Search Index Building with HBase at eBay

Ming Ma

12/02/2011

Agenda

- About eBay Search Engine
- Intro to HBase
- Why HBase
- System Design
- HBase Enhancements
- Future Works

Search functionality at eBay

seller

Post new items

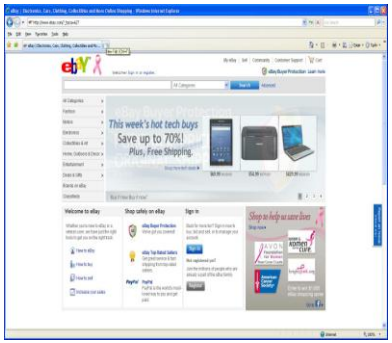

Keyword-based Search

buyer

- It requires a great search engine
  - Fast, scalable, reliable
  - Result relevance
  - Search Experience

Batched Oriented Search Engine

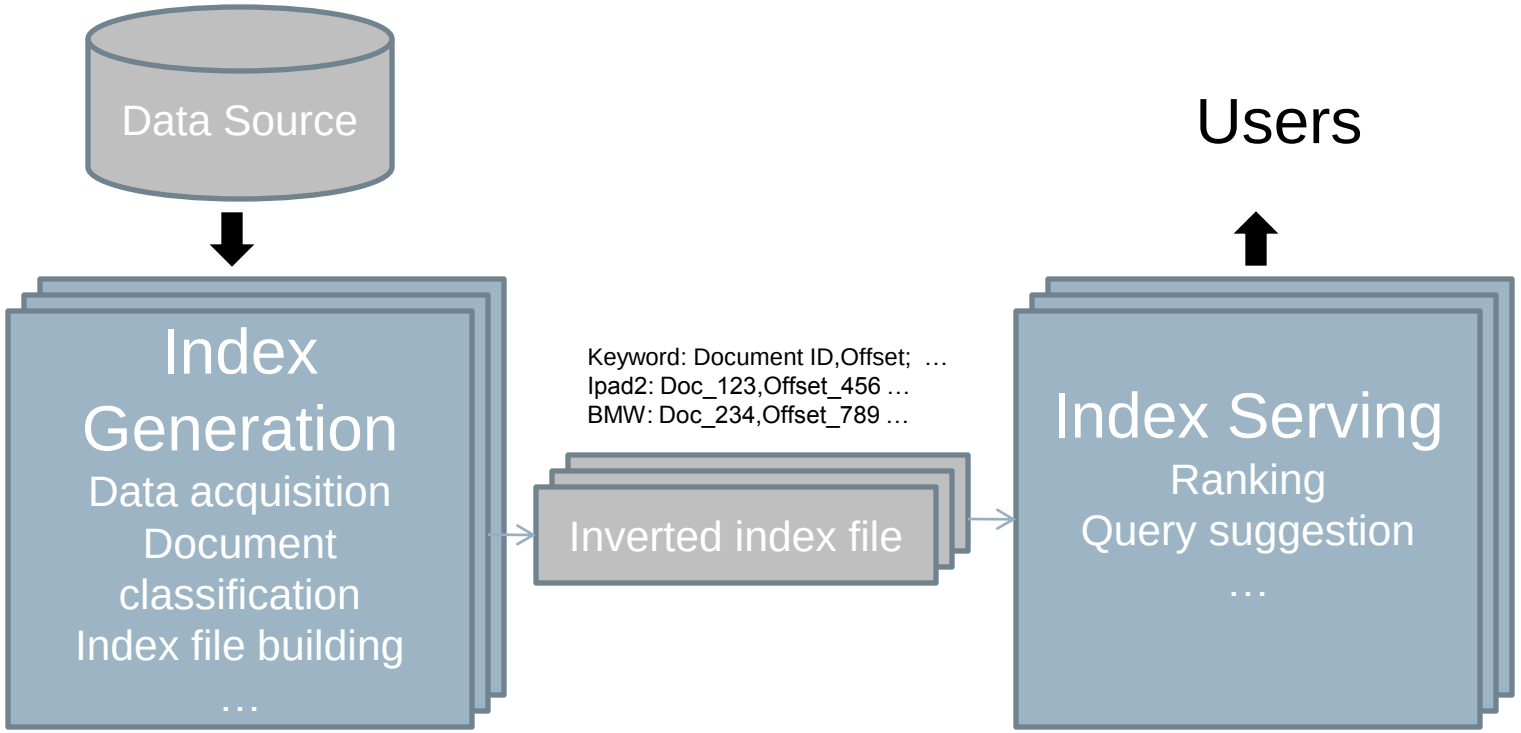

Use Hadoop stack for index building

Hadoop Open Source Stack

- HBase -> Database Storage Engine
- HDFS -> Distributed File System
- MapReduce -> Framework for writing scalable data applications
- Zookeeper -> distributed coordination service
- Oozie -> workflow engine to manage MapReduce jobs

Agenda

- About eBay Search Engine
- Intro to HBase
- Why HBase
- System Design
- HBase Enhancements
- Future Works

What is HBase?

- Column-oriented distributed data store
- Modeled after google's BigTable
- It is open source

HBase Data Model

- HBase tables have rows with primary keys
- Each row could have any number of columns
- Columns are grouped into column families
- Each cell has multiple versions
- Rows are stored in sorted order

Example

Item Id	Partial:BidPrice			Main:Title	Main:Category	…
0	$480	$460	…	Ipad2…	Electronics	
1				BMW…	Motors	
…						

Tables and Regions

- Region
  - A subset of a table's rows, range defined as [startKey, endKey)
  - Dynamically split
  - Hosted by Region Server
- Two special tables used by HBase, "-ROOT-", ".META."
  - Store regions locations of user tables

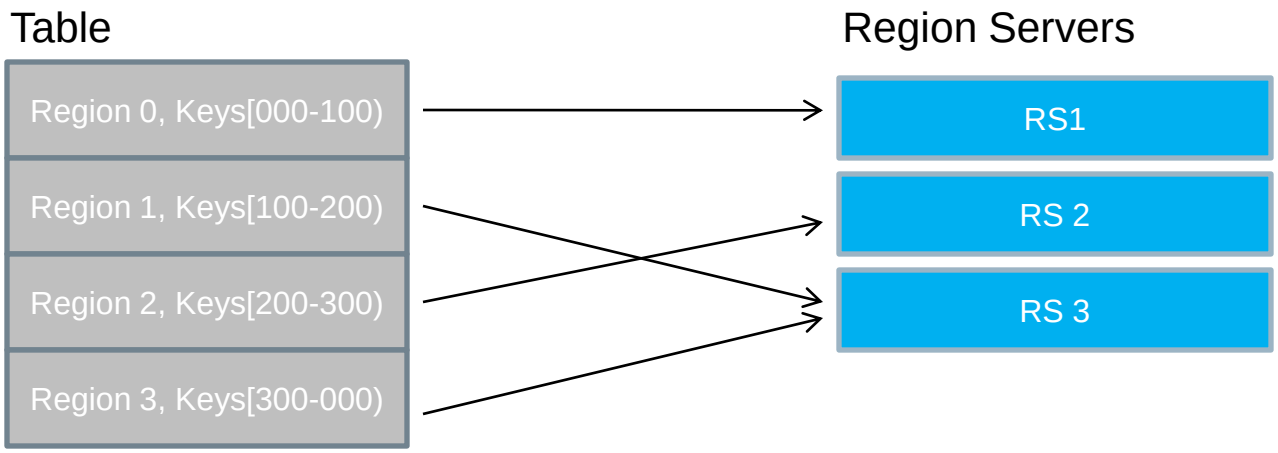

HBase System Overview

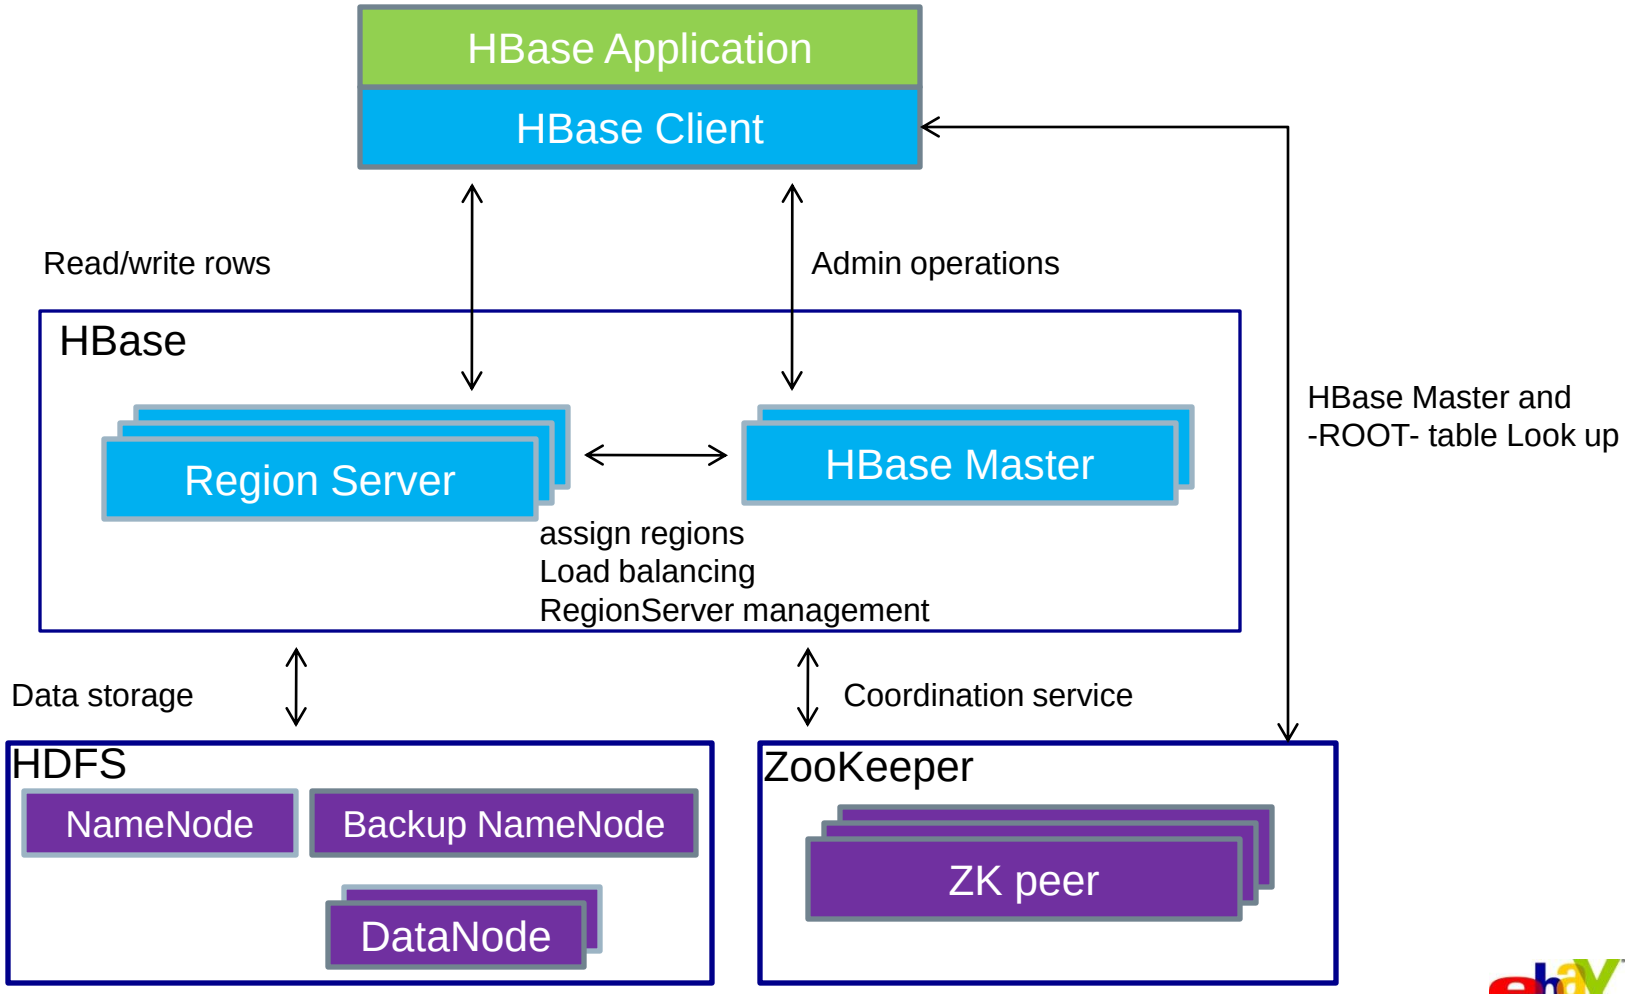

HBase and MapReduce

- HBase MapReduce Adaptors for Input, Output
  - Scenarios: data processing, data import to HBase
- Region Server, Data Node, Task Tracker are collocated.
  - Data locality

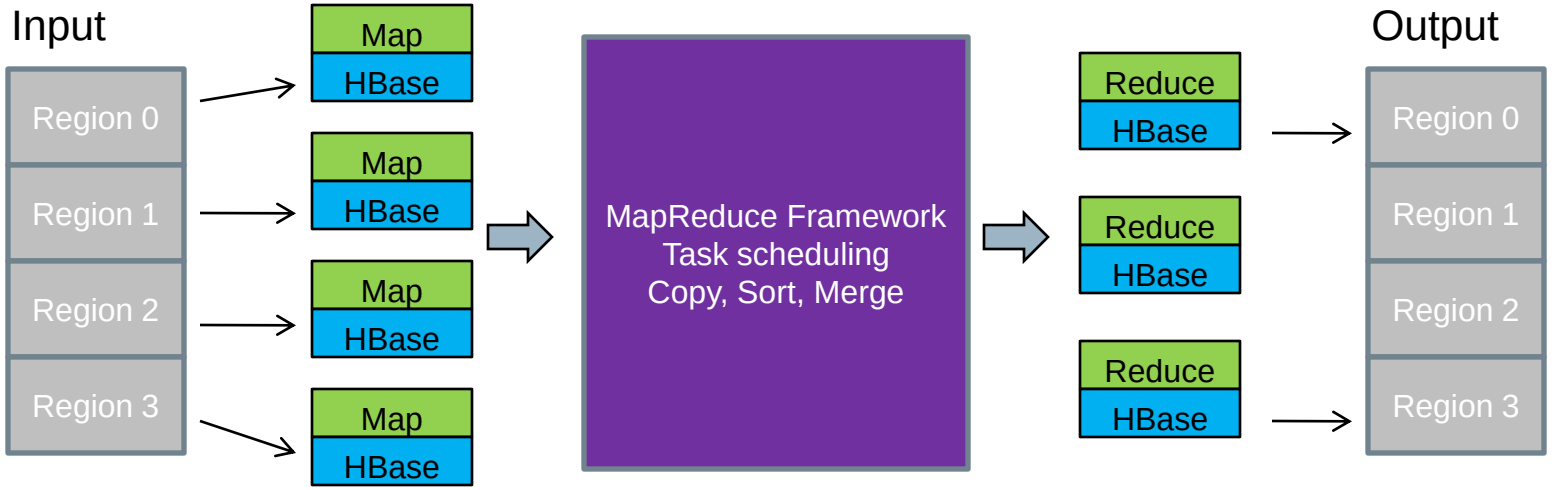

Agenda

- About eBay Search Engine
- Intro to HBase
- Why HBase
- System Design
- HBase Enhancements
- Future Works

Technical Aspects

- HBase
  - Horizontal scalability
  - Automatic failover
  - Strong consistency
  - Efficient at random reads/writes
- Layered over HDFS
  - Get all benefits of scalable, reliable storage
- Integration with Map Reduce framework

Other Considerations

- HDFS, Map Reduce well supported inside eBay
- Great community support
- Evaluation of open source search engines

Agenda

- About eBay Search Engine
- Intro to HBase
- Why HBase
- System Design
- HBase Enhancements
- Future Works

Indexing Building System Overview

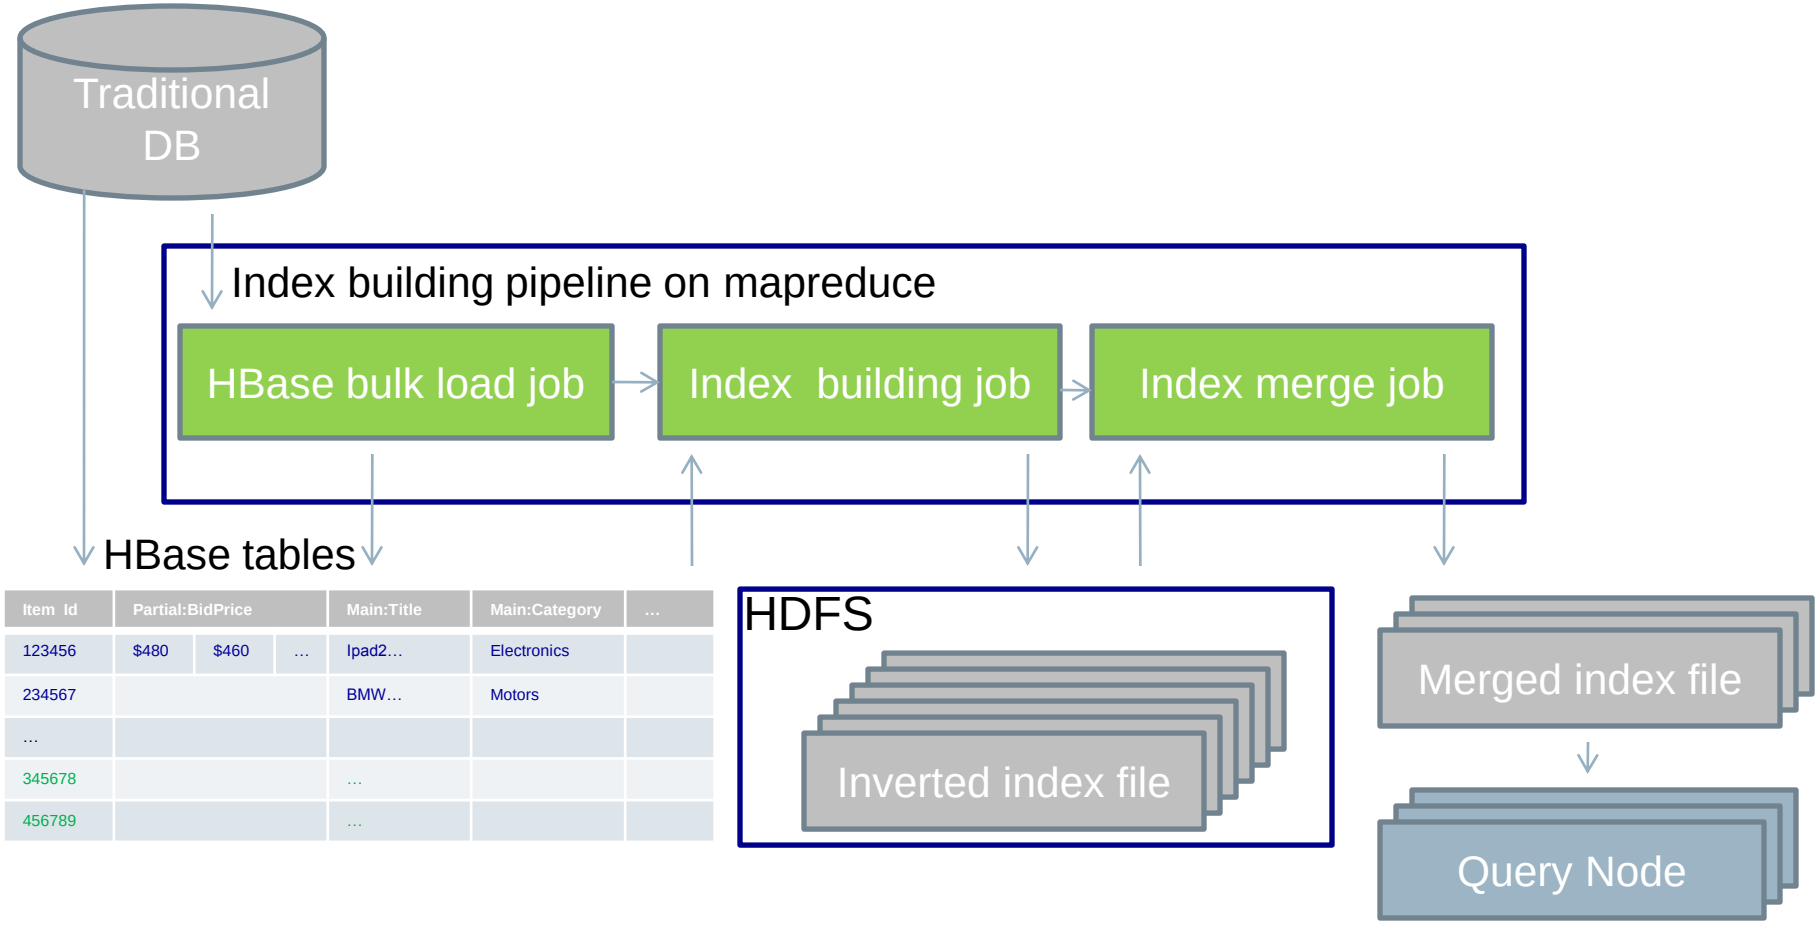

Item Id	Partial:BidPrice			Main:Title	Main:Category	…
123456	$480	$460	…	Ipad2…	Electronics	
234567				BMW…	Motors	
…						
345678				…		
456789				…		

Other Design Considerations

- Use small number of column families
  - Better performance
- Store multiple versions with custom timestamp
  - Deterministic index building
- Time-range scan
  - Near-Real-Time index support
- Row key is 64bit based on item id
  - Uniformed distribution of rows

Agenda

- About eBay Search Engine
- Intro to HBase
- Why HBase
- System Design
- HBase Enhancements
- Future Works

HBase Enhancements

- Metrics
  - HBase Client-side metrics
  - HDFS block data locality
- Reliability
  - “Regions-in-transition” issues
  - Table operations
- Performance
  - Table level load balancing
  - Handling of small map reduce jobs
- Features
  - Multi-column-family bulk load

HBase HFile Data Locality

- HBase table content is stored as HFile in HDFS
- MapReduce Jobs that scan lots of data
  - HFile's HDFS data locality become an issue
  - Solution: balance regions to achieve data locality

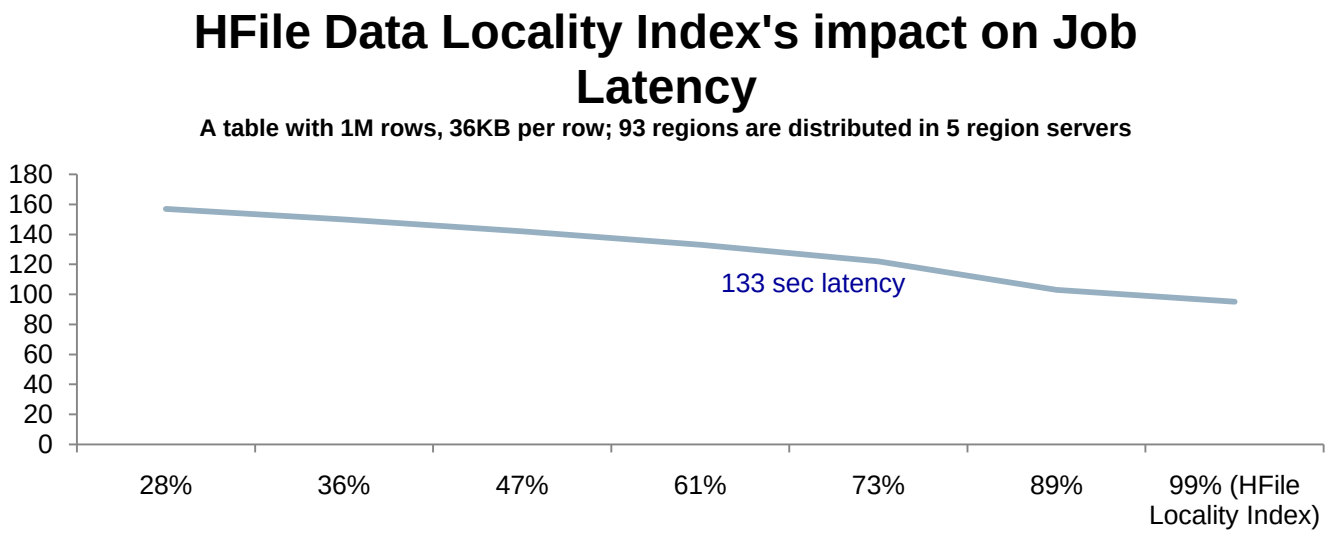

Agenda

- About eBay Search Engine
- Intro to HBase
- Why HBase
- System Design
- HBase Enhancements
- Future Works

Improve Search Index Building

- System Reliability
  - HBase, HDFS, MapReduce
- System Availability
  - HDFS Namenode HA
  - HBase region high availability
- Performance
  - Real time index with coprocessors
  - HDFS efficient merge
  - MapReduce scheduling optimization

Support Diverse Work Loads

- Scalability
  - # of regions, machines
- Shared service
  - Multi-tenancy within HBase
  - Multi-tenancy with HDFS, MapReduce

Thanks! Questions?

http://www.ebaycareers.com/Home.html