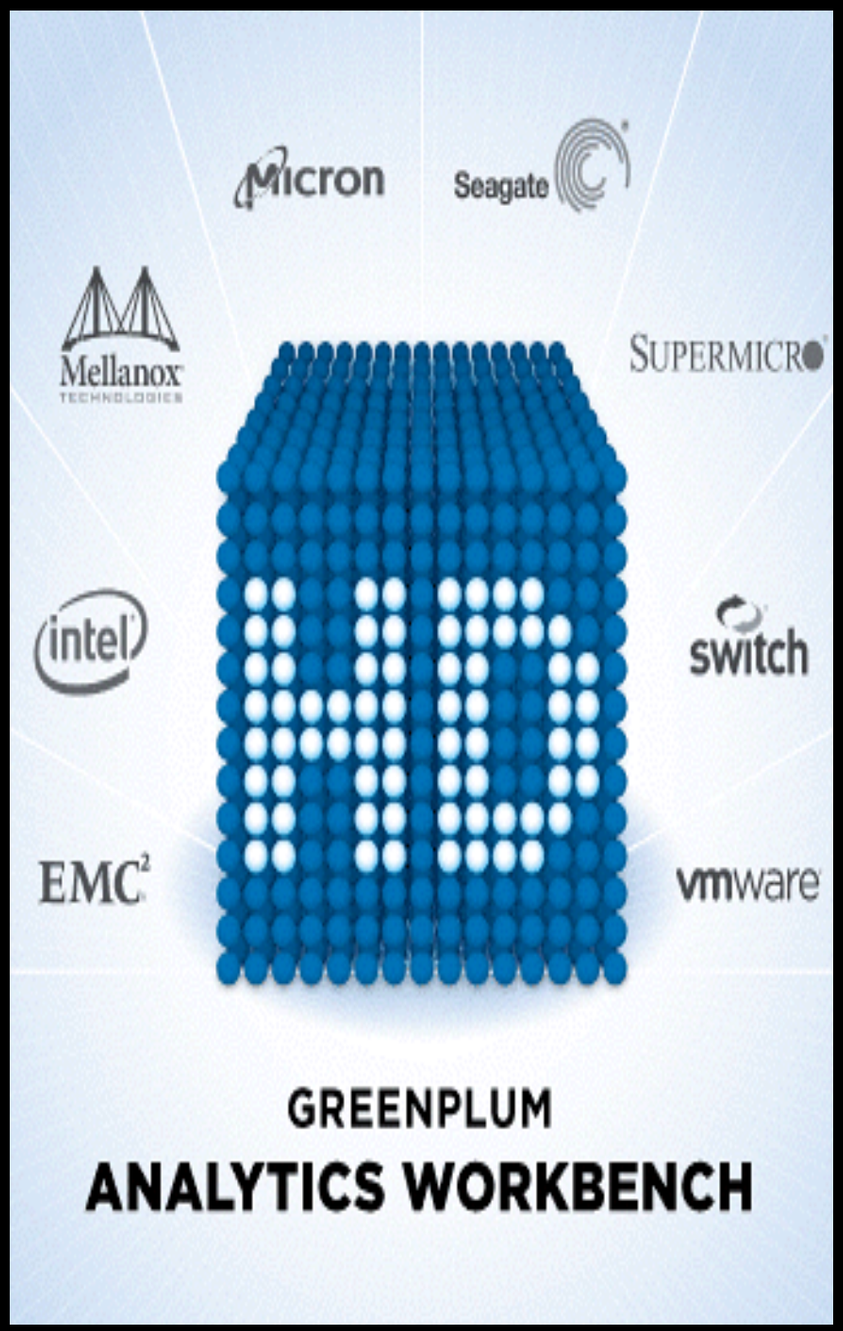

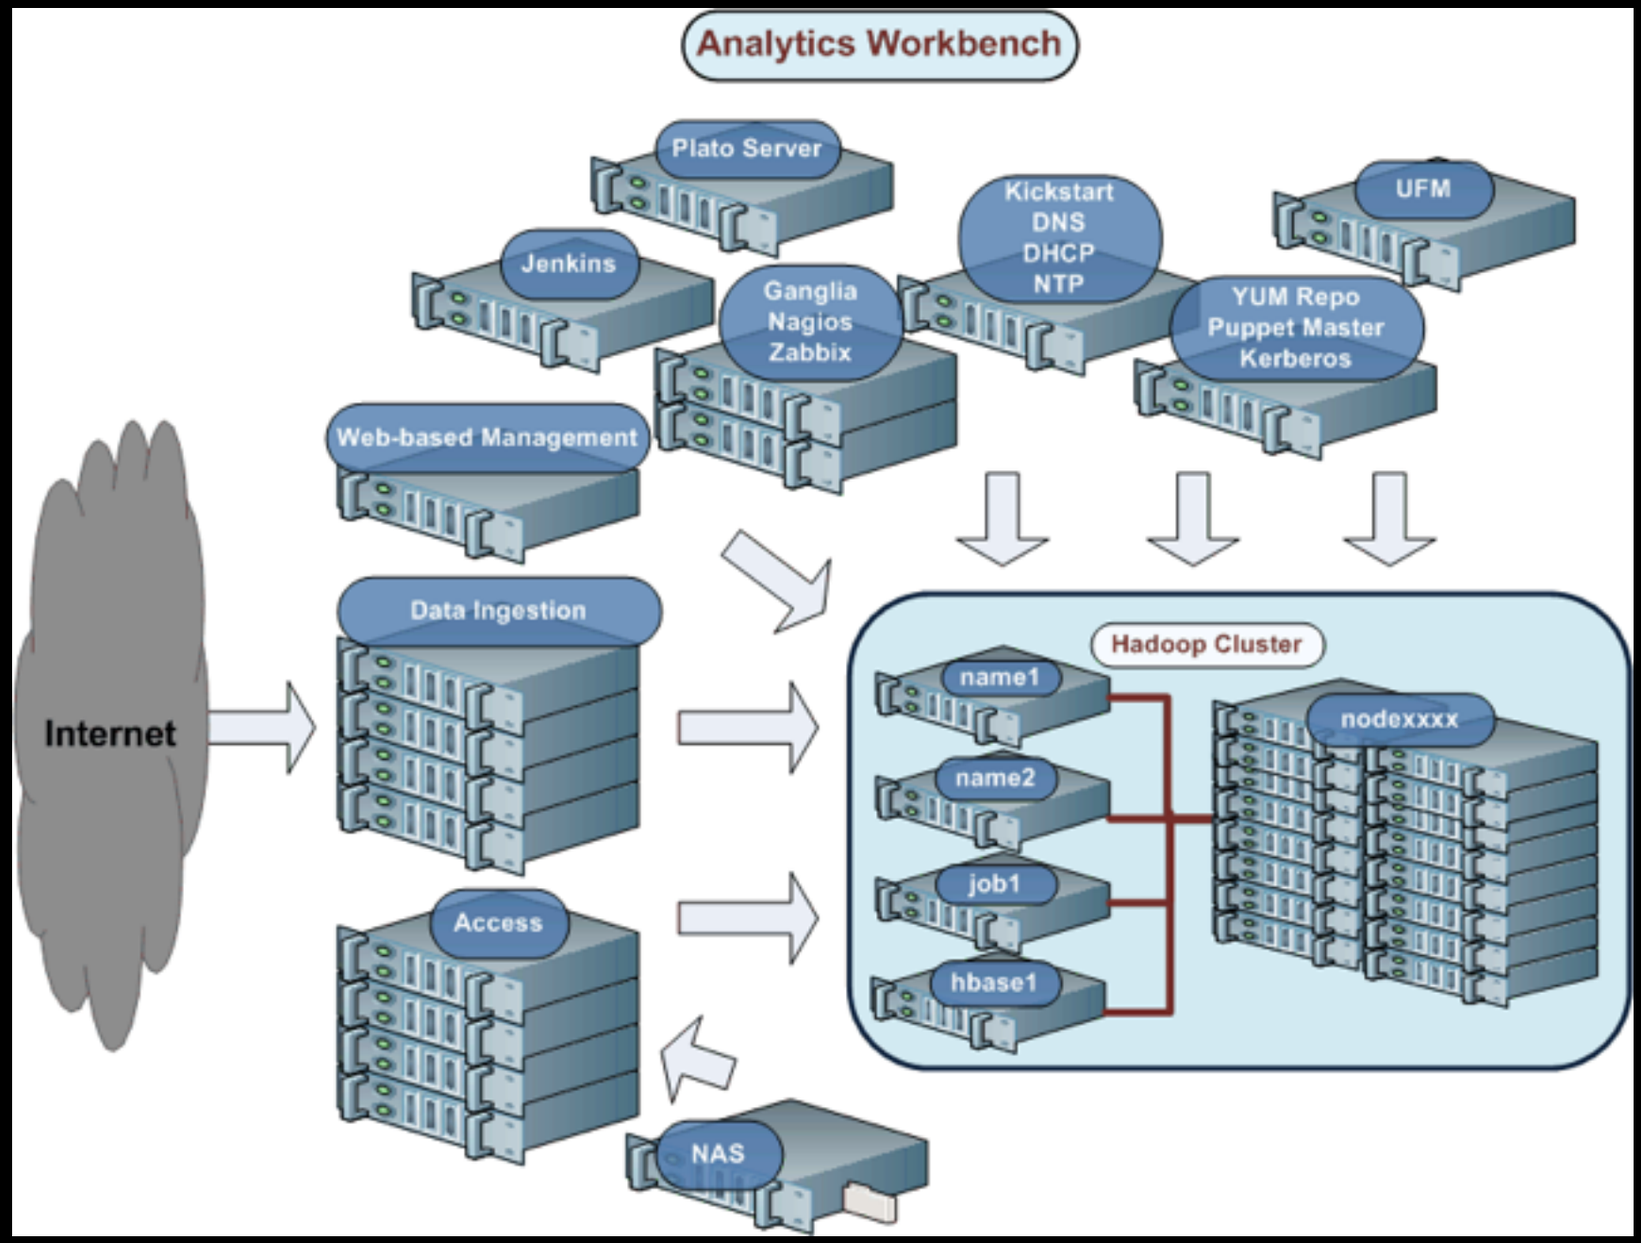

THE WALL STREET JOURNAL.

Why Software Is Eating The World

By MARC ANDREESSEN

70% of data generated by customers

80% of data being stored

3% being prepared for analysis

0.5% being analyzed

<0.5% being operationalized

Average Enterprises

The Big Gap

Example: Healthcare

- In last 5 years
  - 3,573 studies on hospital readmissions
  - 9,745 papers on comparative effectiveness
  - 39,230 studies on drug interactions
  - 132,241 studies on hospital mortality
- Yet, very few models operational

PowerPoint is where Models go to Die.

- Hulya Farinas, Principal Data Scientist, Pivotal

Bits

OCTOBER 9, 2013, 9:00 AM | 12 Comments

G.E.'s 'Industrial Internet' Goes Big

By QUENTIN HARDY

Nathaniel Brooks for The New York Times

Eric Anderson, an engineer at a turbine factory run by General Electric, which has been one of the biggest promoters of the so-called Industrial Internet.

Modernization

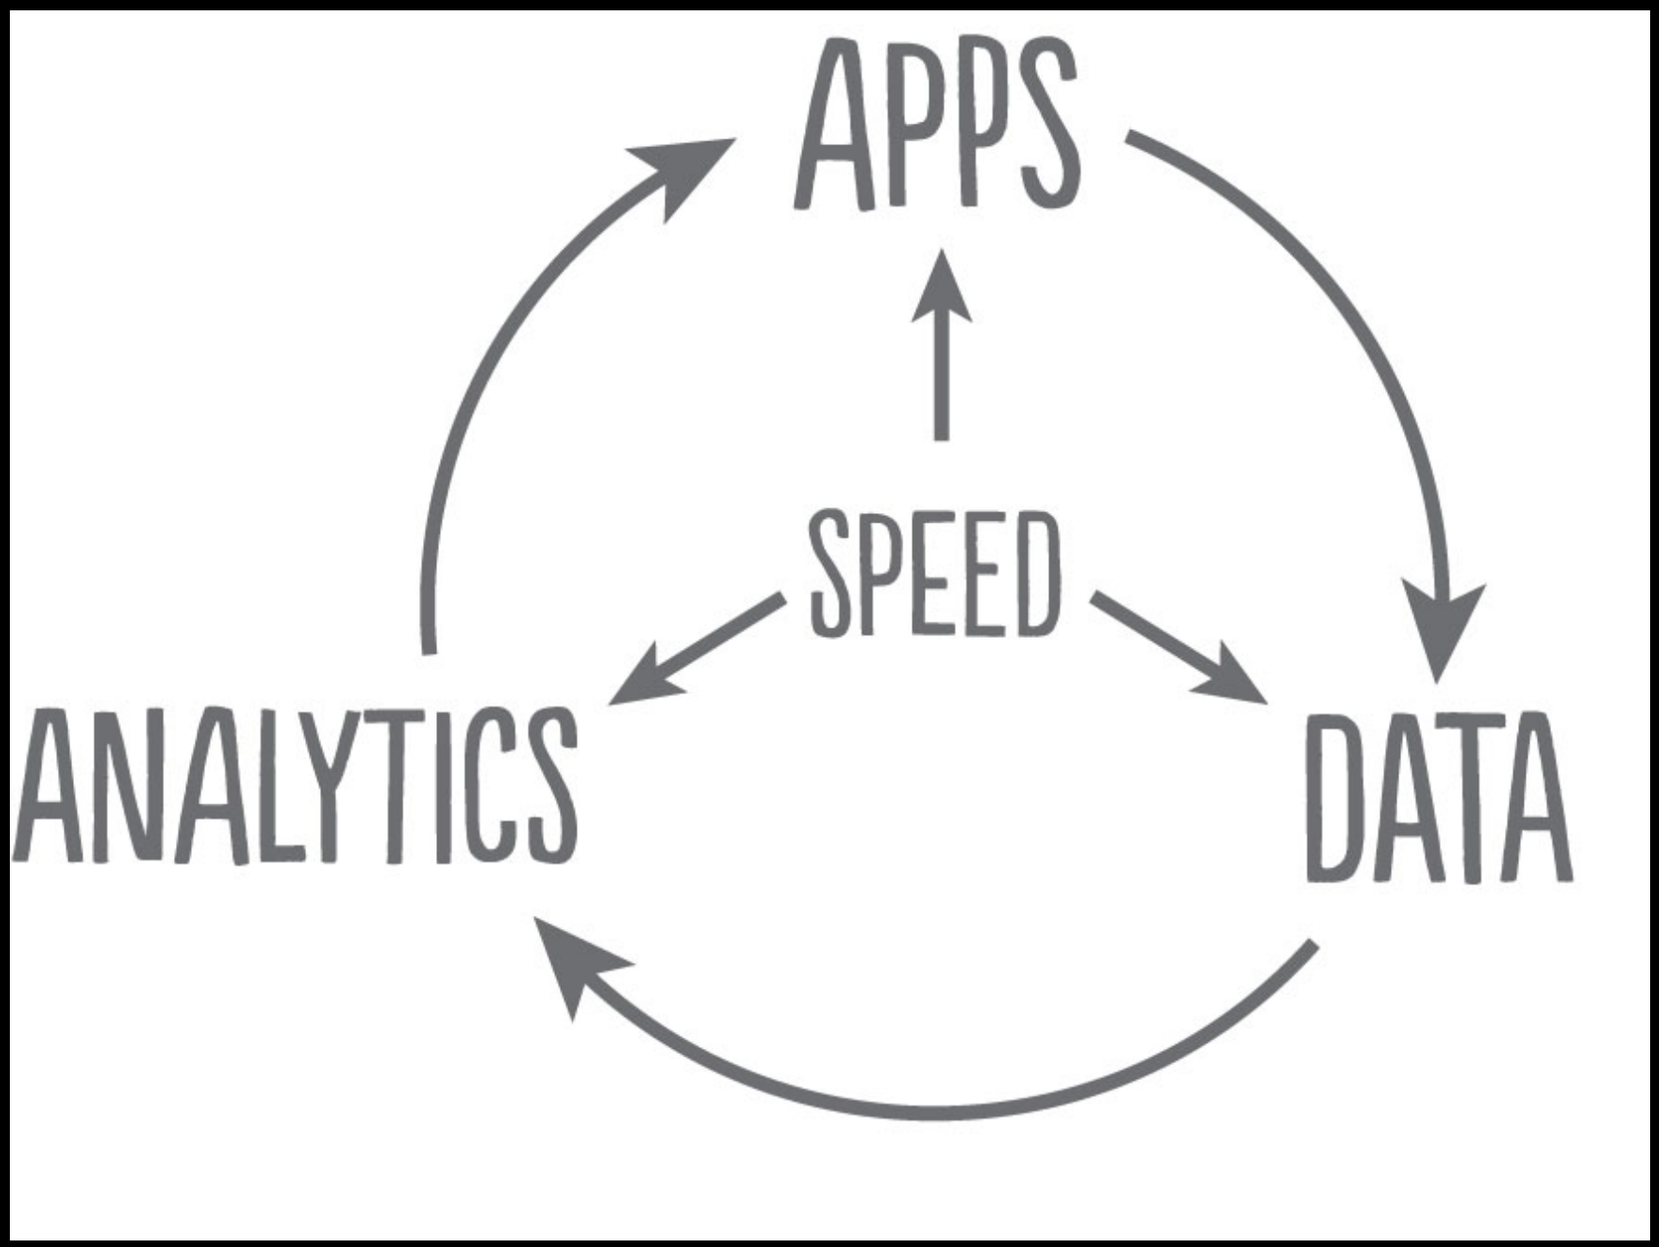

Building Blocks

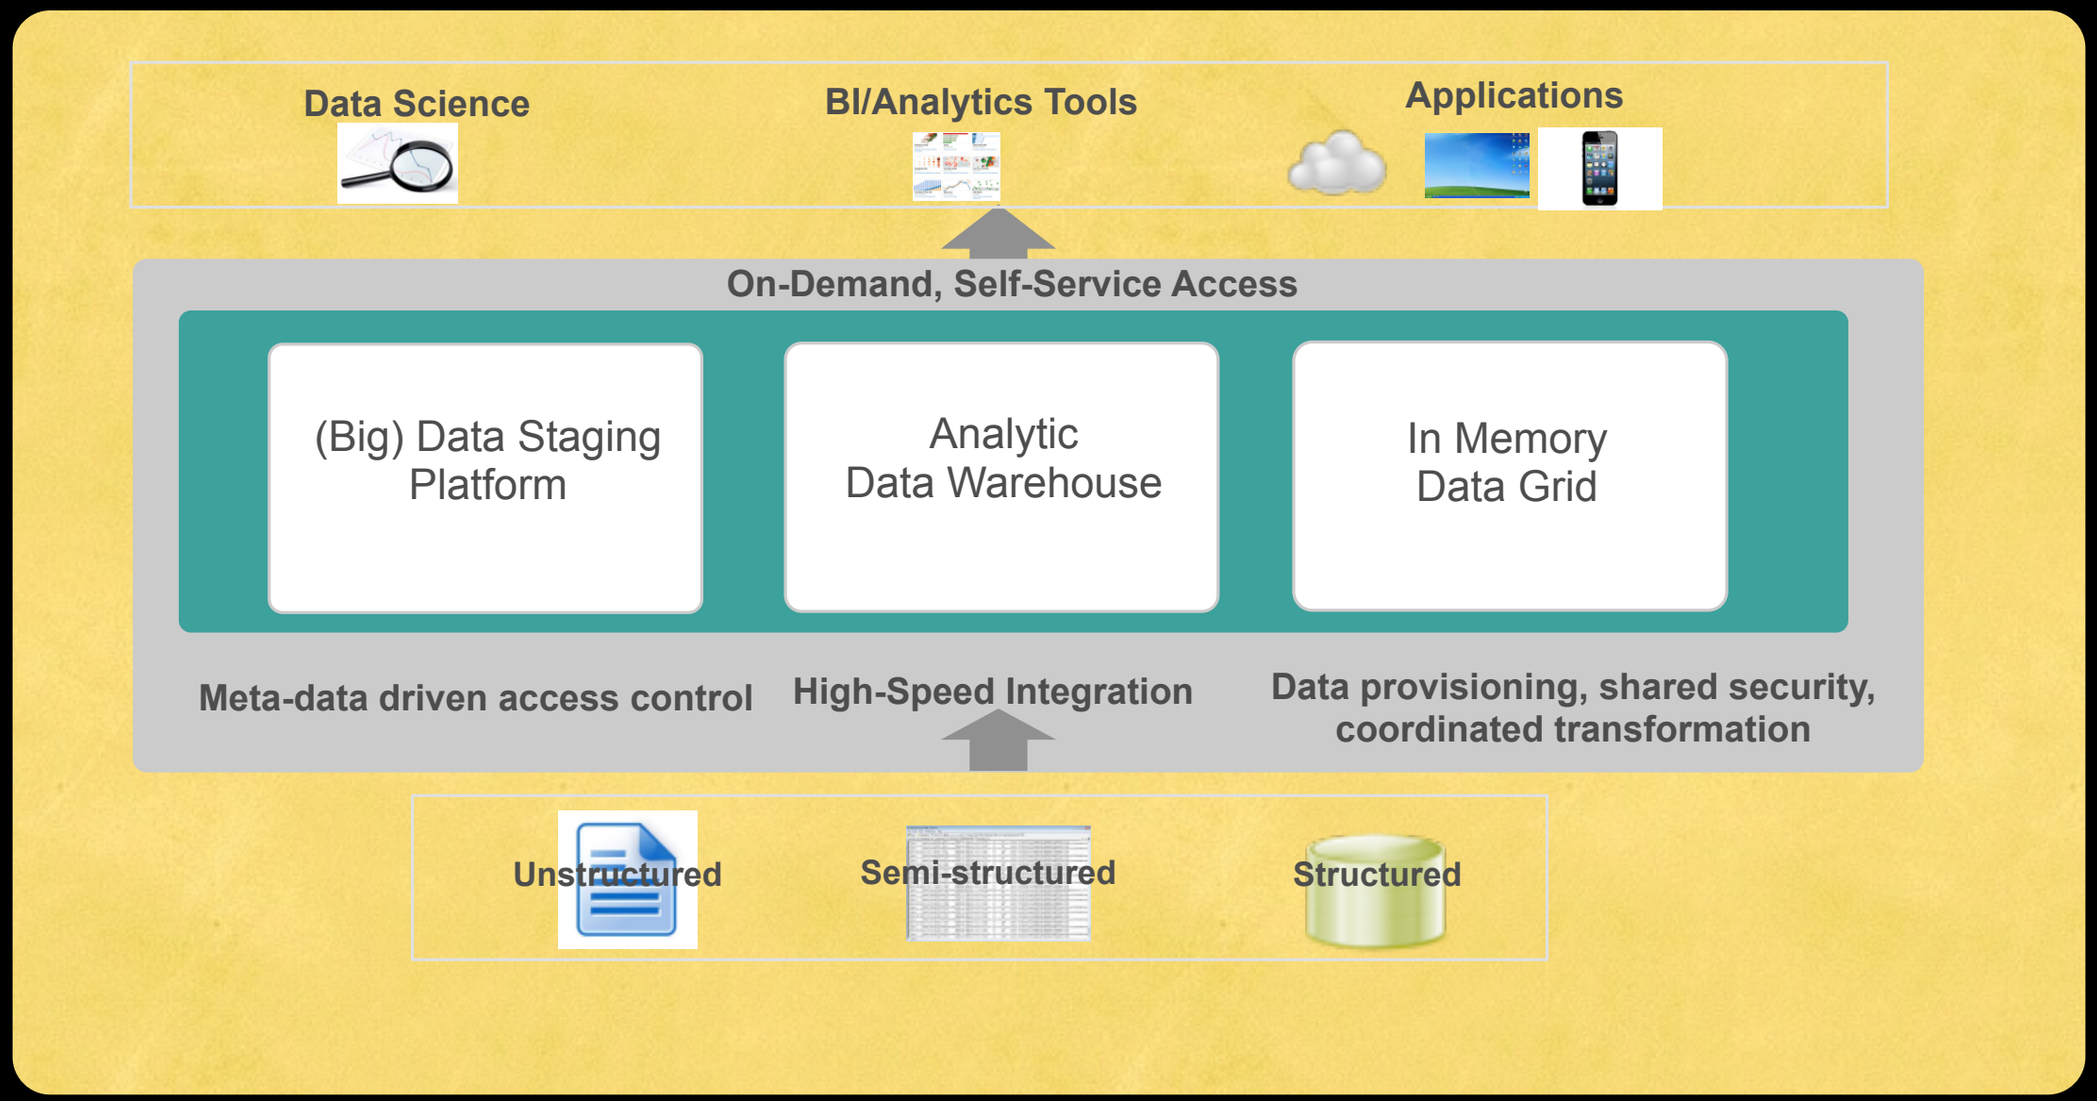

Modern Data Architecture

Data Fabric Requirements

- Store massive & diverse data sets economically
- Integrate and Ingest from legacy & disparate sources
- Ability to rapidly analyze massive data sets
- Control, Auditing, Manageability
- Self-Service

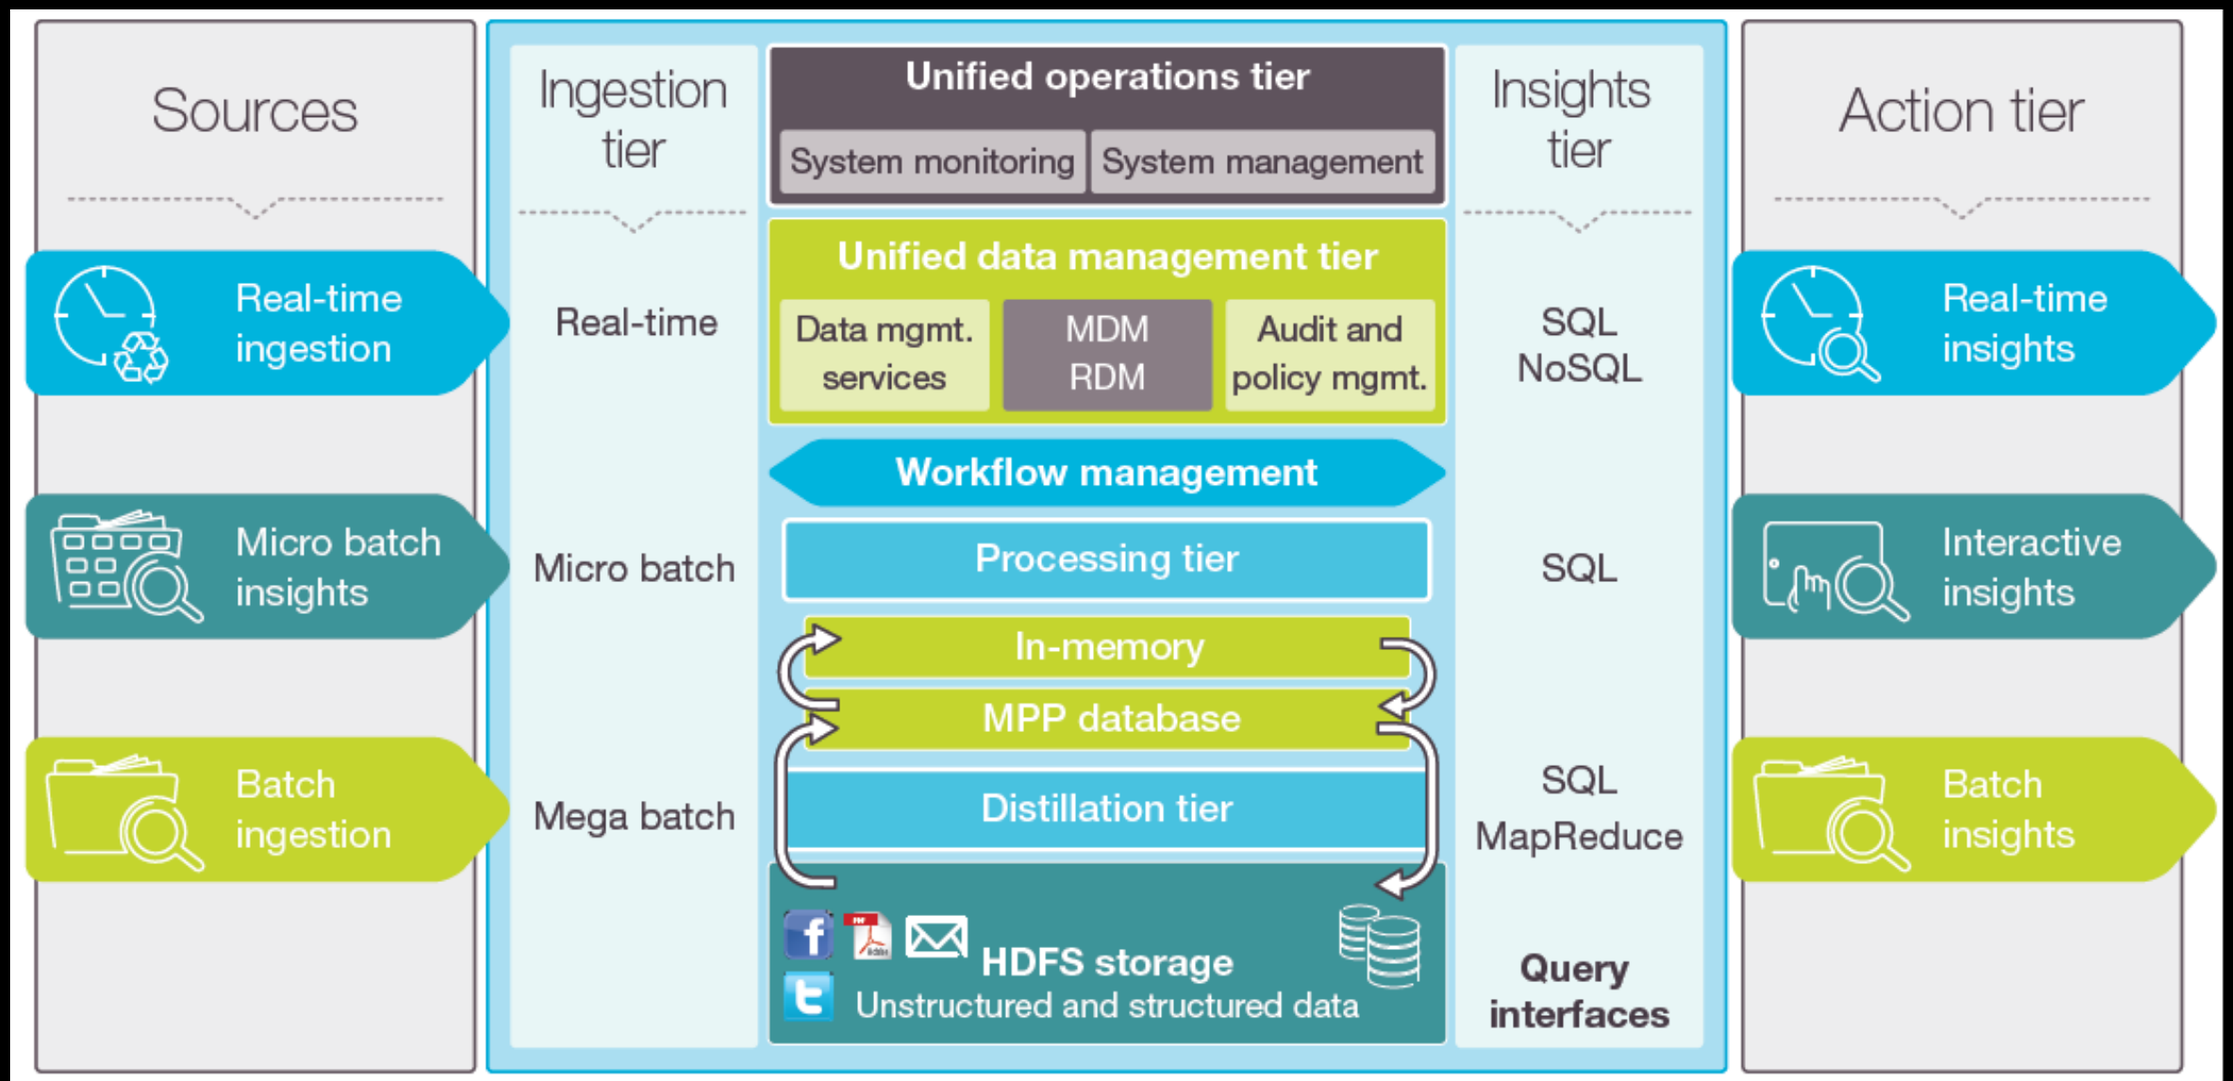

Data Fabric Architecture

Infrastructure-As-A-Service is the new "Hardware"

IAAS: New Hardware

- AWS, GCE, Azure
- vSphere, OpenStack
- Easy Provisioning
- Scalable, Elastic, Ubiquitous
- Needs bundling with Data & Analytics as Services

App Fabric Requirements

- IAAS Cloud-Agnostic
- Rapid provisioning, Elasticity
- Open, No-Lock-In, Data As-A-Service
- Automation for Application Lifecycle Management
- Developer Agility : Eliminate Infrastructure Wiring

CLOUD FOUNDRY™


intel


Bai du 百度


ospero.


Rakuten


COMIC
RELIEF


IBM


vmware


Northern Trust


SAVVIS


Piston
PISTON {CLOUD} COMPUTING


DIEBOLD


NTT Communications

Ecosystem

Broader Ecosystem

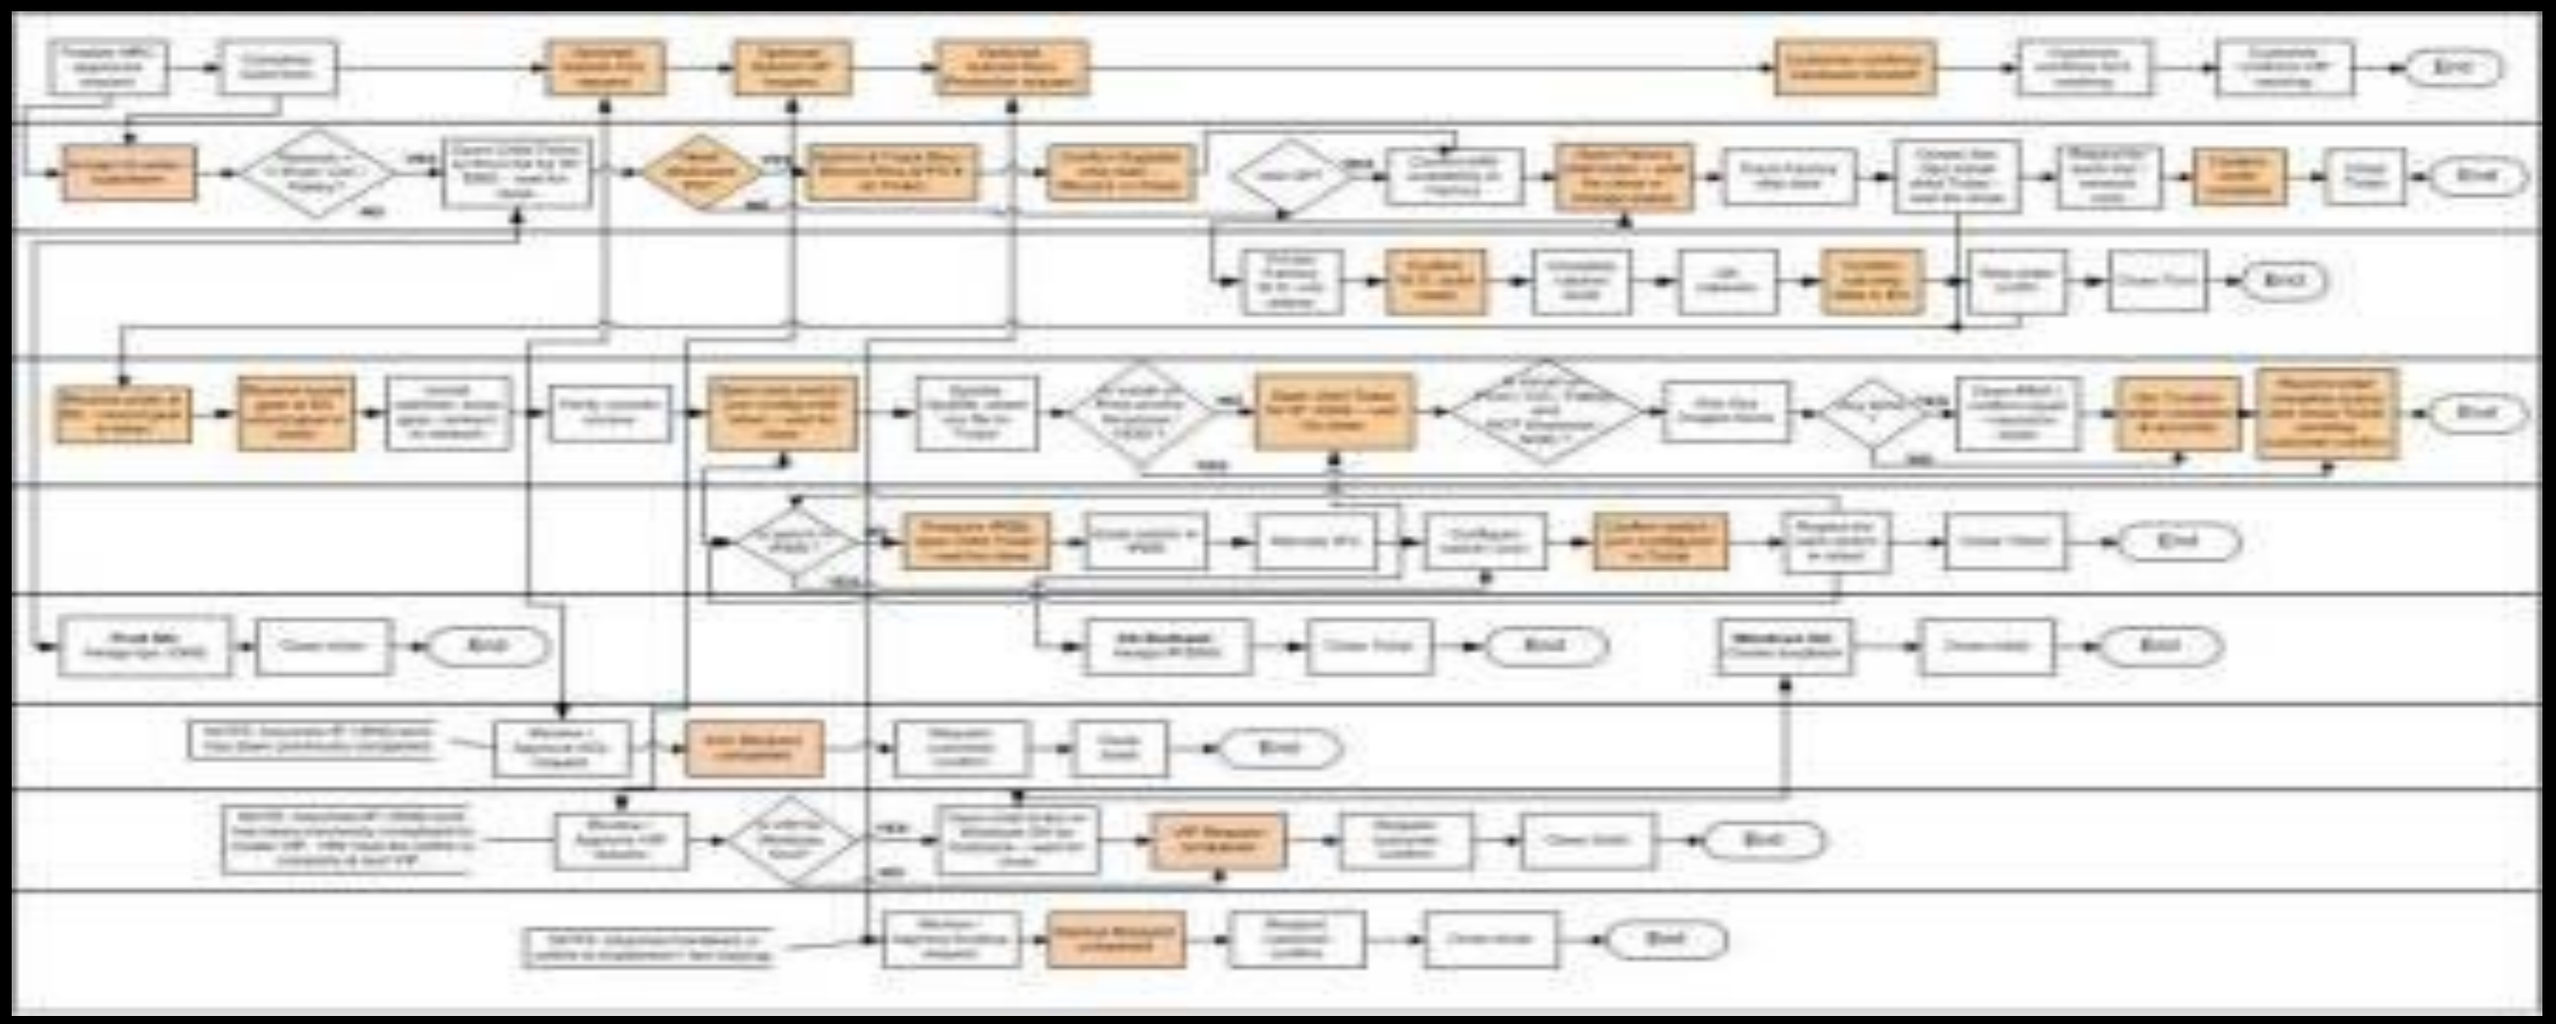

Legacy App Deployment

```
provision <my cloud>
  target <my cloud>
  push <my app>
  bind <my services>
  scale <my app> +100
upgrade <my cloud>
```

Modern App Deployment

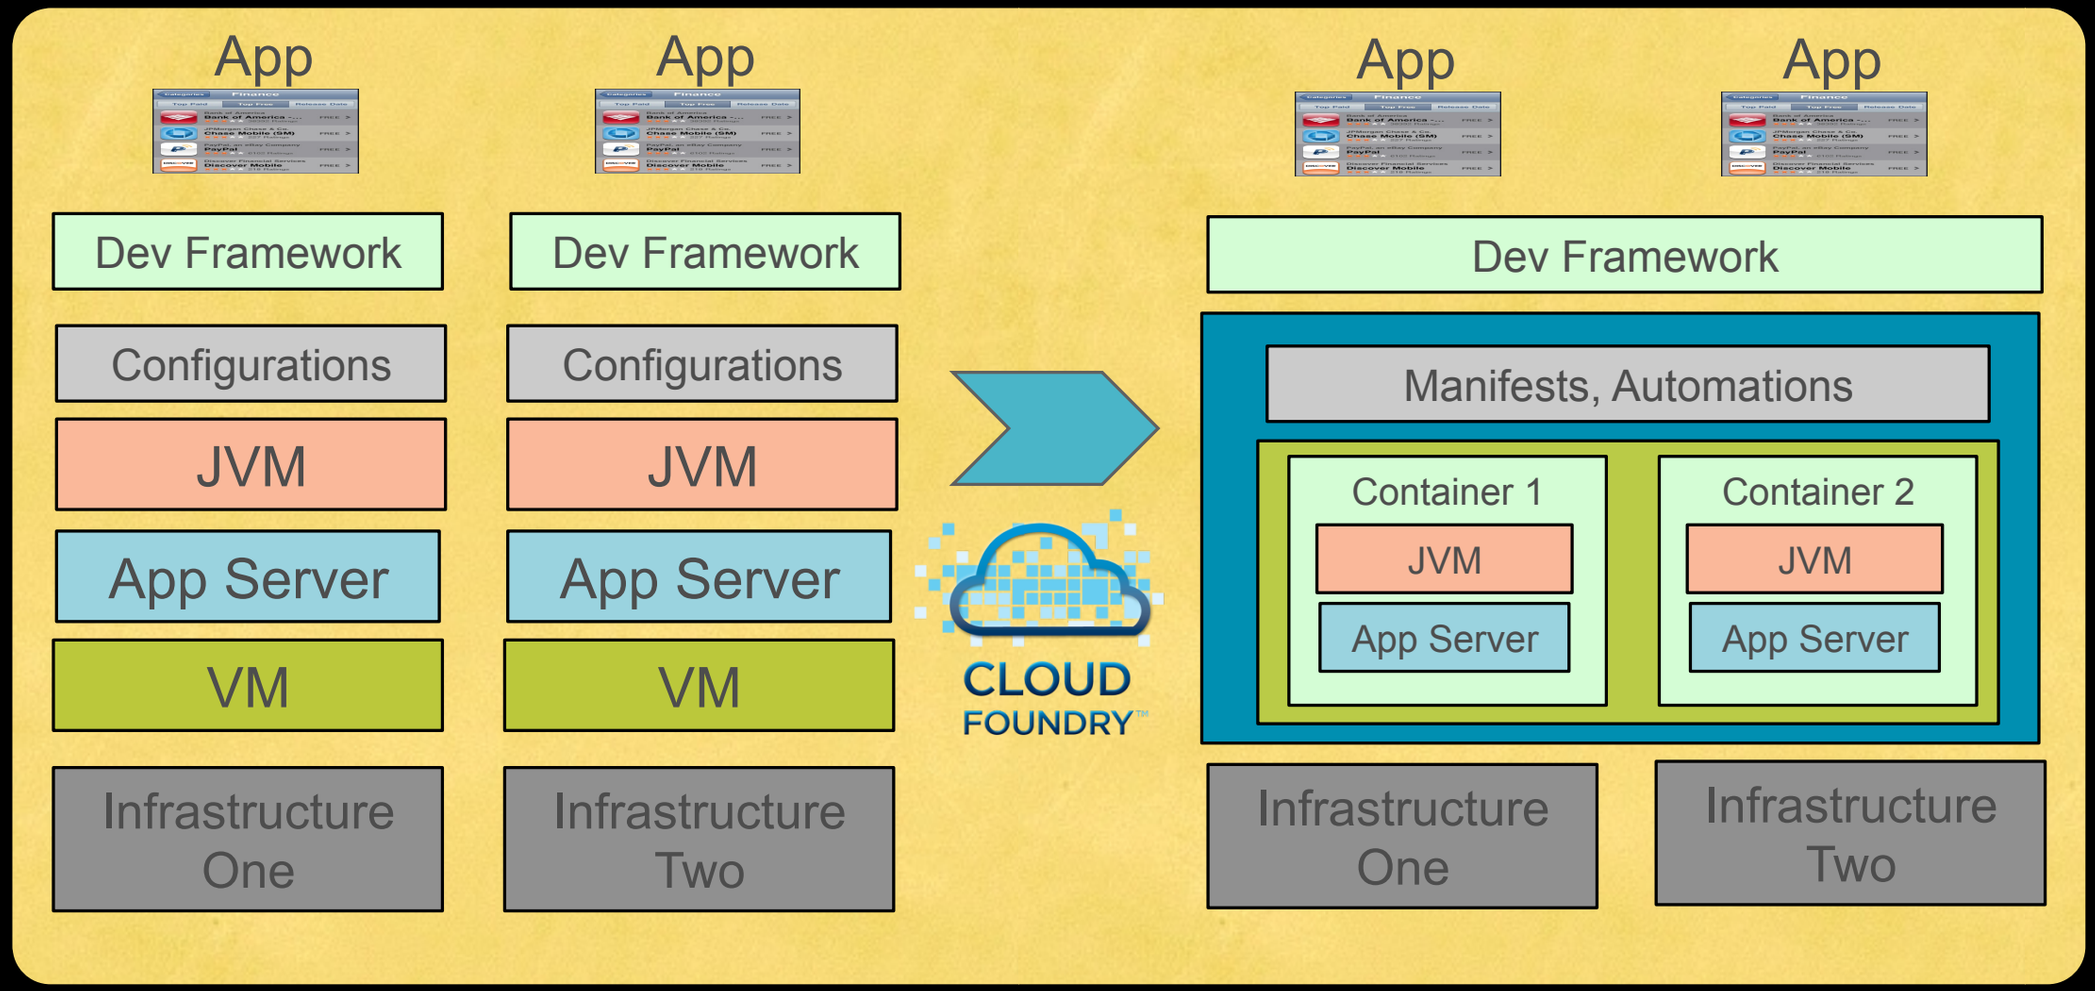

Application As Unit of Deployment

Hadoop's Role in Data Clouds

History (2003-2010)

Google Papers

Inktomi®

+

overture™

=

YAHOO! SEARCH

Yahoo! Search

W-1-W

- WebMap : Graph processing for WWW
- Dreadnaught: Infrastructure for WebMap
- W-1-W: WebMap In One Week
- Juggernaut: Infrastructure for W-1-W
- JFS, JMR, Condor: Abandoned for Hadoop

Lucene

nutch

Lucene, Nutch

Lessons Learned

- Multi-Tenancy from ground-up
- Agility in lieu of Performance
- Provisioning vs Procurement
- "Weird" use cases as learning experience
- Academic collaboration

MapReduce: A major step backwards

on Jan 17 in Database architecture, Database history, Database innovation posted by DeWitt

Major Step Backwards?

MapReduce is the Revenge of System Programmers on Database community.

- Anonymous at XLDB, Stanford, 2010

MapReduce for the Cloud

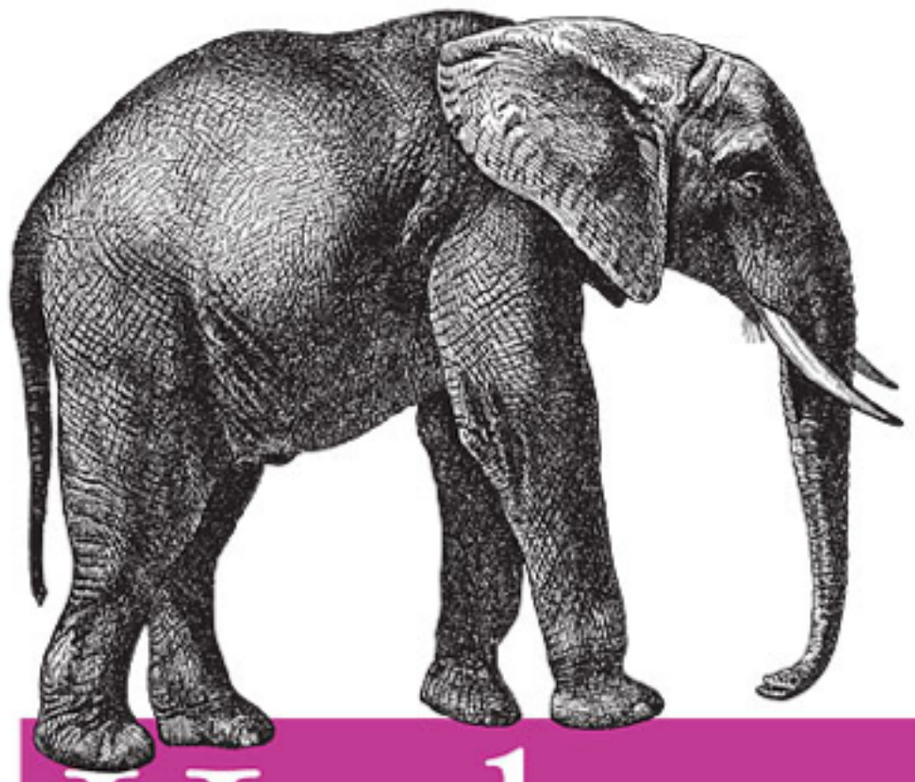

Hadoop

The Definitive Guide

O'REILLY® | YAHOO! PRESS

Tom White

O'Reilly Books 2013

Pivotal™

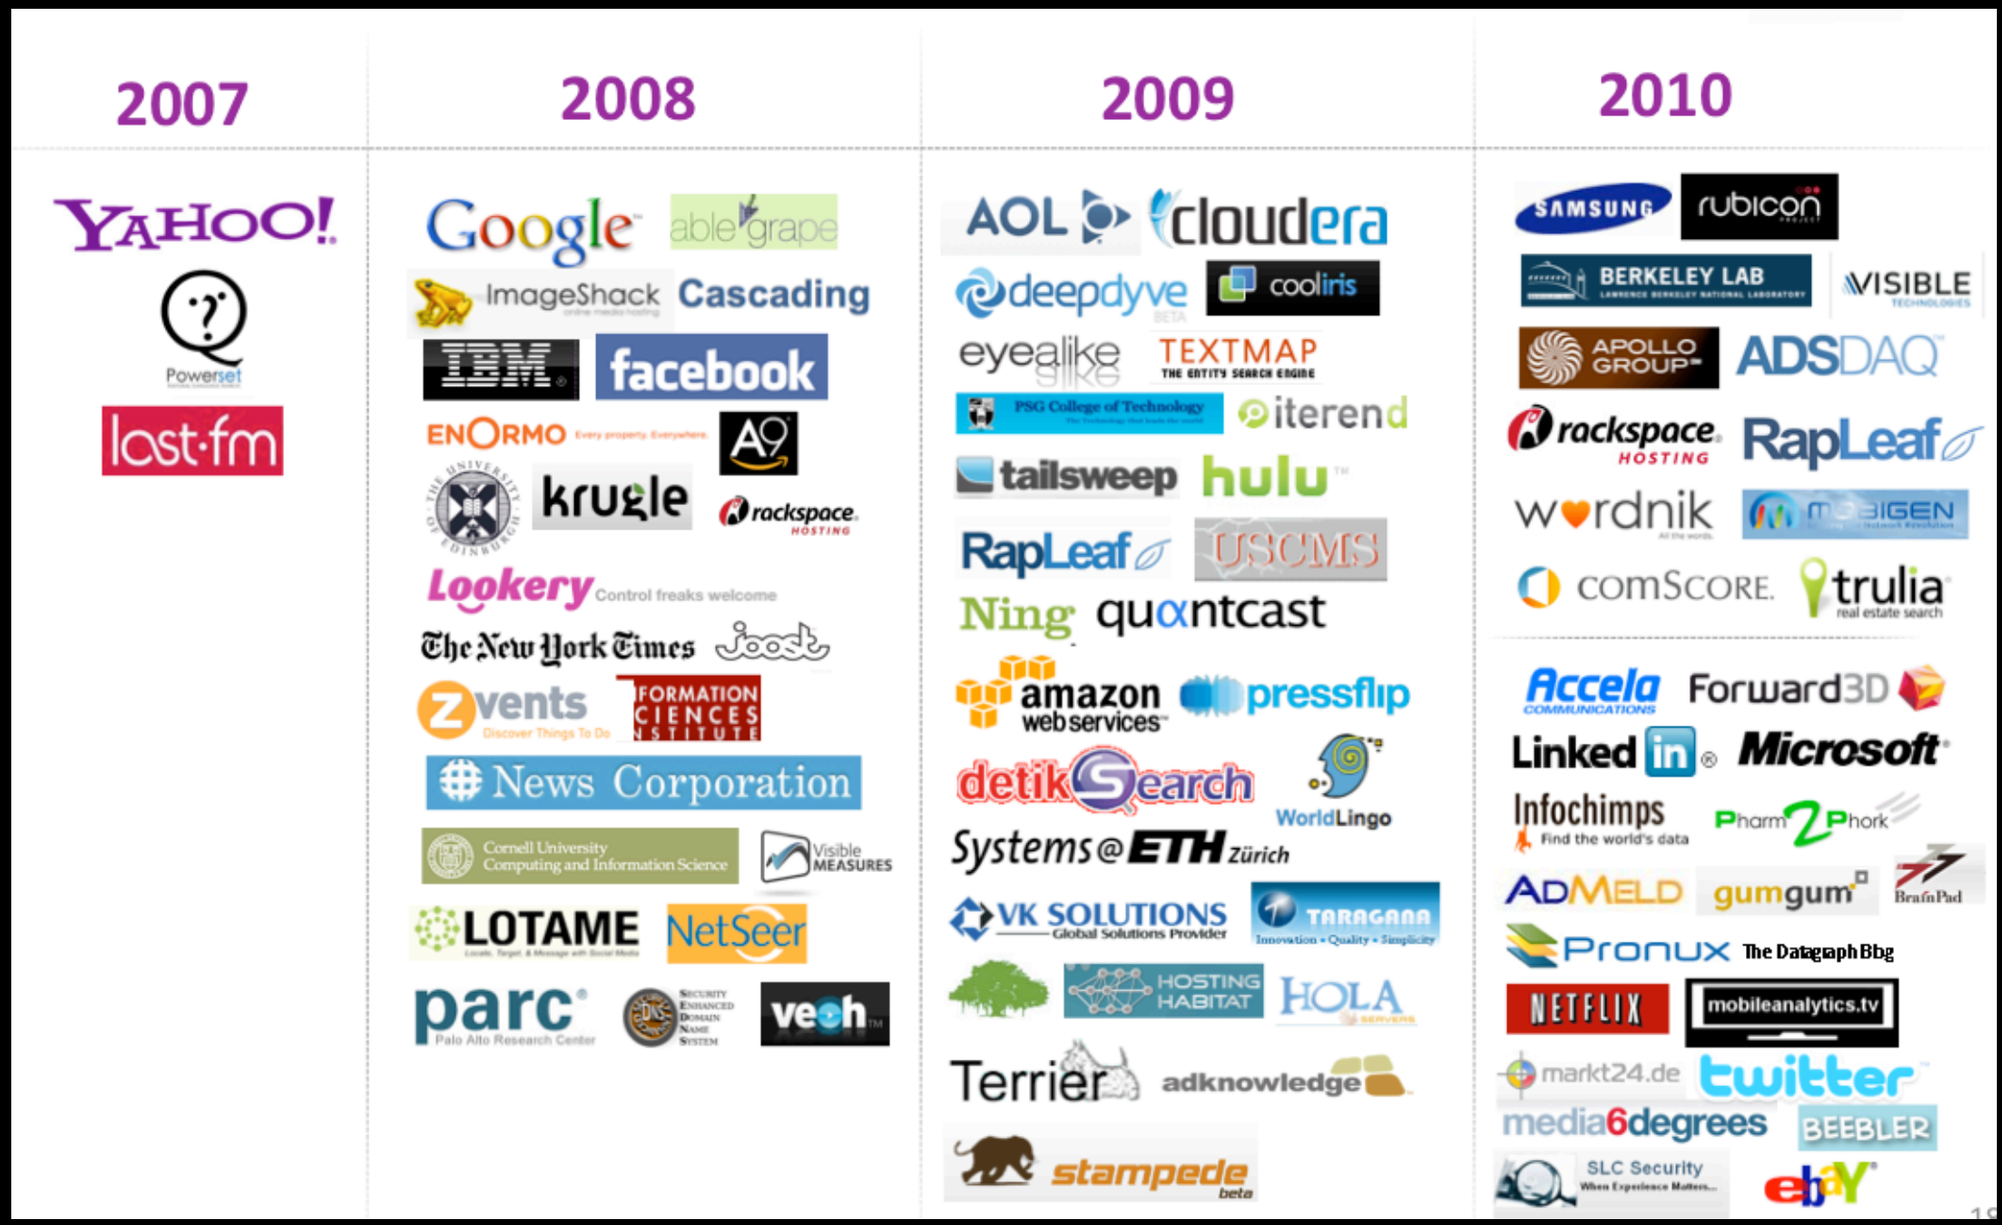

Who Uses Hadoop?

(From Hadoop Summit 2010)

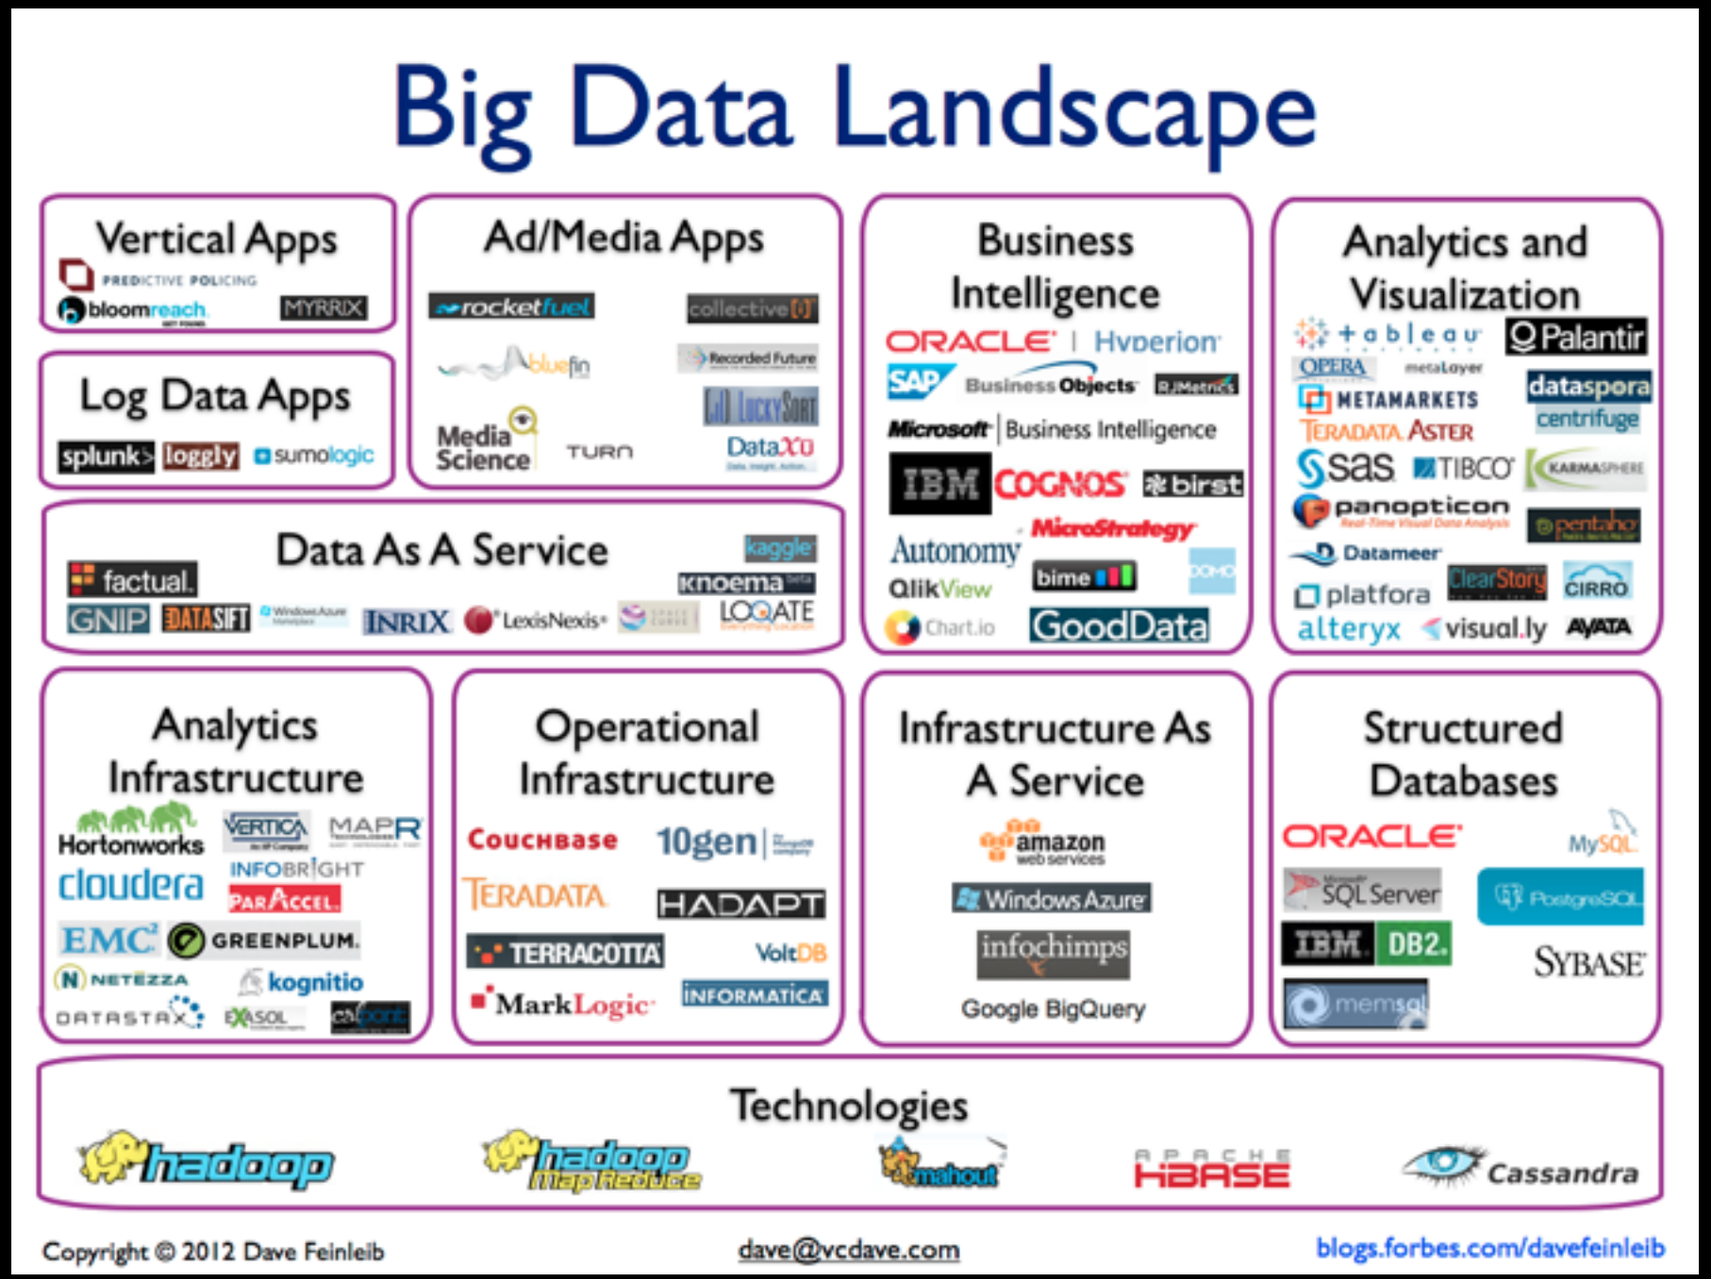

Big Data Landscape - July 2012

http://www.forbes.com/sites/davefeinleib/2012/06/19/the-big-data-landscape/

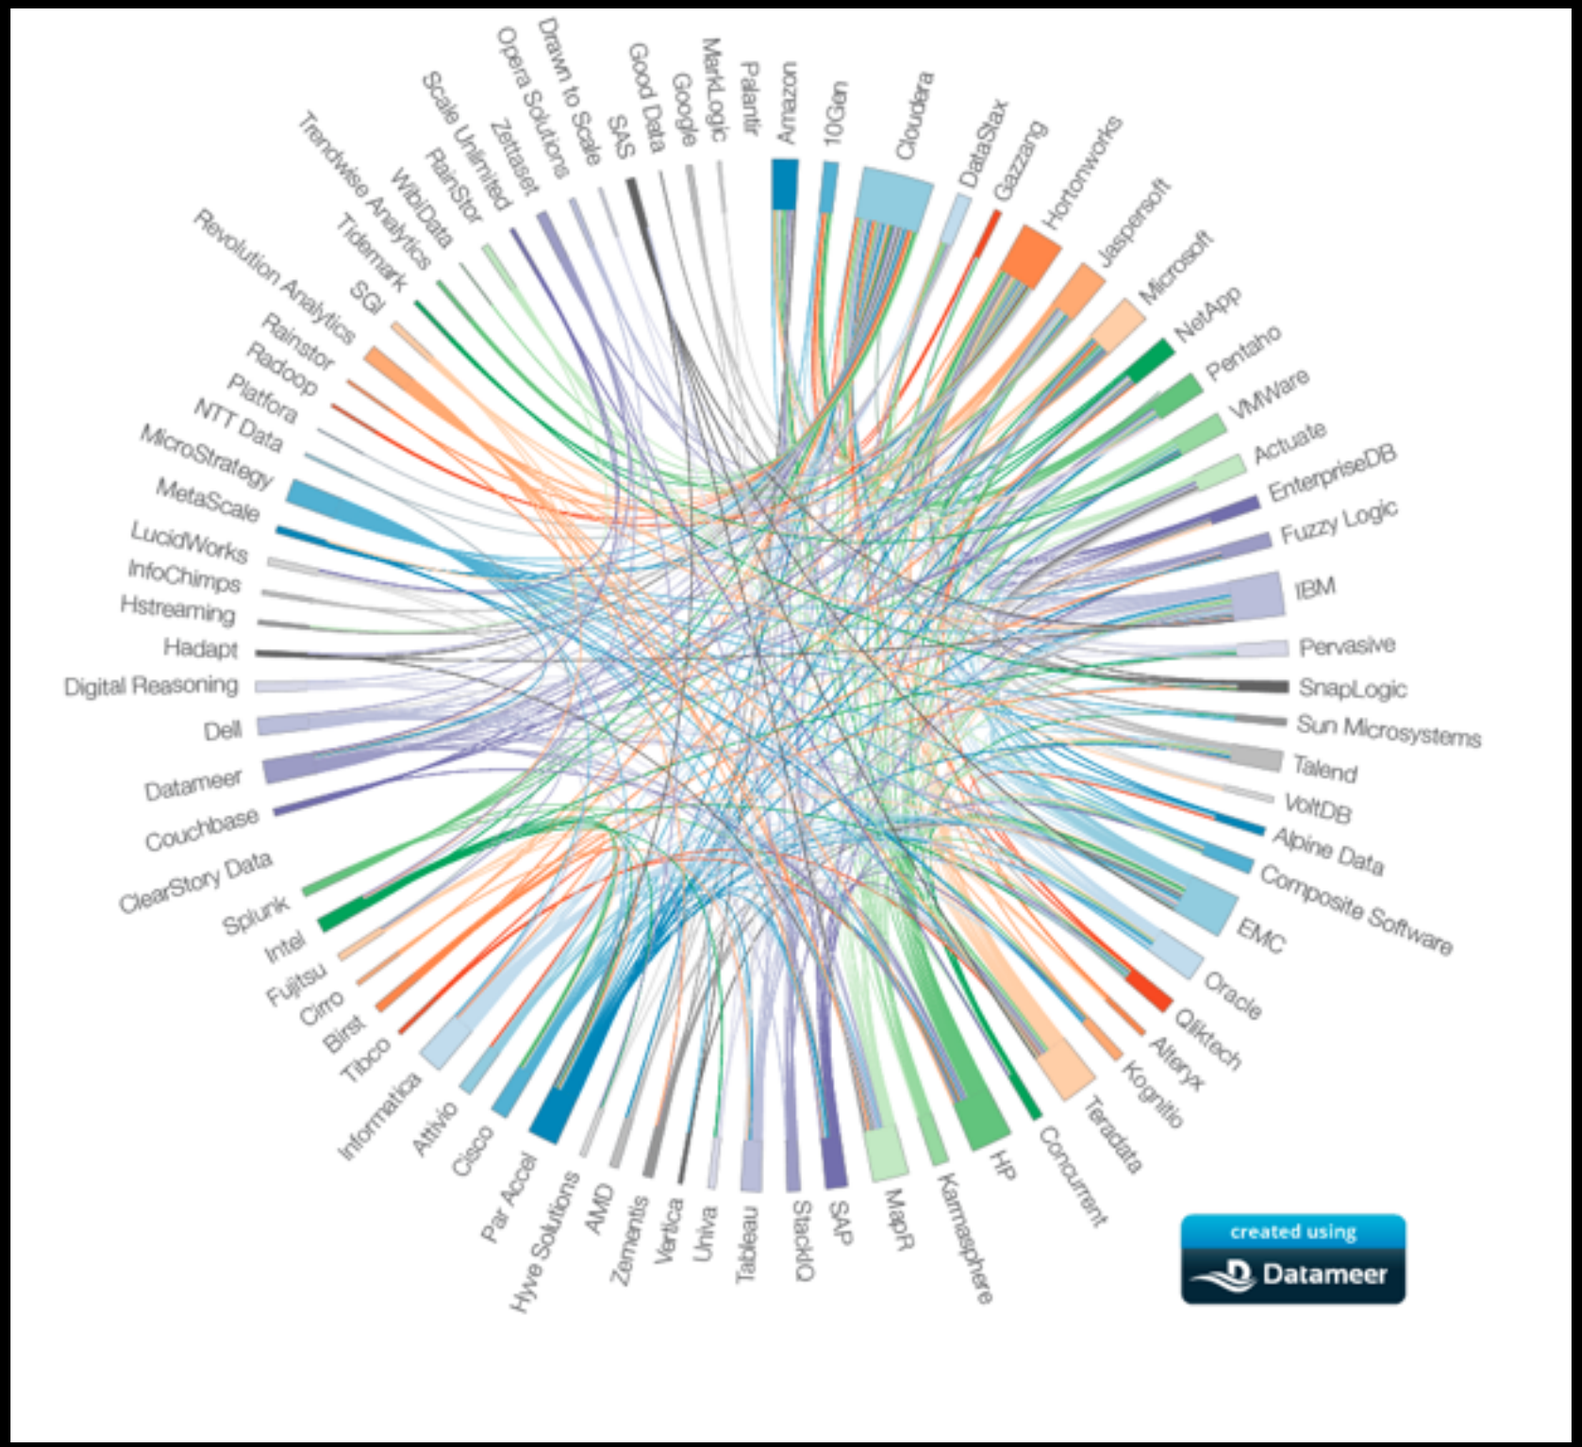

Hadoop Ecosystem (Jan 2013)

http://www.datameer.com/blog/perspectives/hadoop-ecosystem-as-of-january-2013-now-an-app.html

Evolution of Analytics Process (Part 1)

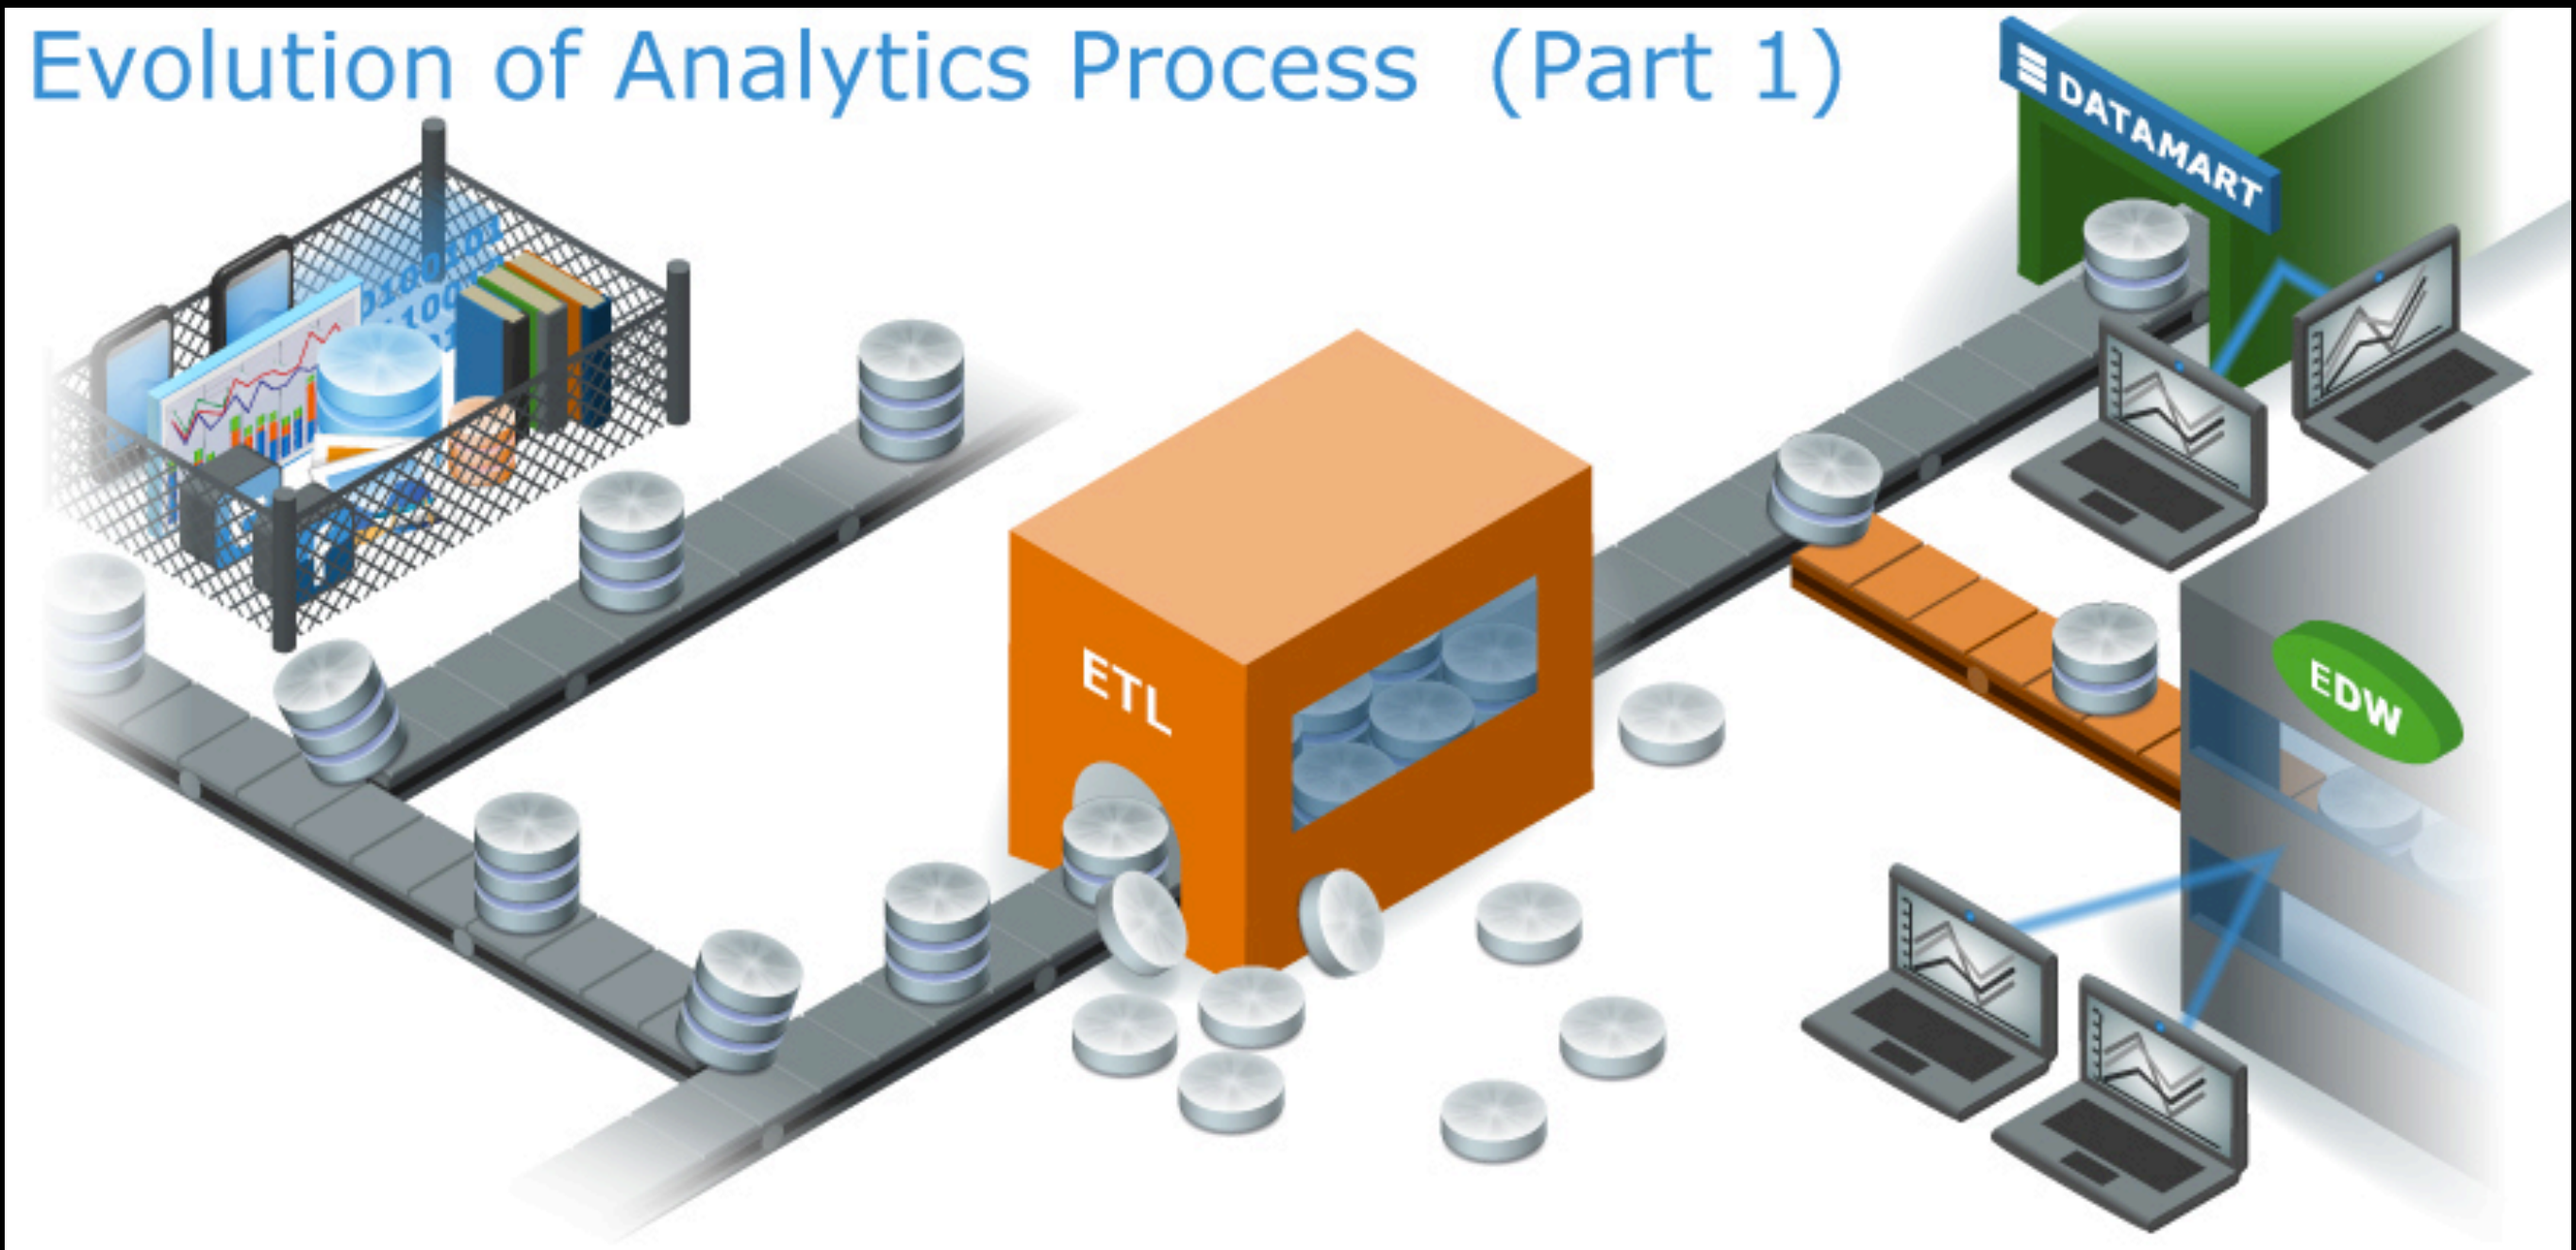

Evolution of Analytics Process (Part 2)

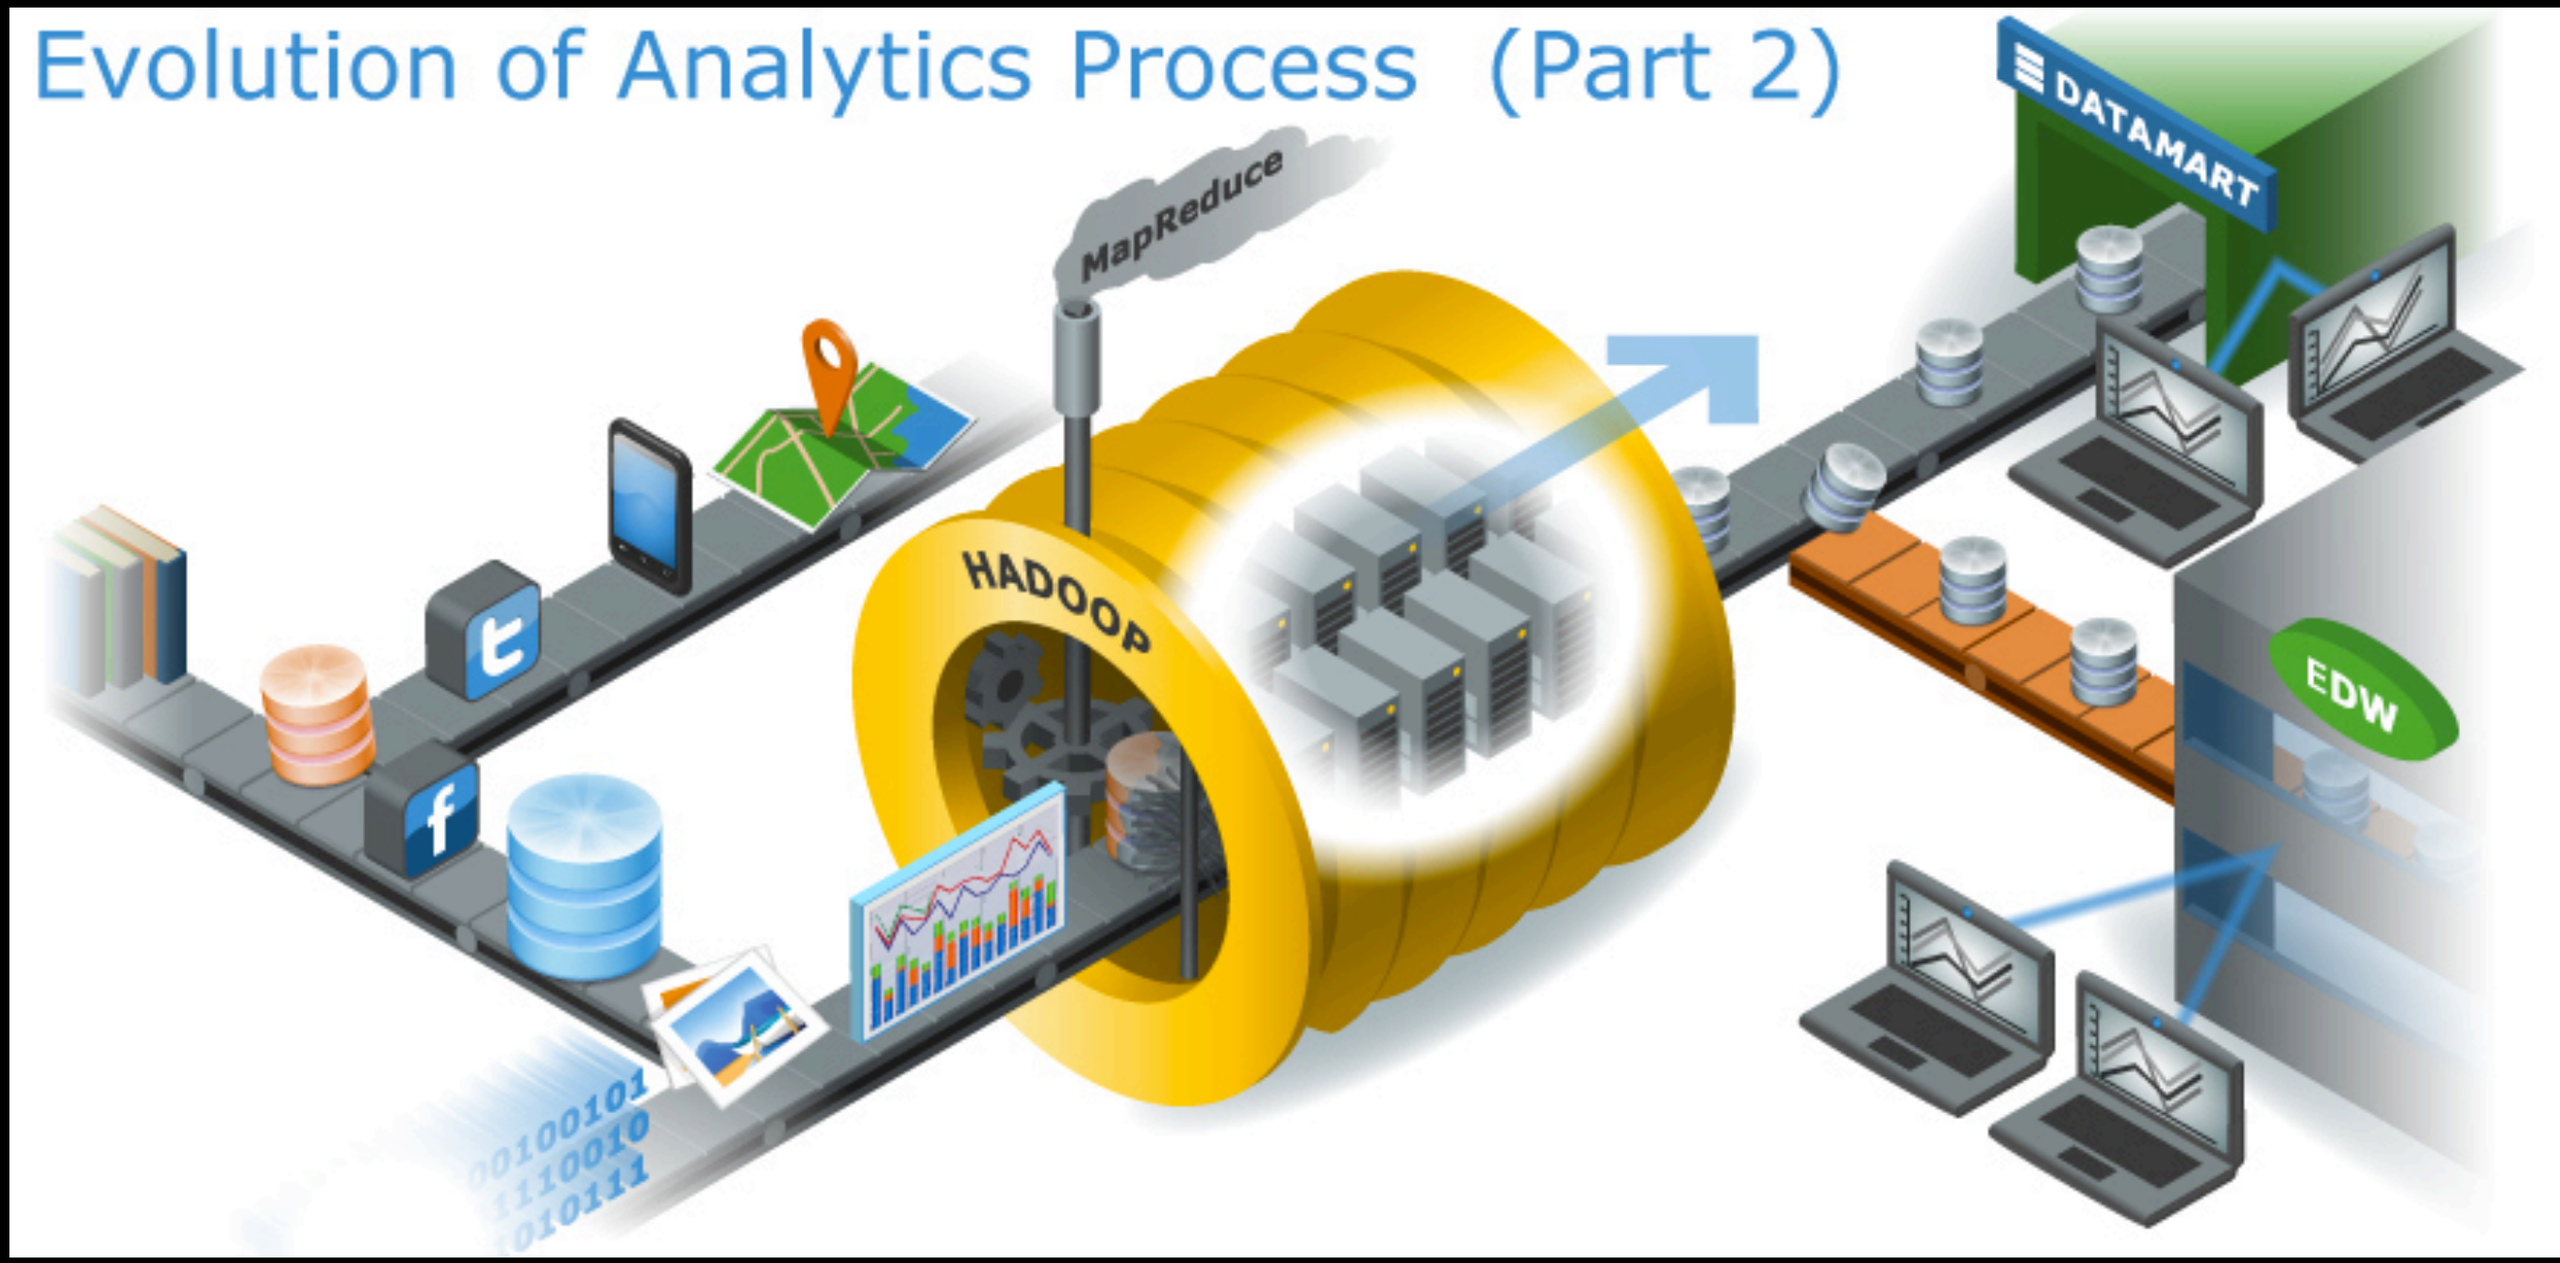

Evolution of Analytics Process (Part 3)

Game Changing Hadoop Economics

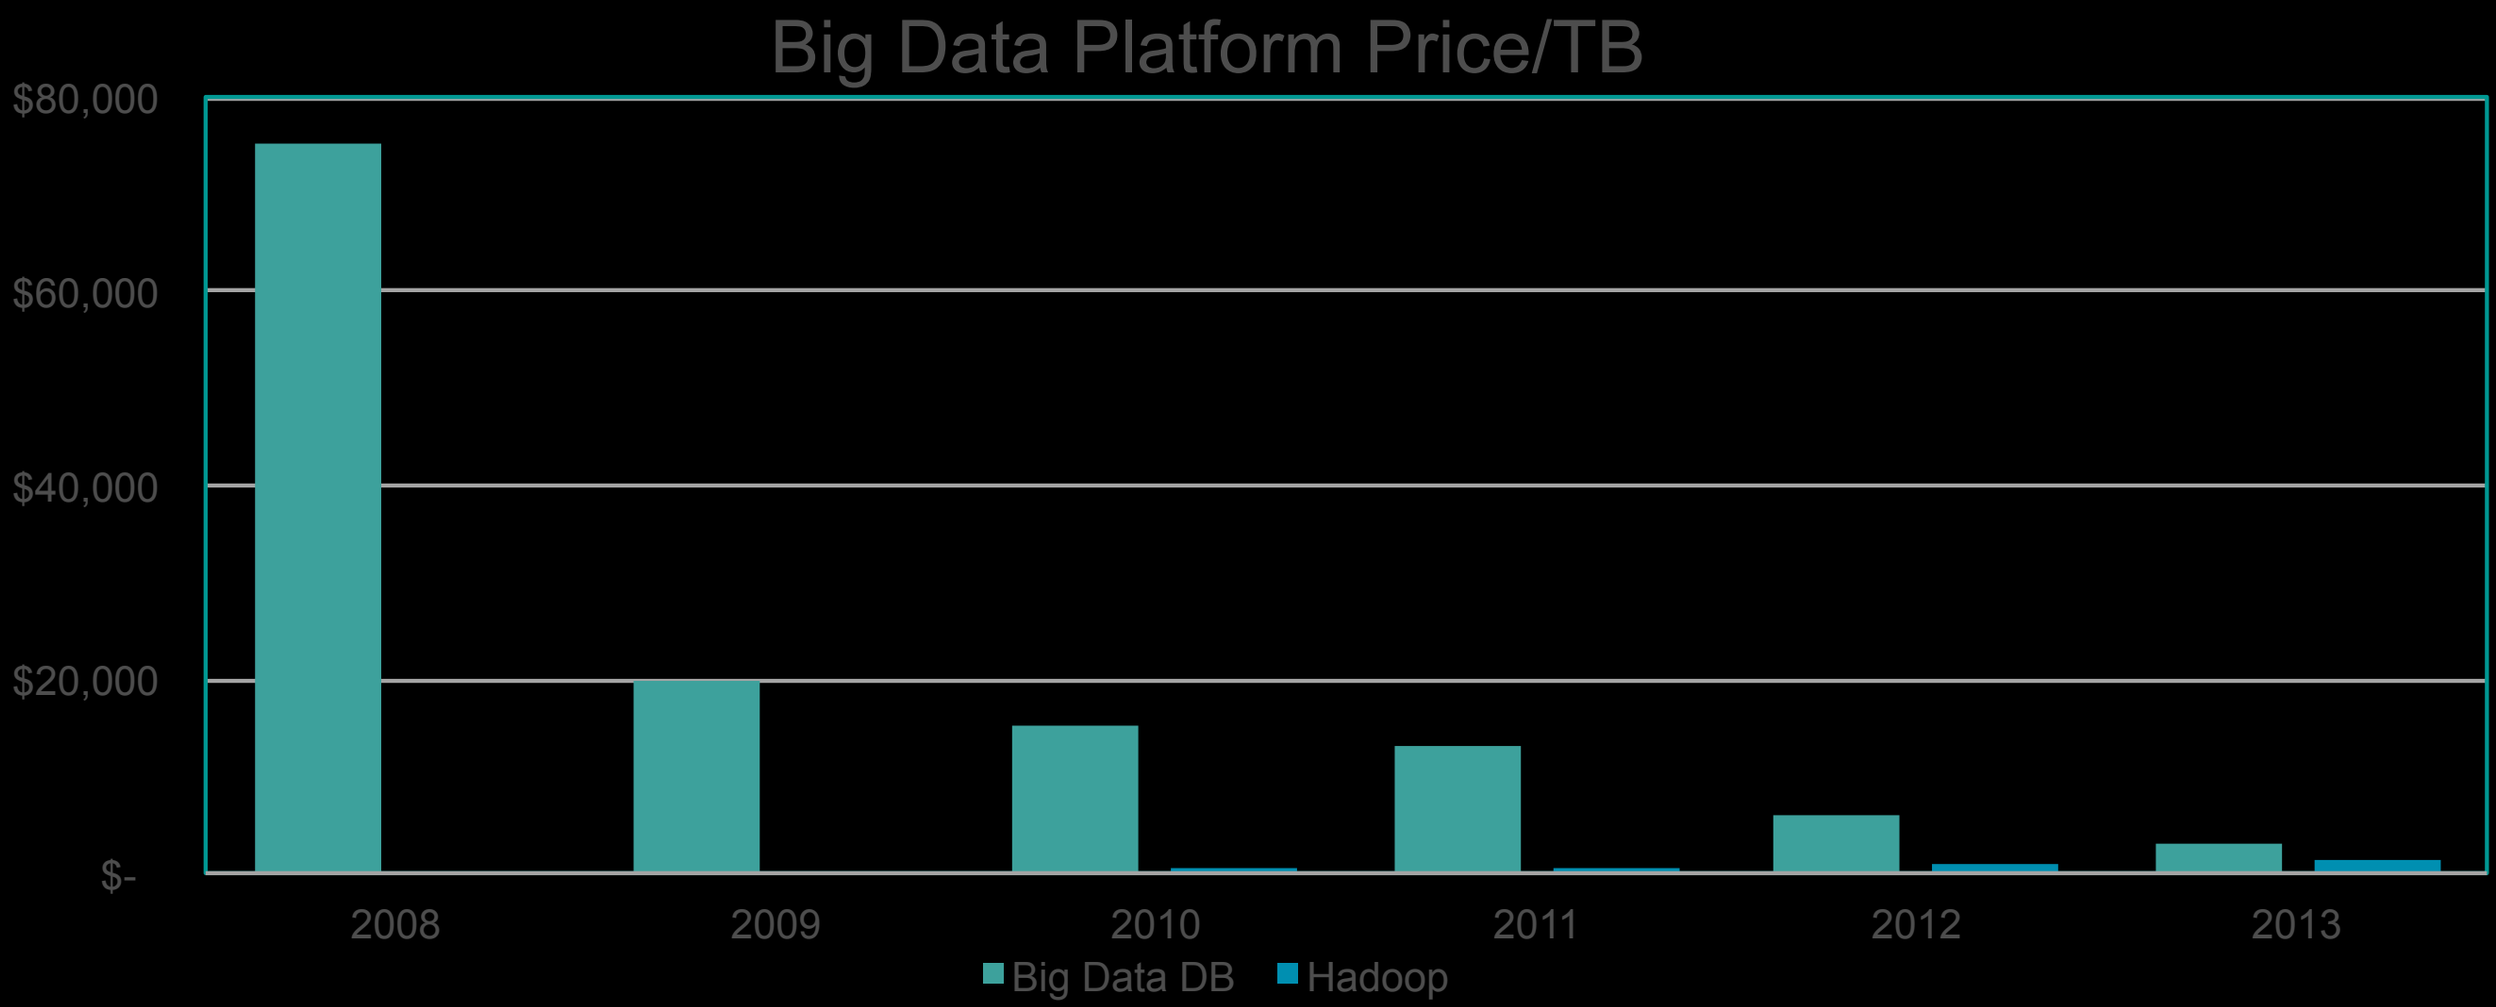

Pivotal™

Hadoop Use Cases

- Data Lake
- ETL Offload
- Deep Analytics

Deep Analytics

- "User" Modeling
- Objective: Determine User-Interests by mining user-activities
- Large dimensionality of possible user activities
- Typical user has sparse activity vector
- Event attributes change over time

Retail

- User = Customer
- Activities
  - Online: Purchase, Ad click, FB Likes
  - Offline : Brick-and-mortar purchases, returns, coupon clipping, gift cards
- Goal: Personalized Product Recommendation

IT Infrastructure

- "User" = HW & SW Components
- Activities
  - Log messages, Metrics, connectivity, communication events
- Goal: Proactive alerting of imminent failures

Healthcare

- User = Patient
- Activities
  - Doctor Visits, Medicine refills, Medical History
  - 3G/WiFi-enabled Pillbox...
- Goal: Prevent Hospital Readmissions

Telecom

- User: Subscriber
- Activities
  - Calls made, duration, calls dropped, locations, ...
  - “social” graph, status updates
- Goal: Reduce customer churn

Ad-Supported Web

- User = User :-)
- Activities
  - Clicks on content, Likes, Repost
  - Search Queries, Comments, Participation
- Goal: Increase Engagement, Increase Clicks on revenue-generating content (ads/ premium content)

Deep Analytics Pipeline

- Sessionization
- Feature and Target Generation
- Model Training
- Offline Scoring & Evaluation
- Batch Scoring & Upload to serving

What's Next ?

Hype Cycle for Emerging Technologies, 2013

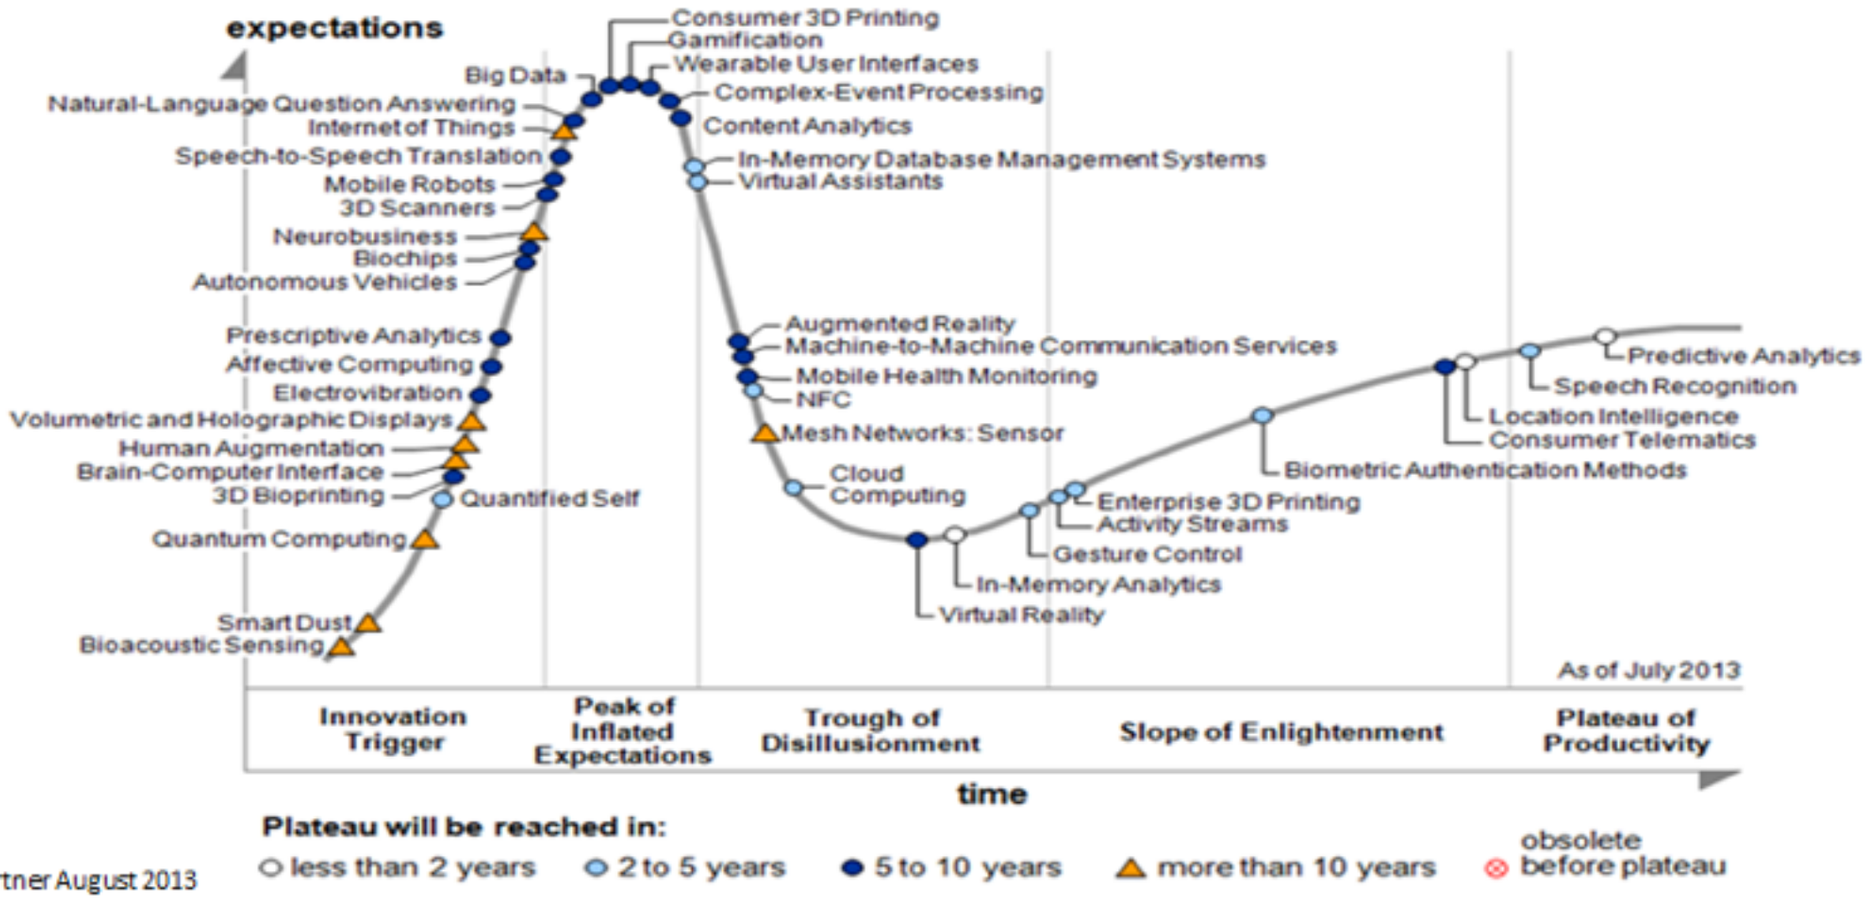

Source: Gartner August 2013

www.facebook.com/EMCacademicalliance

EMC²

Trough of Disillusionment ?

Pivotal

News

Hadoop will be in most advanced analytics products by 2015, Gartner says

Or, Hadoop Everywhere?

Storage Options

- HDFS, MapR, Quantcast QFS
- EMC Isilon, NetApp, IBM GPFS, PanFS, PVFS, Lustre
- Amazon S3, EMC Atmos, OpenStack Swift
- GlusterFS, Ceph
- EMC ViPR

SQL-on-Hadoop

- Pivotal HAWQ
- Cloudera Impala, Facebook Presto, Apache Drill, Cascading Lingual, Optiq, Hortonworks Stinger
- Hadapt, Jethrodata, IBM BigSQL, Microsoft PolyBase
- More to come...

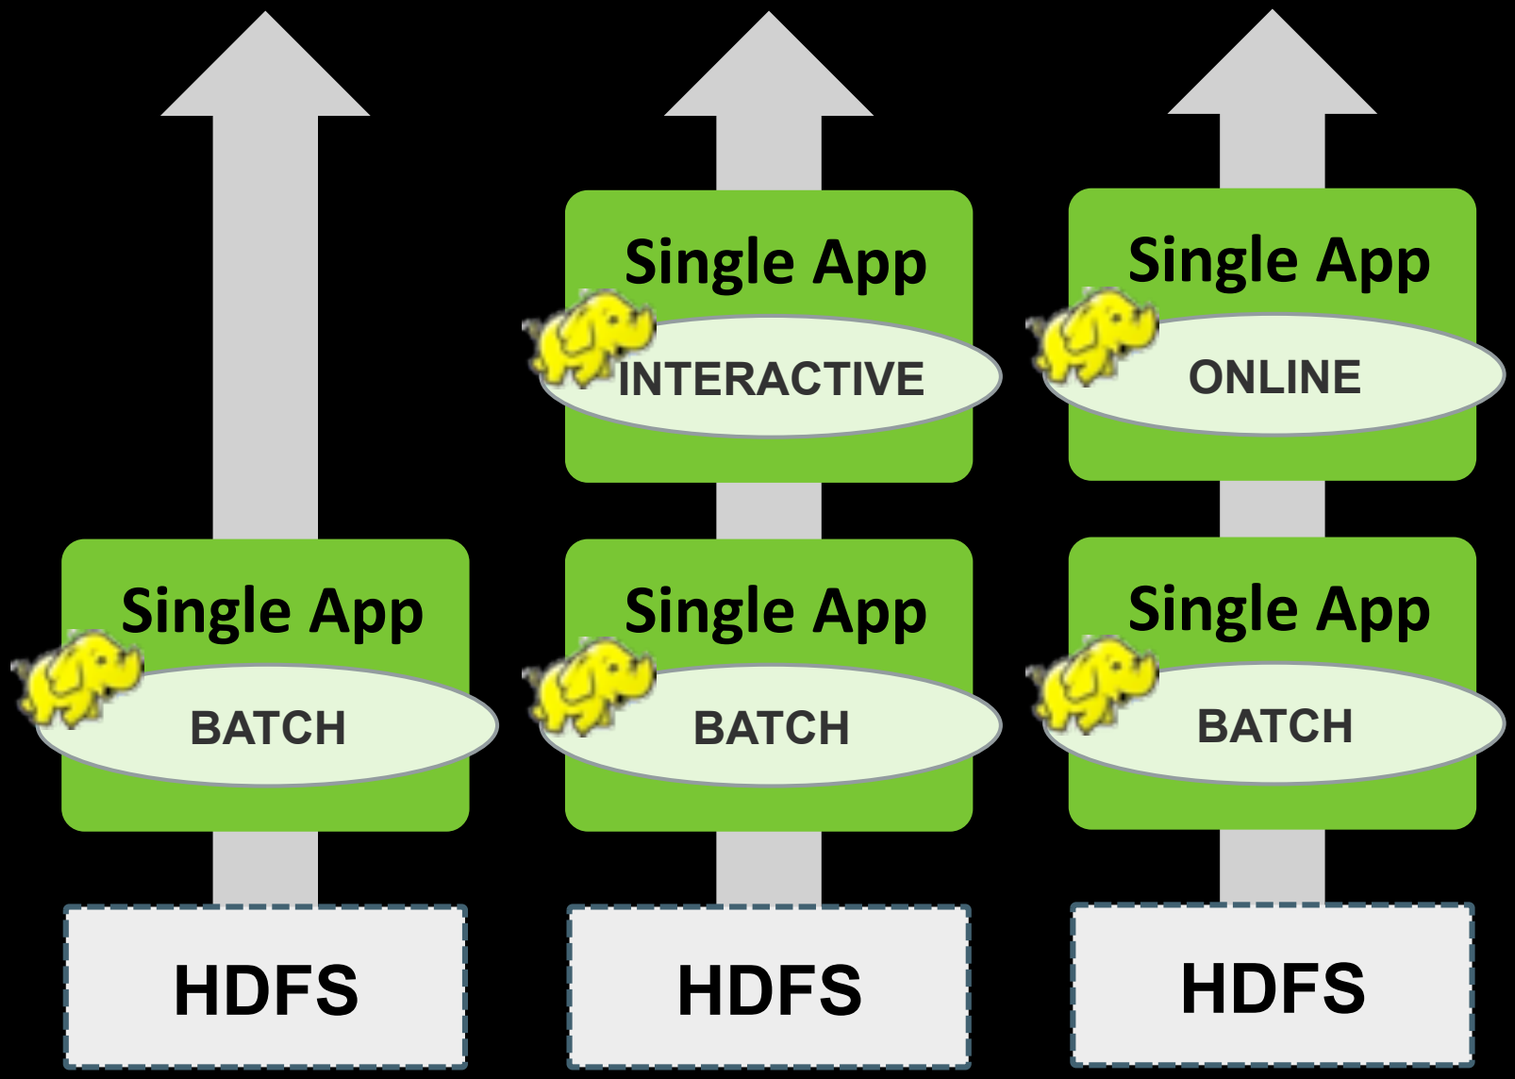

Hadoop 1.0

(Image Courtesy Arun Murthy, Hortonworks)

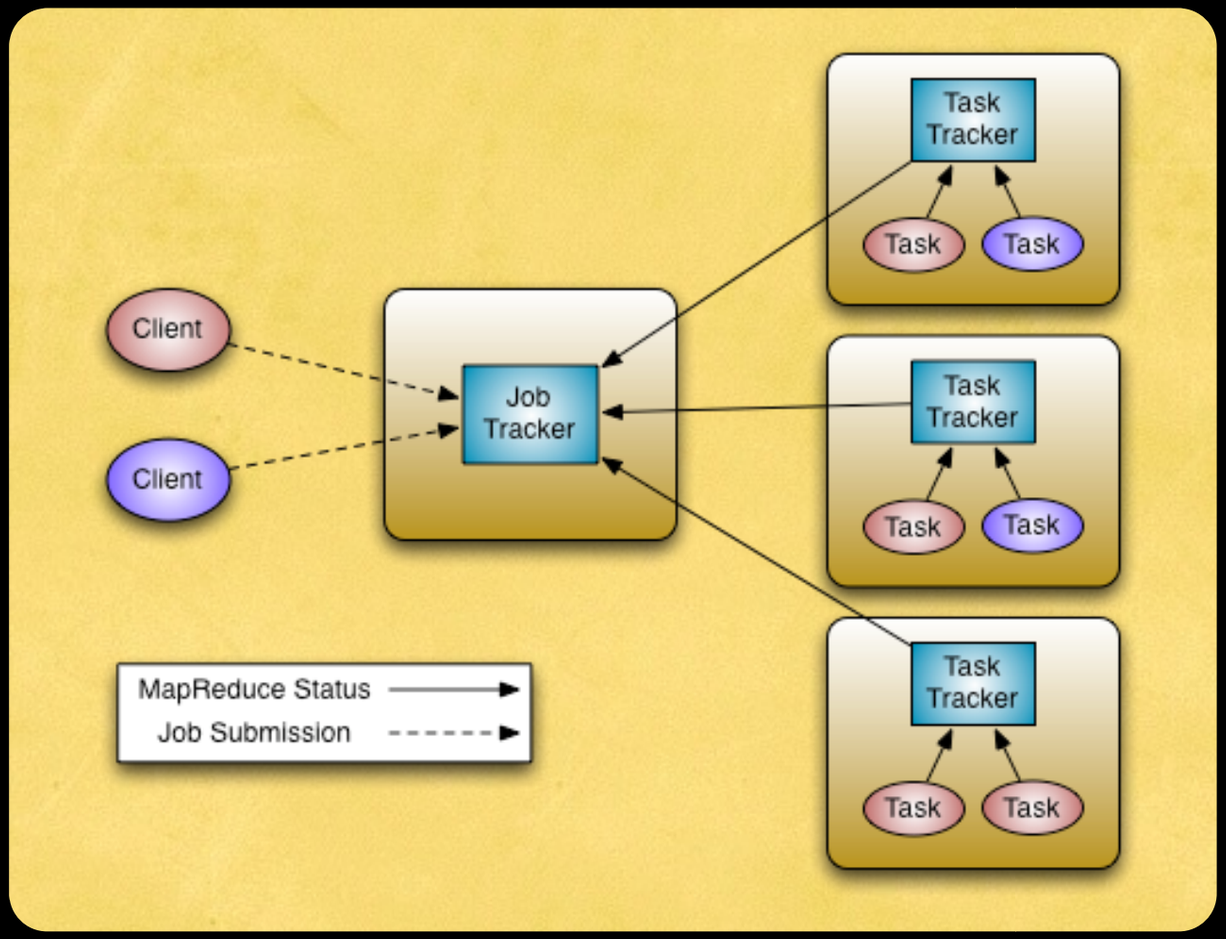

MapReduce 1.0

(Image Courtesy Arun Murthy, Hortonworks)

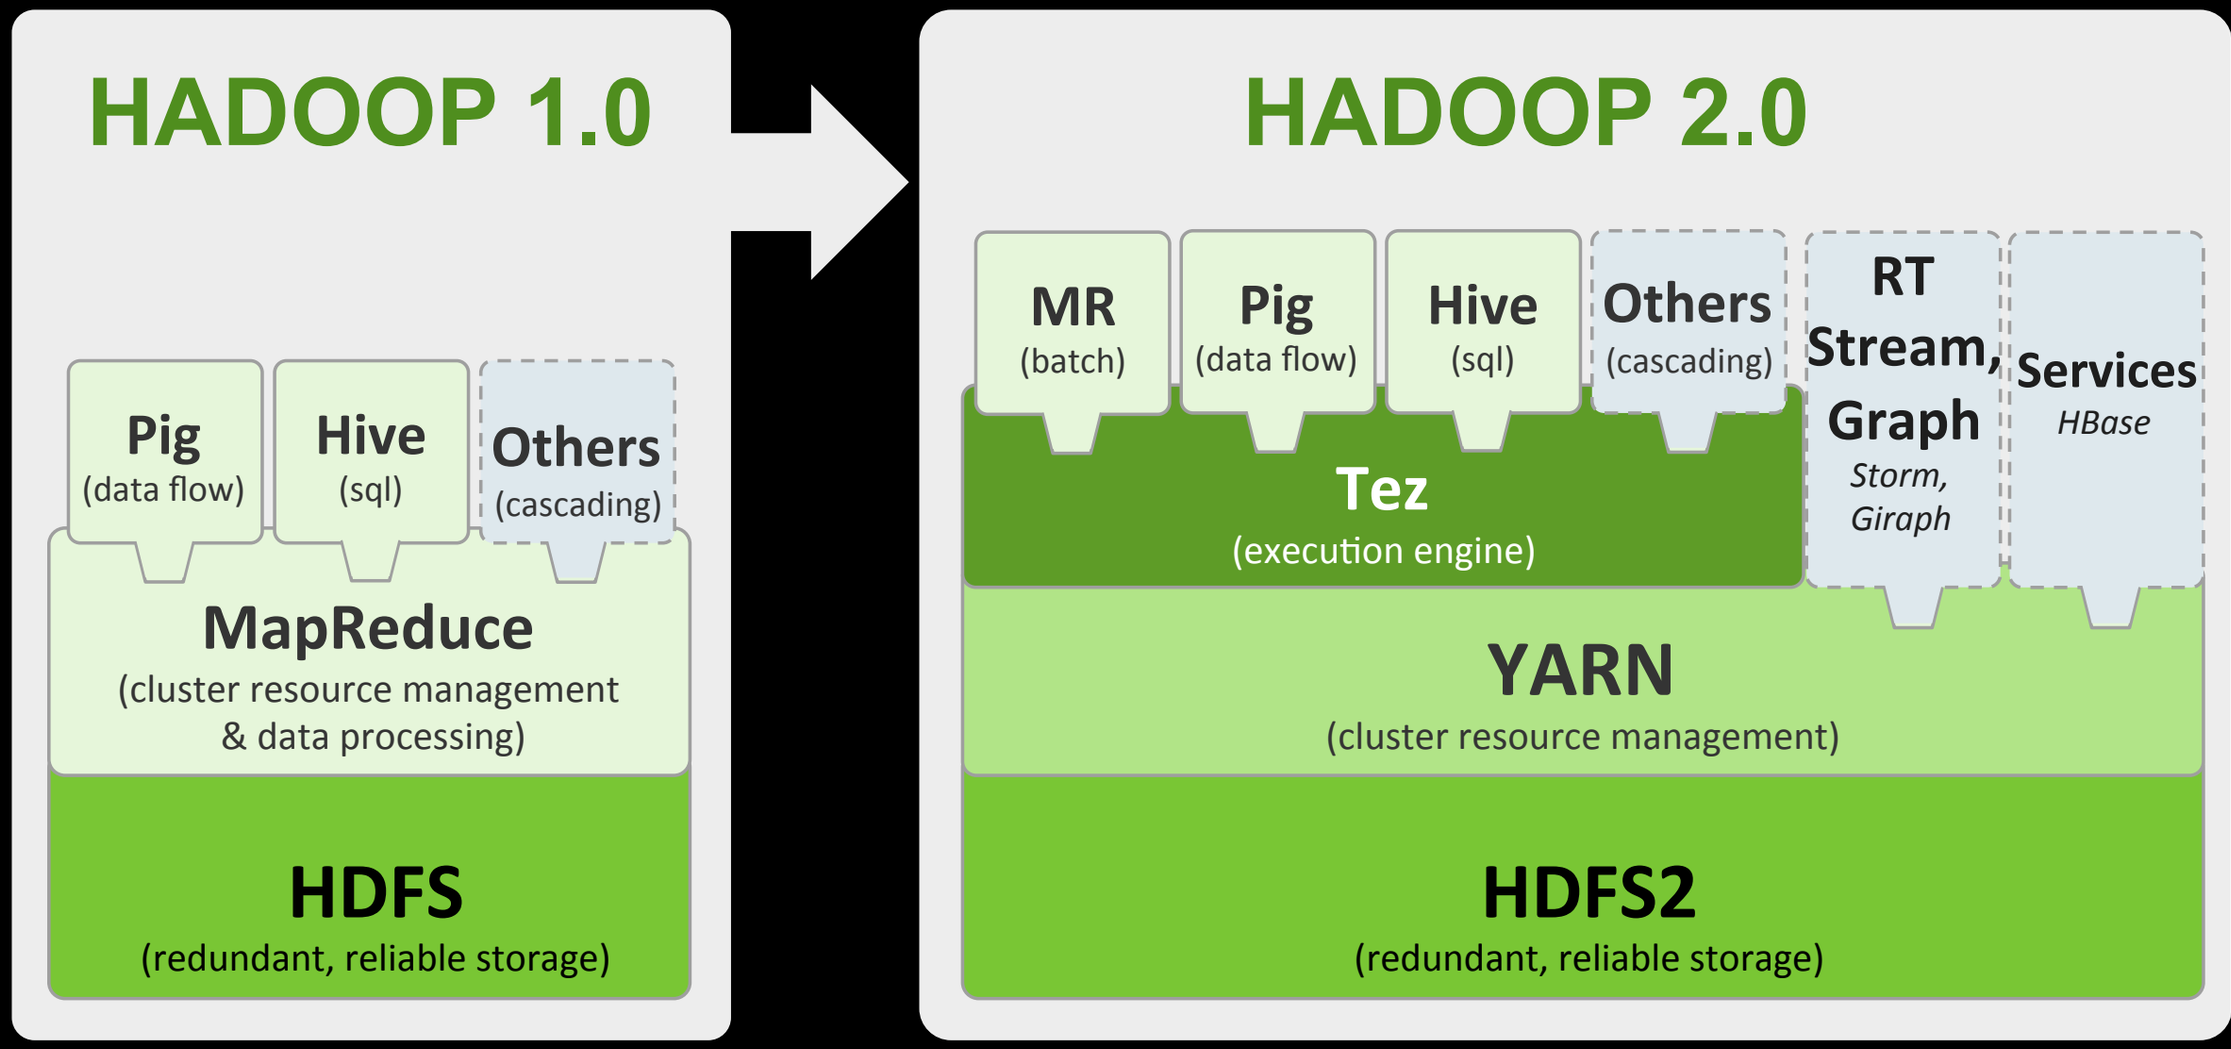

Hadoop 2.0

(Image Courtesy Arun Murthy, Hortonworks)

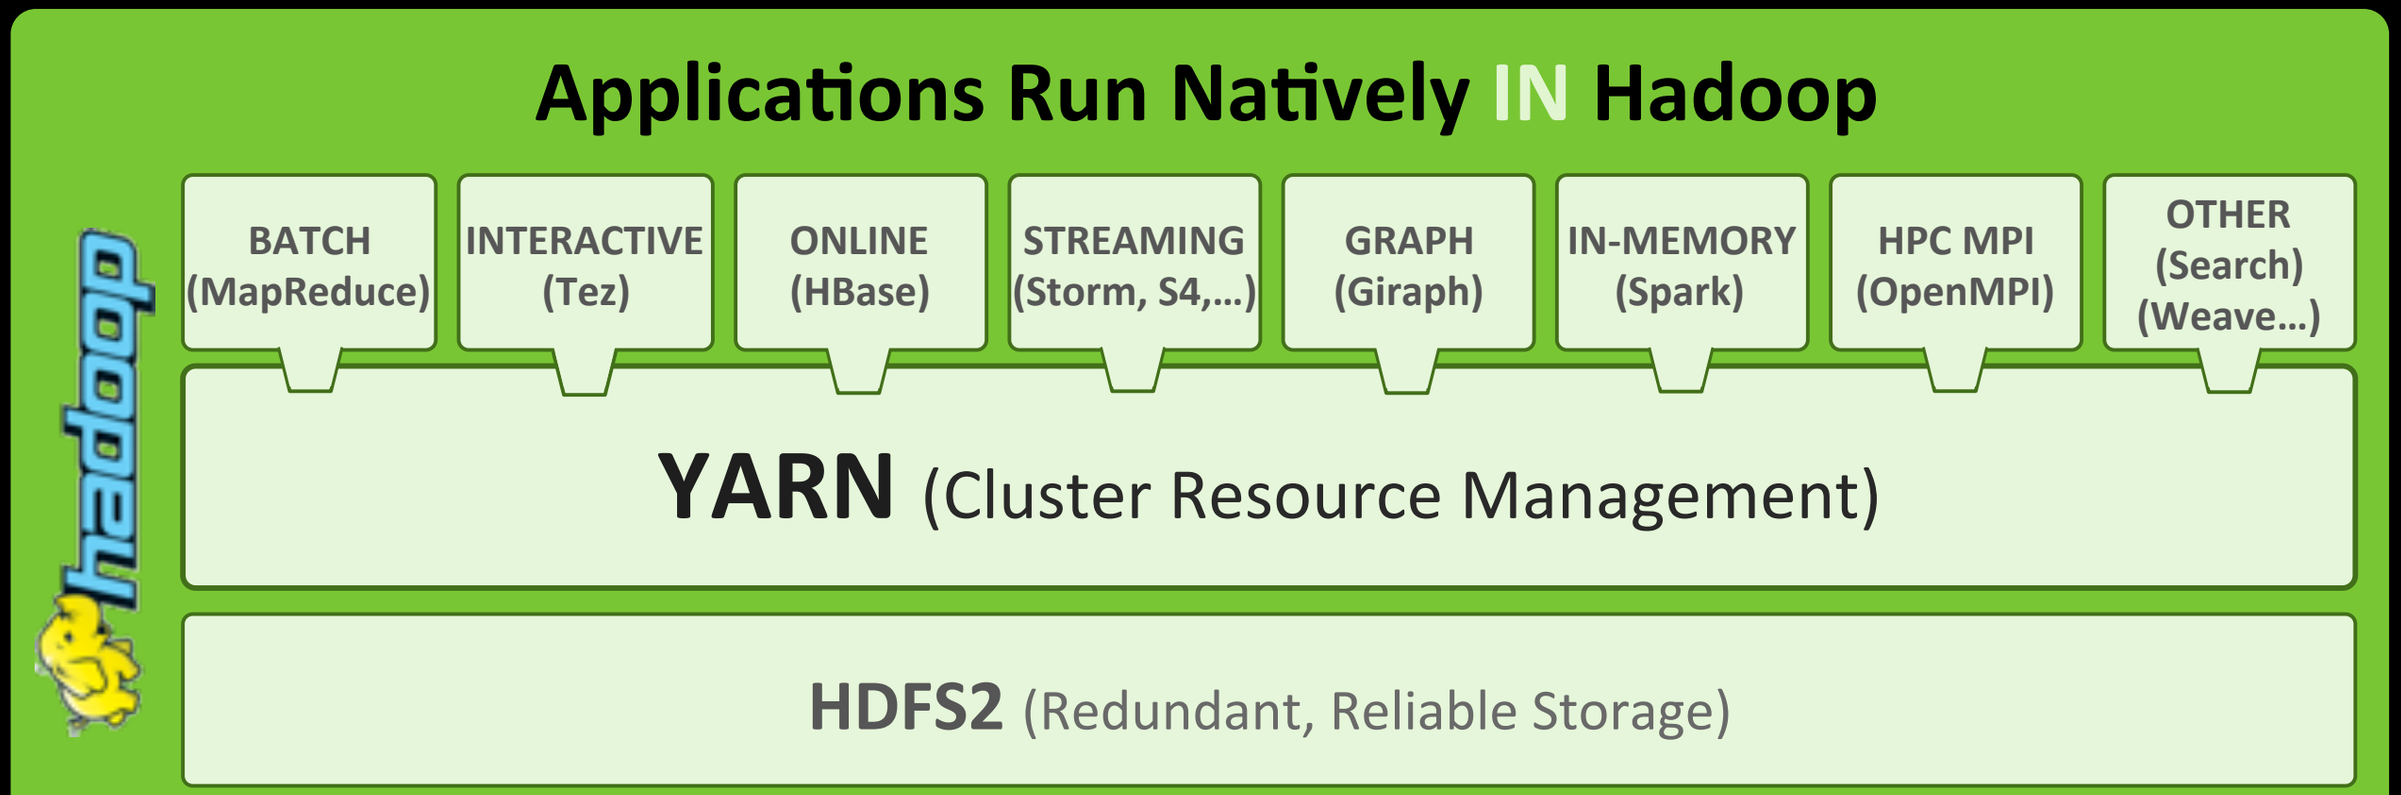

YARN Platform

(Image Courtesy Arun Murthy, Hortonworks)

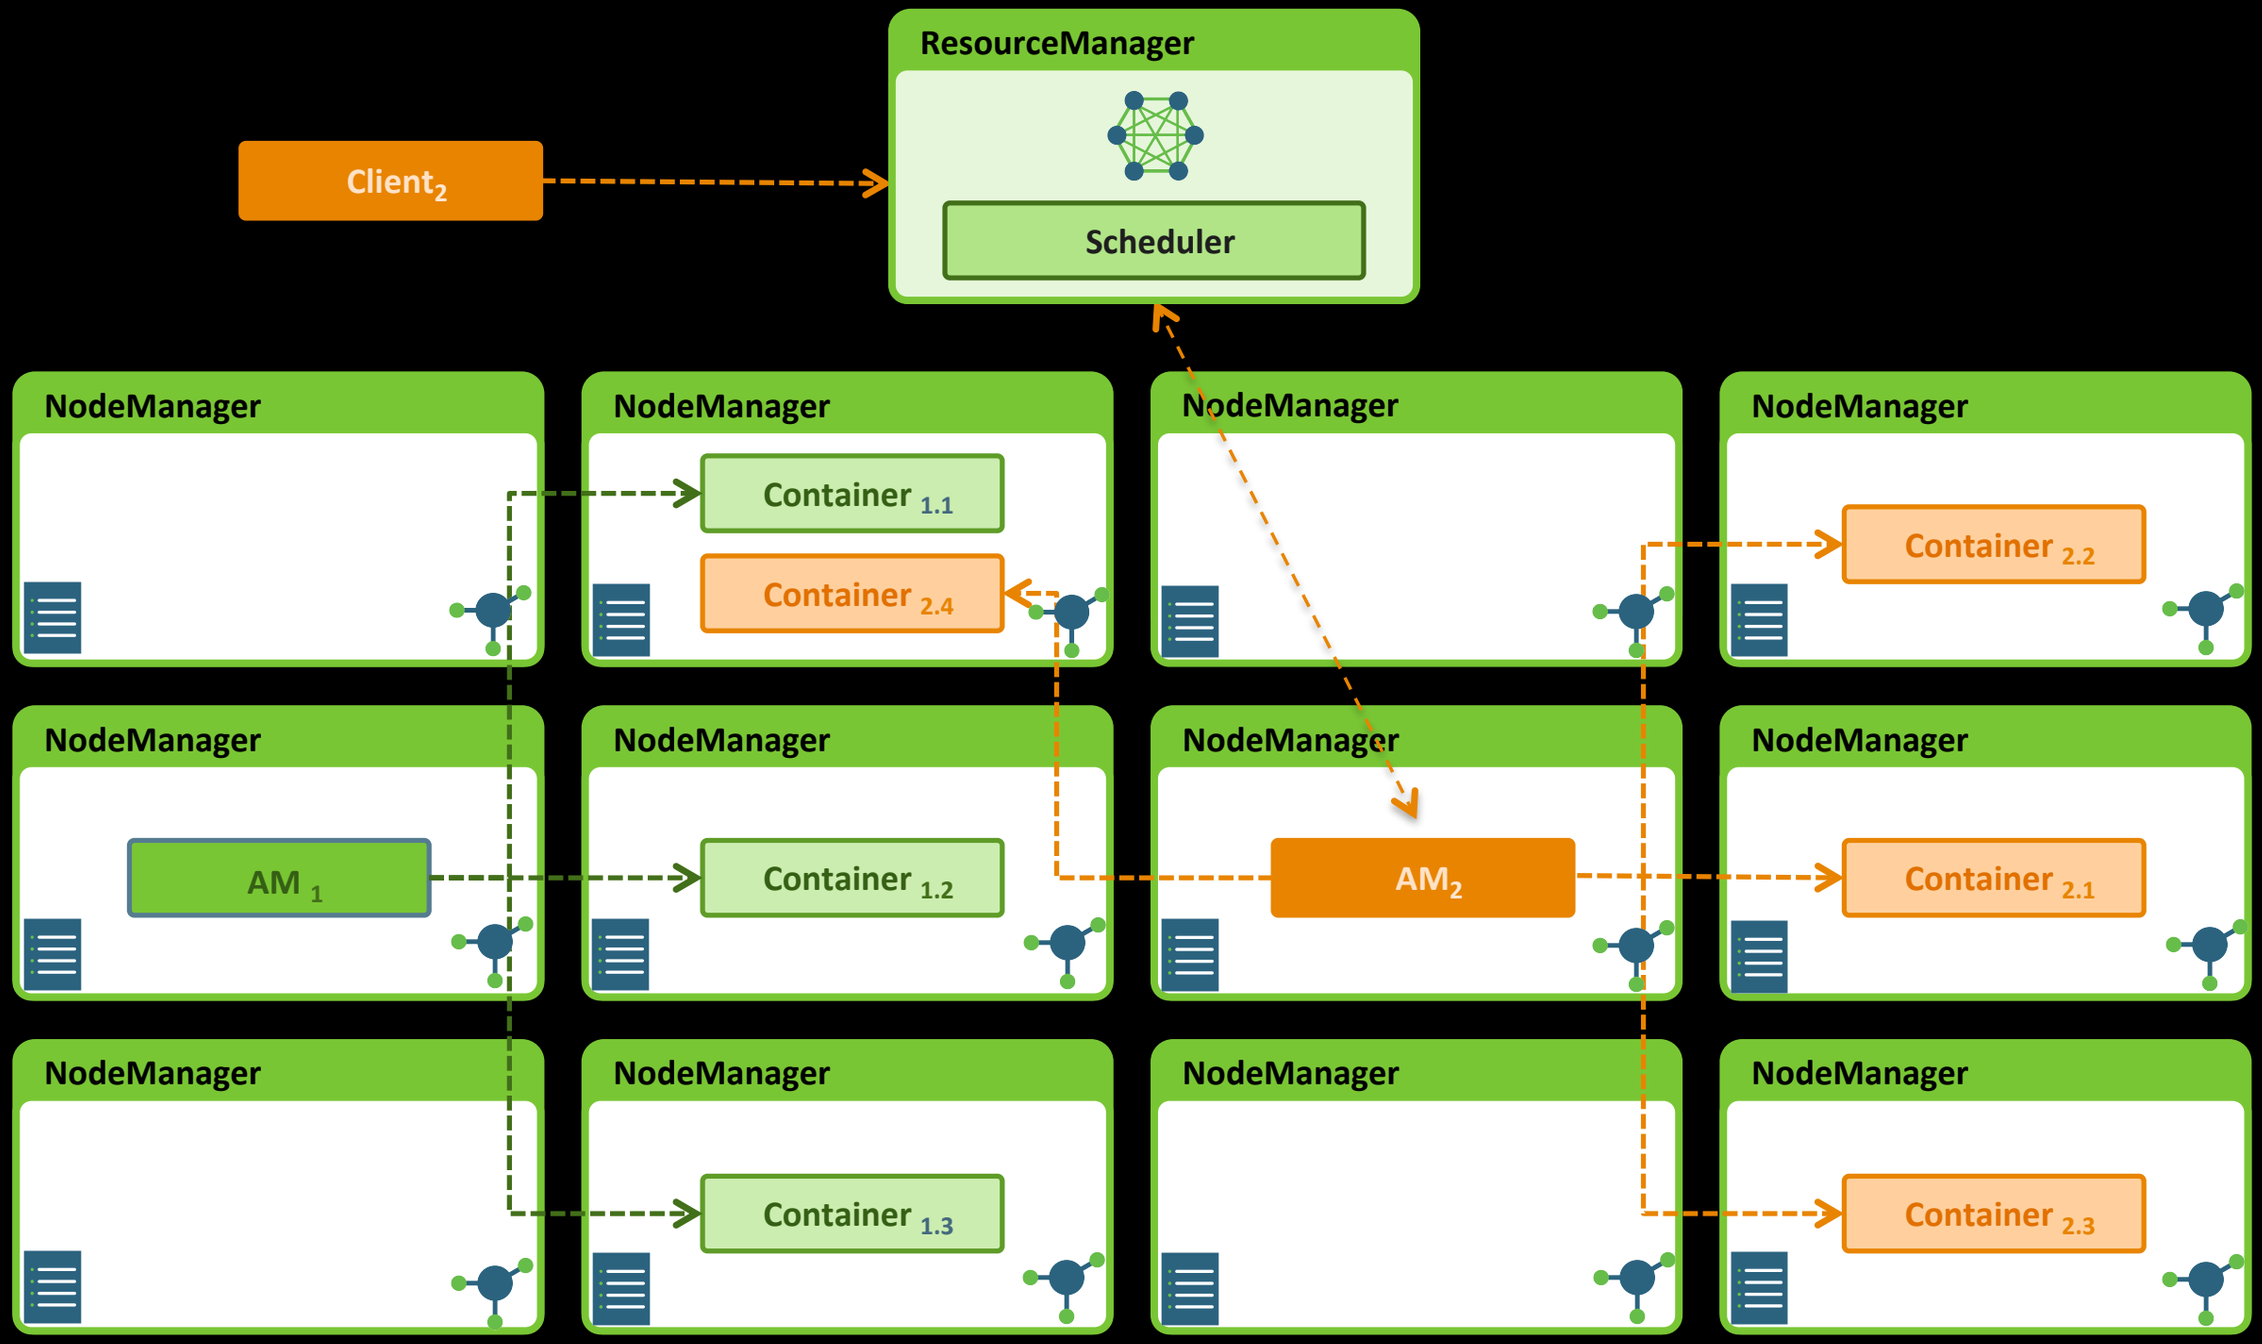

YARN Architecture

(Image Courtesy Arun Murthy, Hortonworks)

YARN

- Yet Another Resource Negotiator
- Resource Manager
- Node Managers
- Application Masters
  - Specific to paradigm, e.g. MR Application master (aka JobTracker)

Beyond MapReduce

- Apache Giraph - BSP & Graph Processing
- Storm on Yarn - Streaming Computation
- HOYA - HBase on Yarn
- Hamster - MPI on Hadoop
- More to come ...

Hamster

- Hadoop and MPI on the same cluster
- OpenMPI Runtime on Hadoop YARN
- Hadoop Provides: Resource Scheduling, Process monitoring, Distributed File System
- Open MPI Provides: Process launching, Communication, I/O forwarding

Hamster Components

- Hamster Application Master
- Gang Scheduler, YARN Application Preemption
- Resource Isolation (lxc Containers)
- ORTE: Hamster Runtime
  - Process launching, Wireup, Interconnect

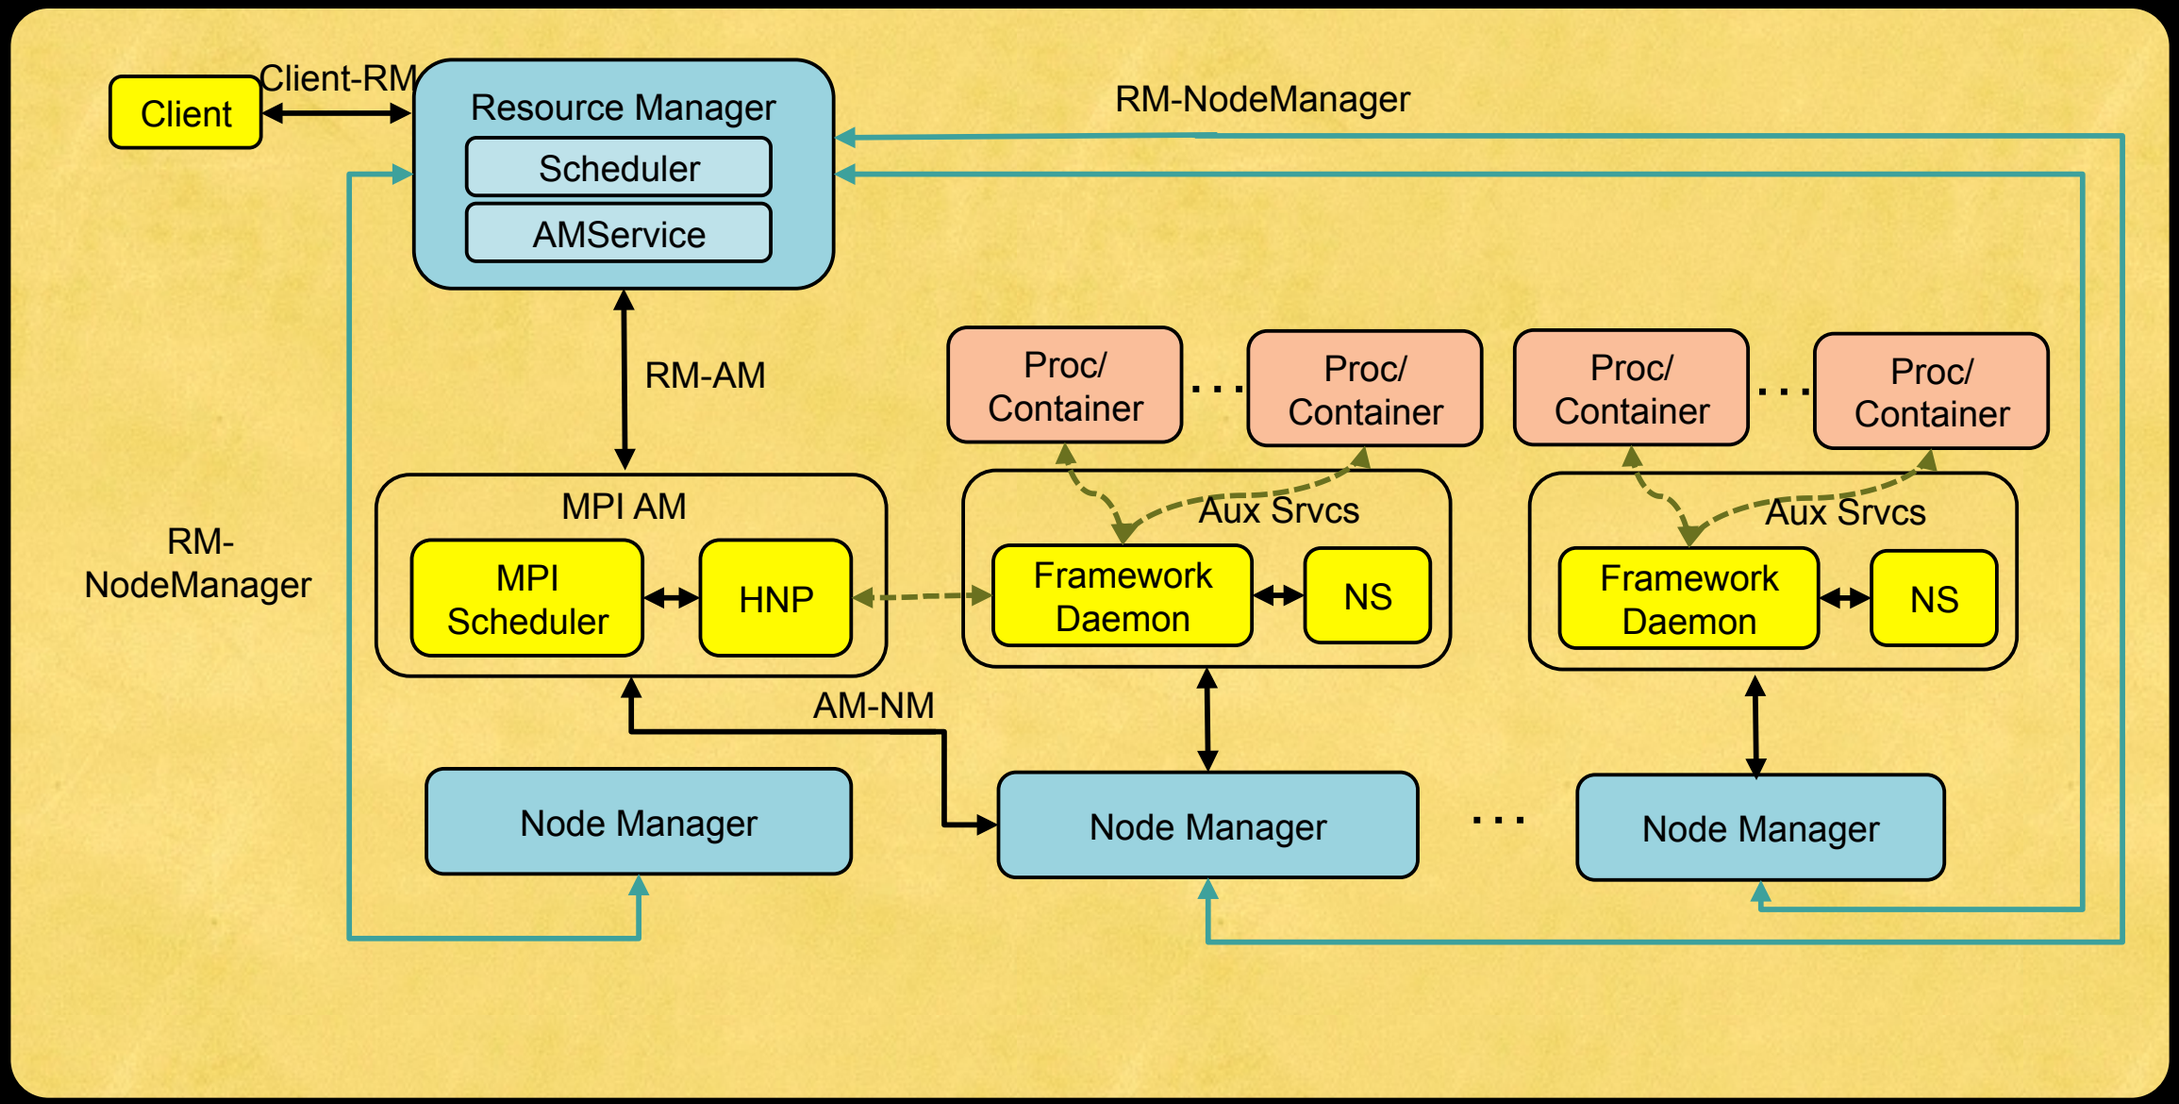

Hamster Architecture

Hamster Scalability

- Sufficient for small to medium HPC workloads
- Job launch time gated by YARN resource scheduler

	Launch	WireUp	Collectives	Monitor
OpenMPI	O(logN)	O(logN)	O(logN)	O(logN)
Hamster	O(N)	O(logN)	O(logN)	O(logN)

GraphLab

GraphLab + Hamster on Hadoop

About GraphLab

- Graph-based, High-Performance distributed computation framework
- Started by Prof. Carlos Guestrin in CMU in 2009
- Recently founded Graphlab Inc to commercialize Graphlab.org

GraphLab Features

- Topic Modeling (e.g. LDA)
- Graph Analytics (Pagerank, Triangle counting)
- Clustering (K-Means)
- Collaborative Filtering
- Linear Solvers
- etc...

Only Graphs are not Enough

- Full Data processing workflow required ETL/ Postprocessing, Visualization, Data Wrangling, Serving
- MapReduce excels at data wrangling
- OLTP/NoSQL Row-Based stores excel at Serving
- GraphLab should co-exist with other Hadoop frameworks

Call To Action

Prepare for Convergence

- HPC: Cache Coherence, Prefetching, Zero-copy, Low-contention locks
- "Big Data": Caching, Mirroring, Sharding (various flavors), relaxed consistency
- Databases: Indexing, MVCC, Columnar storage/processing, Cost-based optimization

Convergence

- Resource Allocation, Scheduling, Lifecycle Management
- Compute, Storage, and Communication isolation, Multi-tenancy, Performance SLAs
- Auth & Auth, Data/System Provisioning and Management, Monitoring, Metadata Management, Metering

New Hardware Platforms

- Mellanox - Hadoop Acceleration through Network-assisted Merge
- RoCE - Brocade, Cisco, Extreme, Arista...
- ARM - Low power Hadoop servers
- SSD - Velobit, Violin, FusionIO, Samsung..
- Niche - Compression, Encryption...

Data Cloud of Future?

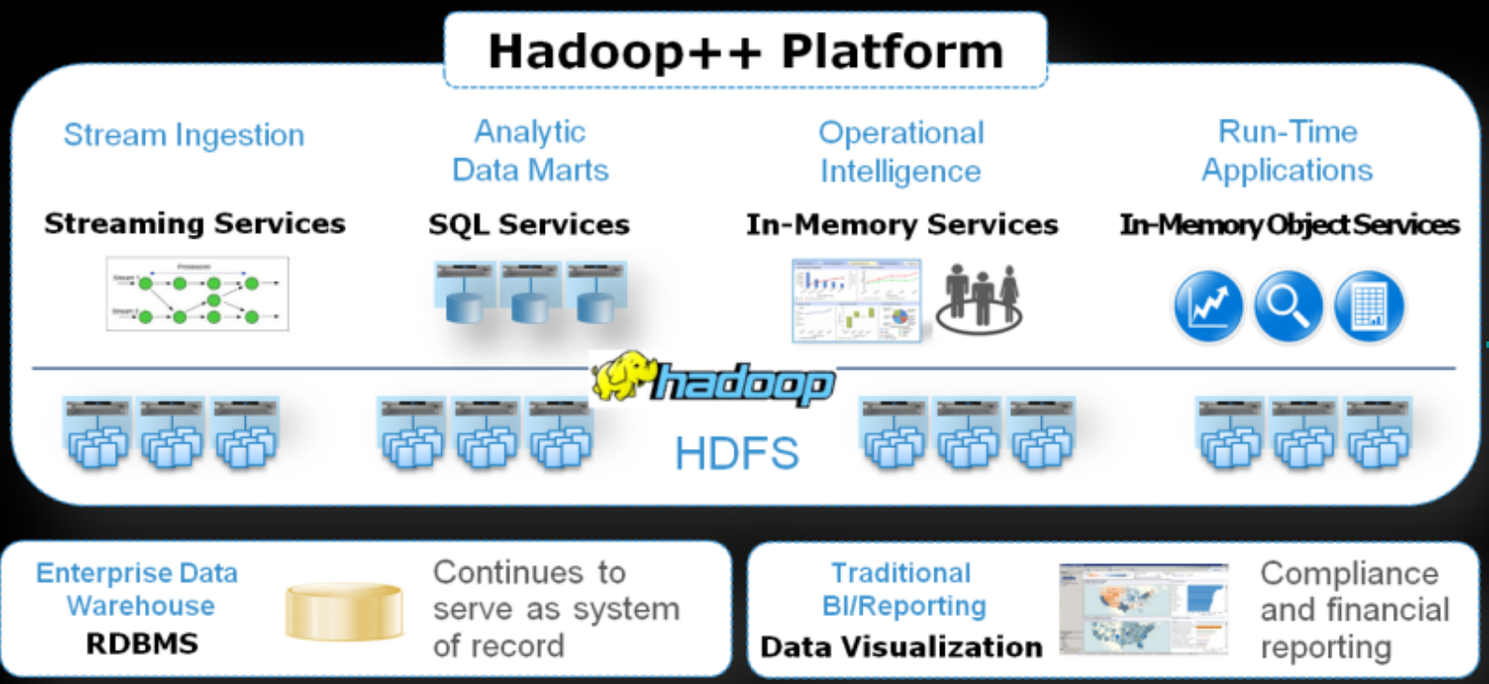

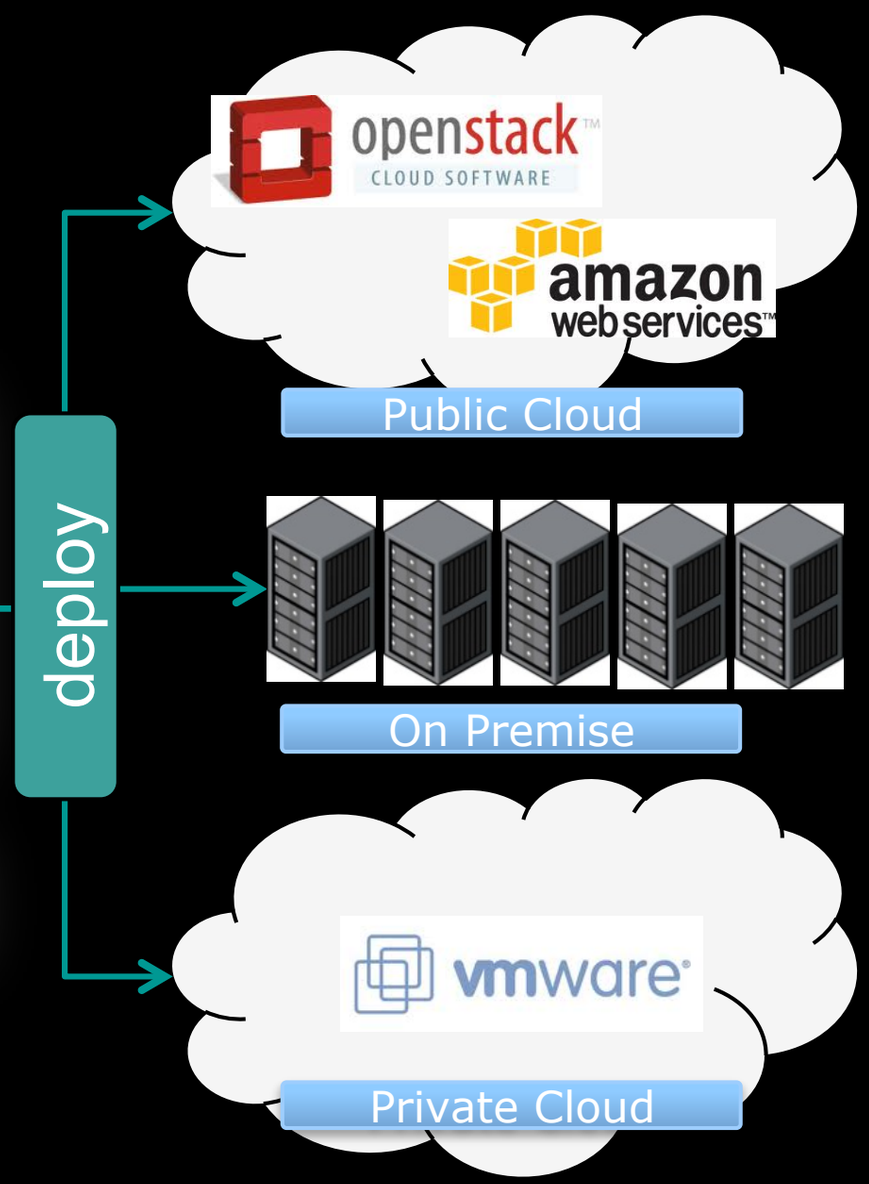

Questions?