Karlsruhe Institute of Technology

Cloud Computing: MapReduce and Hadoop

June 2010

Marcel Kunze, Research Group Cloud Computing

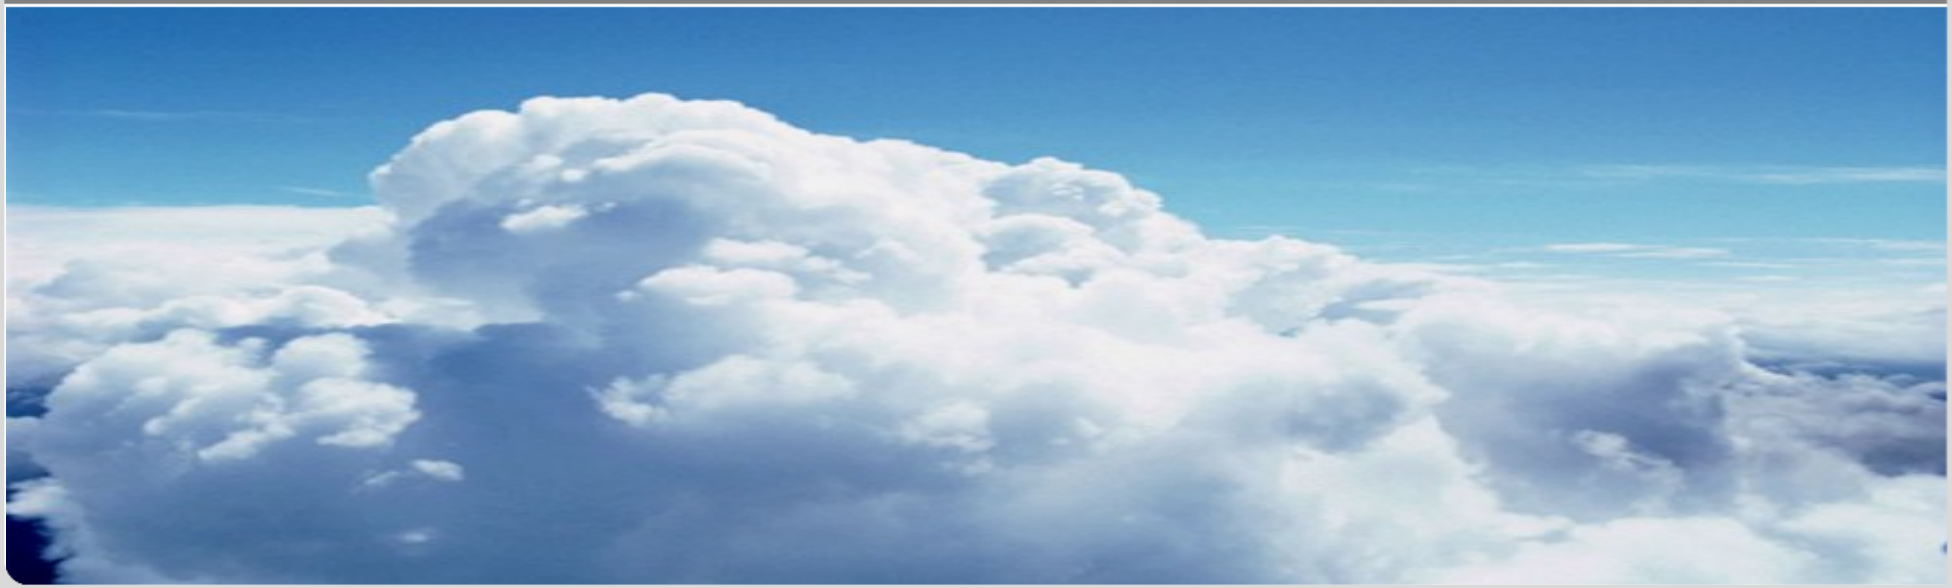

Motivation: Large Scale Data Processing

- MapReduce: Algorithm for large scale parallel data processing
  - Want to process lots of data (> 1 TB)
  - Want to parallelize across hundreds/thousands of CPUs
  - … Want to make this easy
- Potential fields of application
  - Web indexing
  - Data mining
  - Log file analysis
  - Machine learning
  - Scientific simulation
  - Bioinformatics research
- Google uses MapReduce everywhere, e.g. to run PageRank

Google Search: PageRank Algorithm

- Find the most popular page for a specific set of key words
- If a user starts at a random web page and surfs by clicking links and randomly entering new URLs, what is the probability that s/he will arrive at a given page?
- The PageRank of a page captures this notion
  - More “popular” or “worthwhile” pages get a higher rank

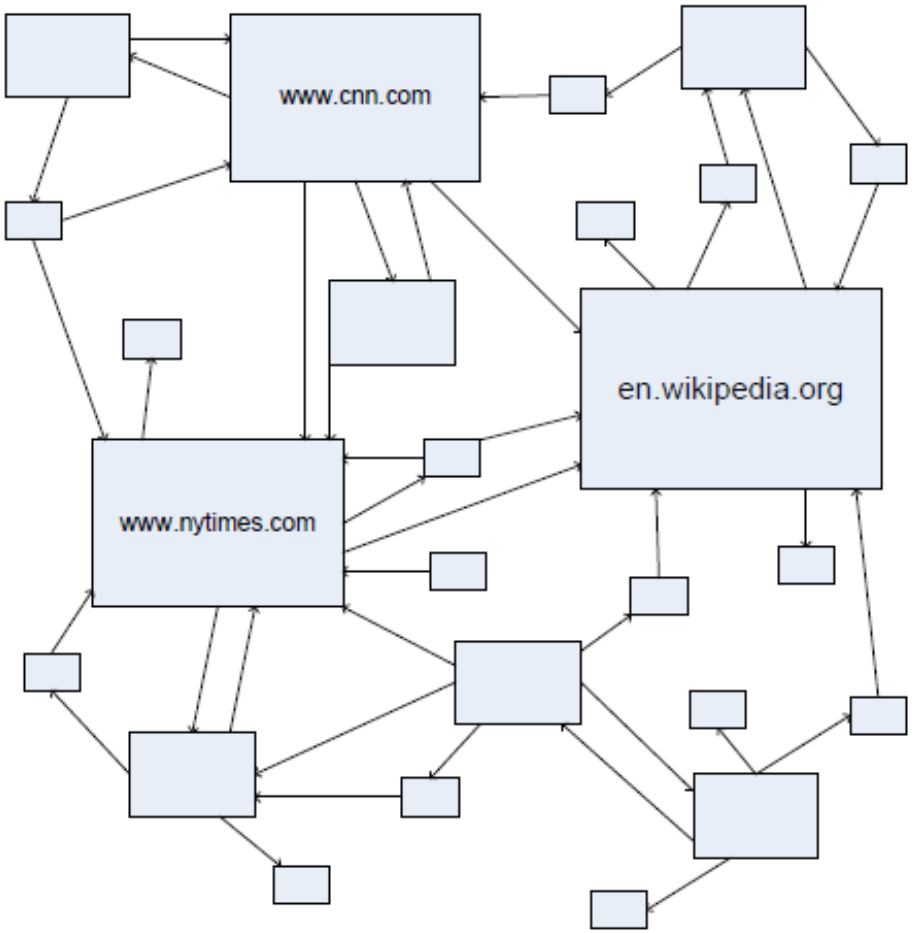

Source: A.Kimball

PageRank: Formula

- Given page A, and pages T_1 through T_n linking to A, PageRank is defined as:

$$PR(A) = (1-d) + d\,(PR(T_1)/C(T_1) + \ldots + PR(T_n)/C(T_n))$$

C(P) is the cardinality (out-degree) of page P

d is the damping (“random URL”) factor

Source: A.Kimball

PageRank: Intuition

- Calculation is iterative: PR_{i+1} is based on PR_i
- Each page distributes its PR_i to all pages it links to. Linkees add up their awarded rank fragments to find their PR_{i+1}
- d is a tunable parameter (usually = 0.85) encapsulating the "random jump factor"

$$PR(A) = (1-d) + d\ (PR(T_1)/C(T_1) + ... + PR(T_n)/C(T_n))$$

Source: A.Kimball

PageRank: Iteration

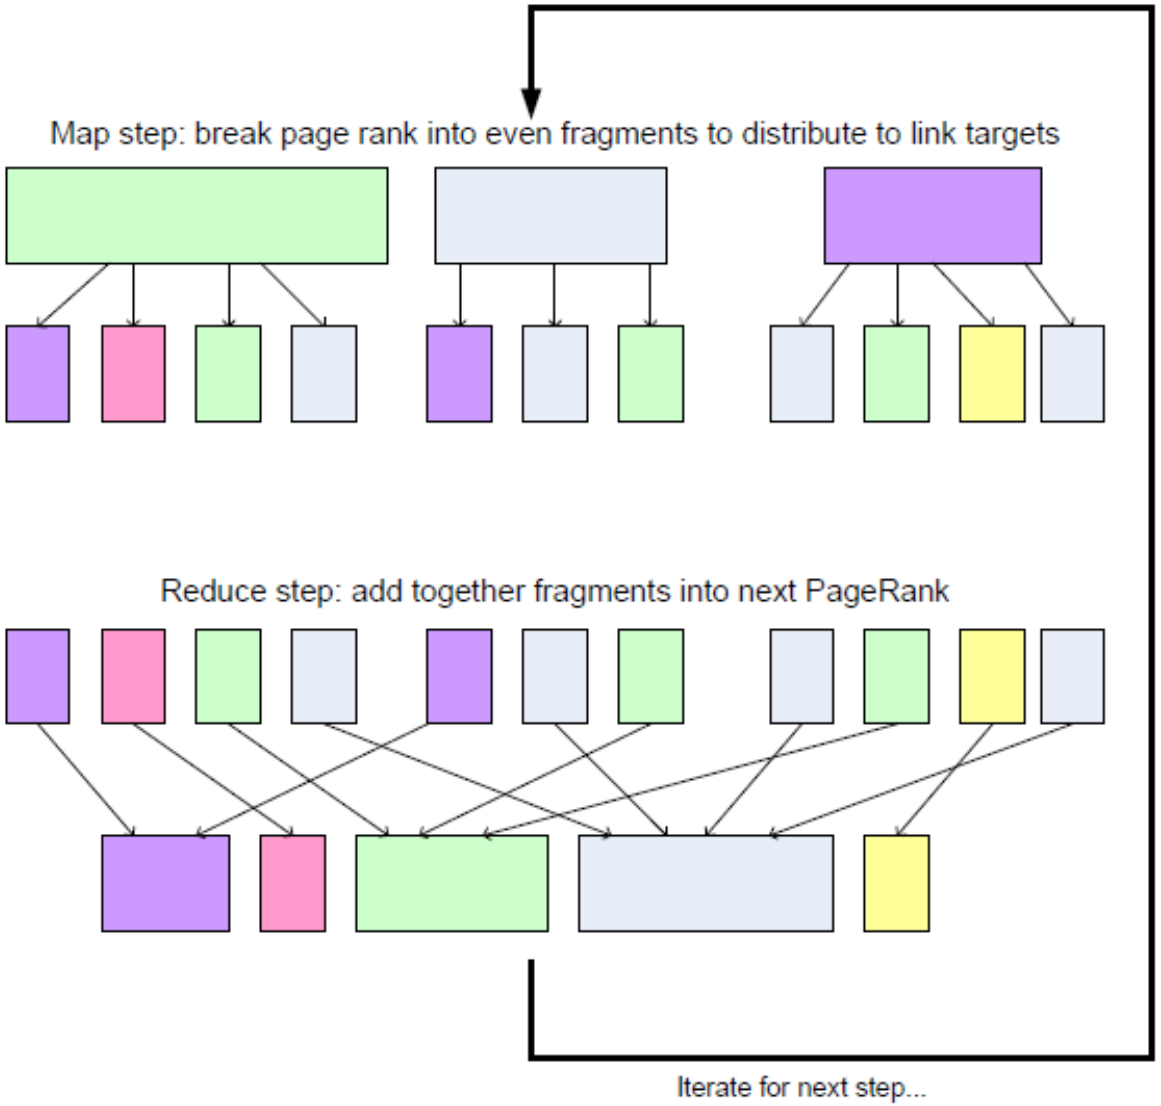

Source: A.Kimball

Google File System (GFS)

- A GFS cluster has one master and many chunkservers
- Files are divided into 64 MB chunks
- Chunks are replicated and stored in the Unix file systems of the chunkservers
- The master holds metadata
- Clients get metadata from the master, and data directly from chunkservers

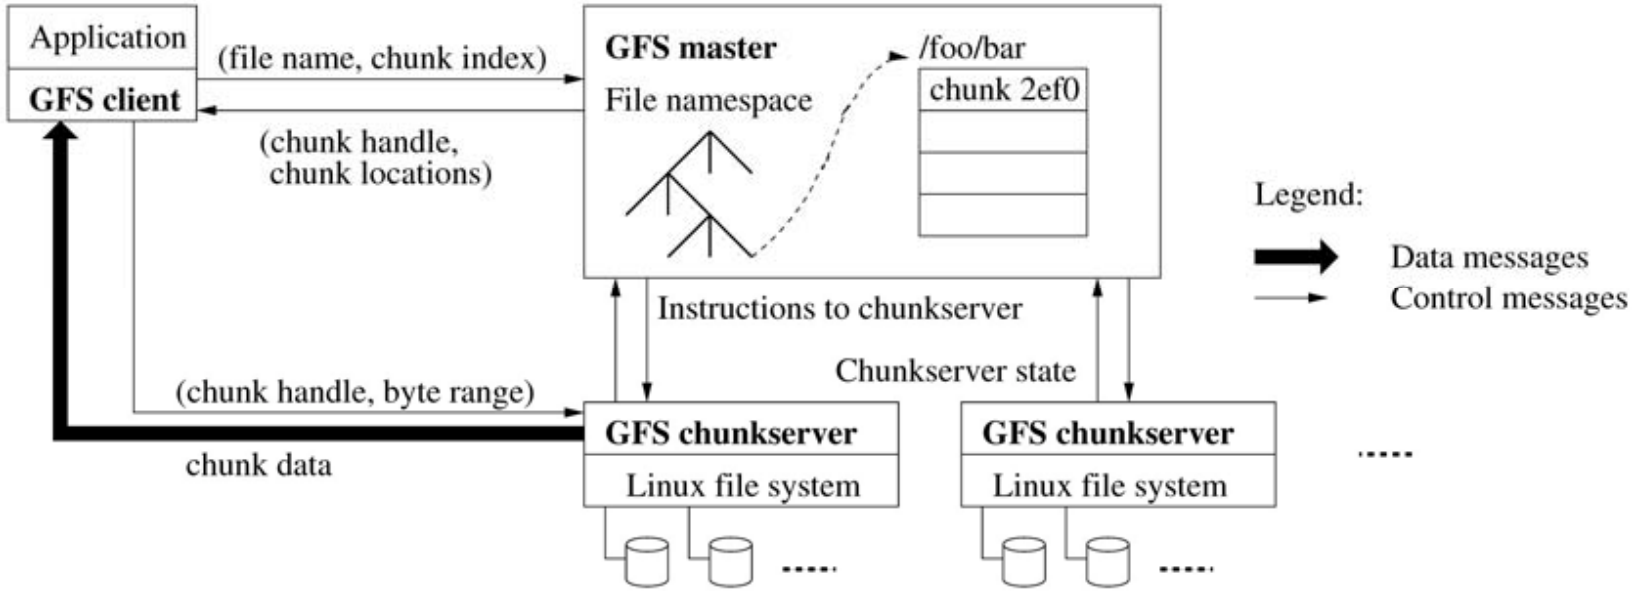

MapReduce

- Jeffrey Dean and Sanjay Ghemawat
  MapReduce: Simplified Data Processing on Large Clusters
  OSDI'04: Sixth Symposium on Operating System Design and Implementation, San Francisco, CA, December, 2004.
  see http://labs.google.com/papers/mapreduce.html

- Google Tutorial at
  http://code.google.com/intl/de-DE/edu/submissions/mapreduce-minilecture/listing.html

Functional Programming Review

- Functional operations do not modify data structures: They always create new ones
- Original data still exists in unmodified form
- Data flows are implicit in program design
- Order of operations does not matter

- History: LISP programming

Functional Programming Review

```
fun foo(l: int list) =
  sum(l) + mul(l) + length(l)
```

Order of sum() and mul(), etc does not matter – they do not modify l

Map

map f lst: (’a->’b) -> (’a list) -> (’b list)

Creates a new list by applying f to each element of the input list; returns output in order.

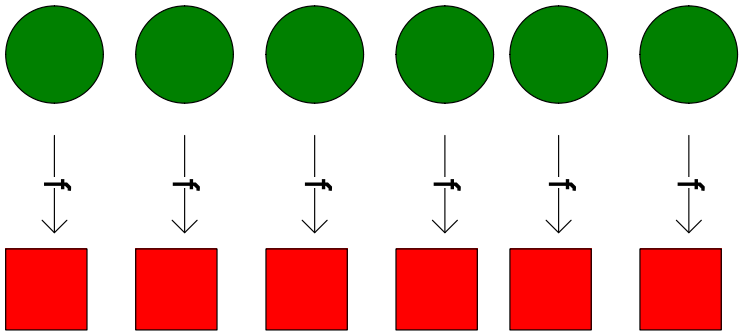

Fold

fold f x_0 lst: ('a*'b->'b)->'b->('a list)->'b

Moves across a list, applying *f* to each element plus an *accumulator*. f returns the next accumulator value, which is combined with the next element of the list

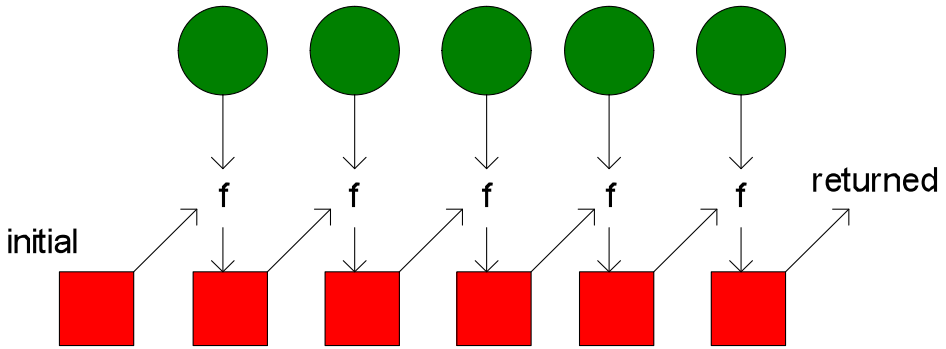

Implicit Parallelism in map

- In a purely functional setting, elements of a list being computed by map cannot see the effects of the computations on other elements
- If order of application of *f* to elements in list is *commutative*, we can reorder or parallelize execution
- This is the “secret” that MapReduce exploits

MapReduce Programming Model

- Borrows from functional programming
- Implements two basic functions:
  - map (in_key, in_value) -> (out_key, intermediate_value) list
  - reduce (out_key, intermediate_value list) -> out_value list

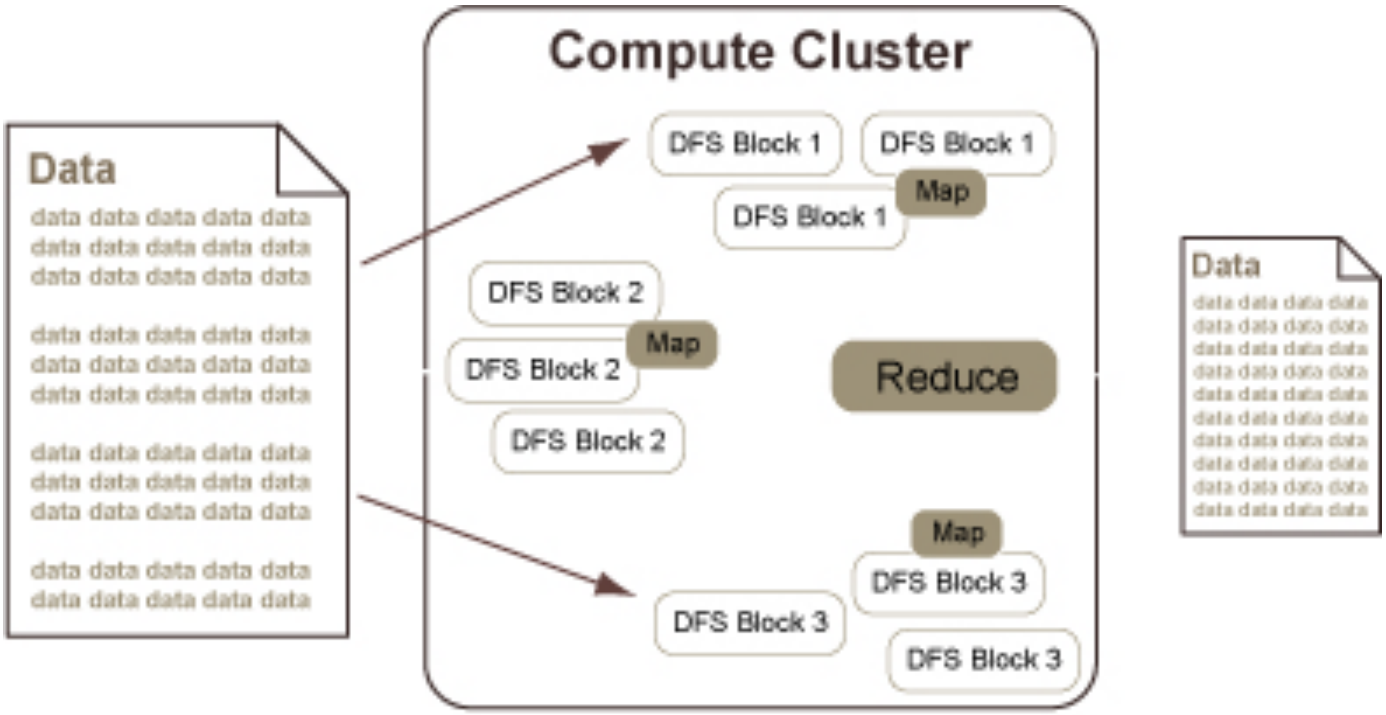

Map

```
map  (in_key, in_value) ->
   (out_key, intermediate_value) list
```

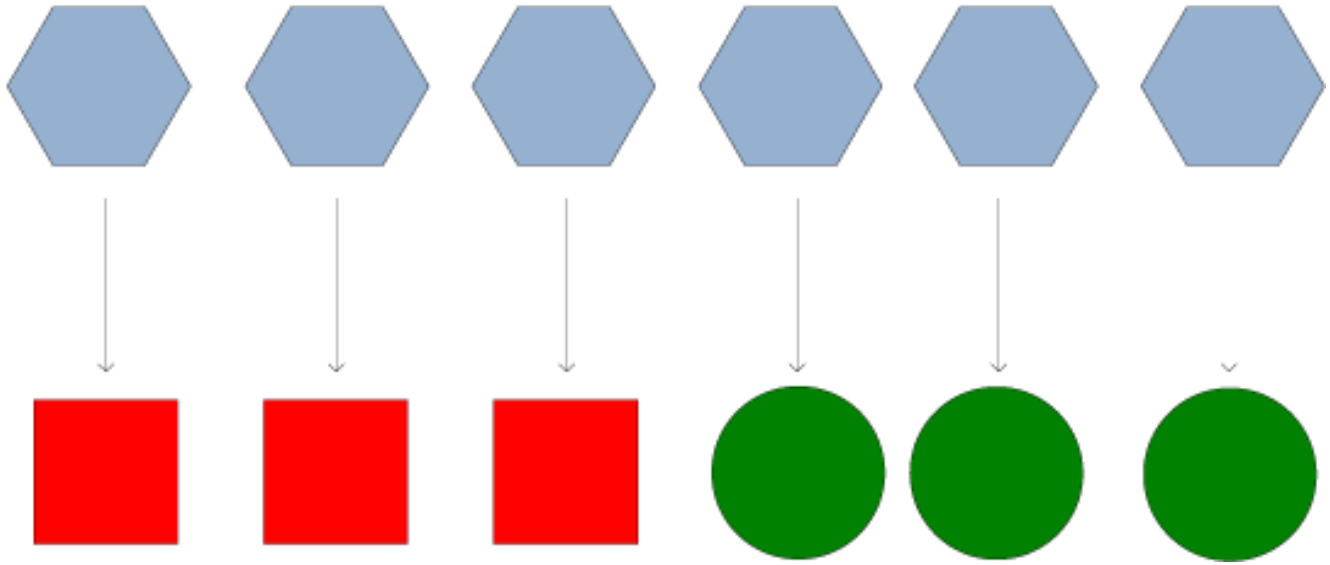

- Records from the data source (lines out of files, rows of a database, etc) are fed into the map function as key*value pairs: e.g., (filename, line).
- map() produces one or more *intermediate* values along with an output key from the input.

Reduce

```
reduce (out_key, intermediate_value list) ->
        out_value list
```

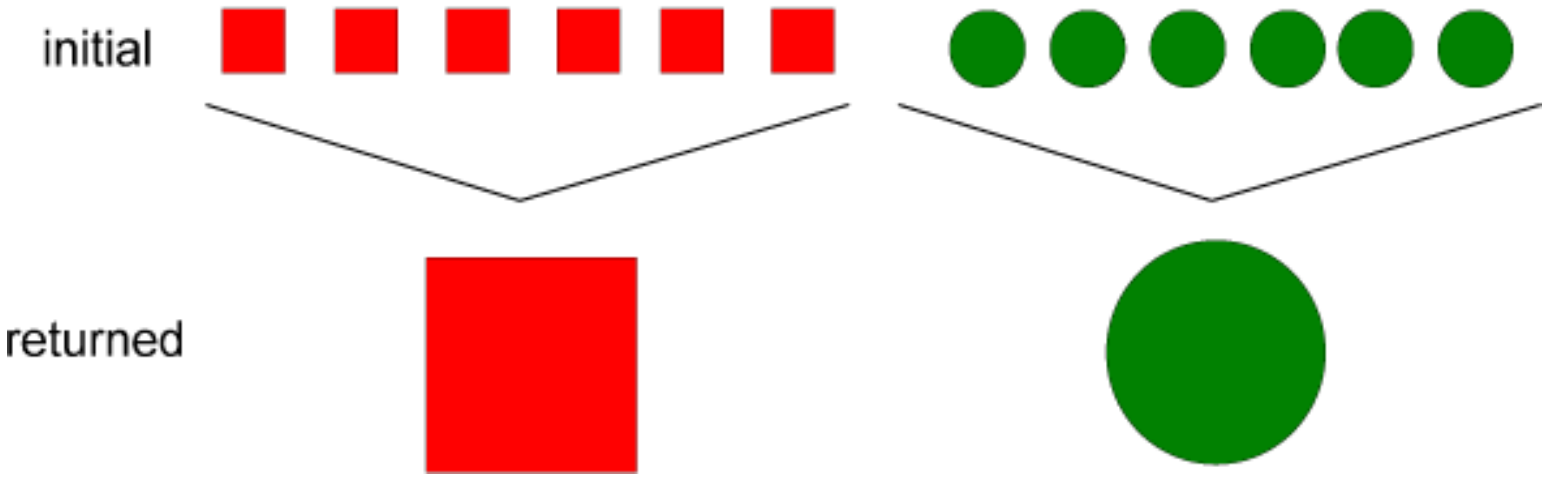

- After the map phase is over, all the intermediate values for a given output key are combined together into a list
- reduce() combines those intermediate values into one or more *final values* for that same output key
- (in practice, usually only one final value per key)

Parallelism

- map() functions run in parallel, creating different intermediate values from different input data sets
- reduce() functions also run in parallel, each working on a different output key
- All values are processed *independently*
- Bottleneck: reduce phase can’t start until map phase is completely finished.

MapReduce Parallel Programming Model

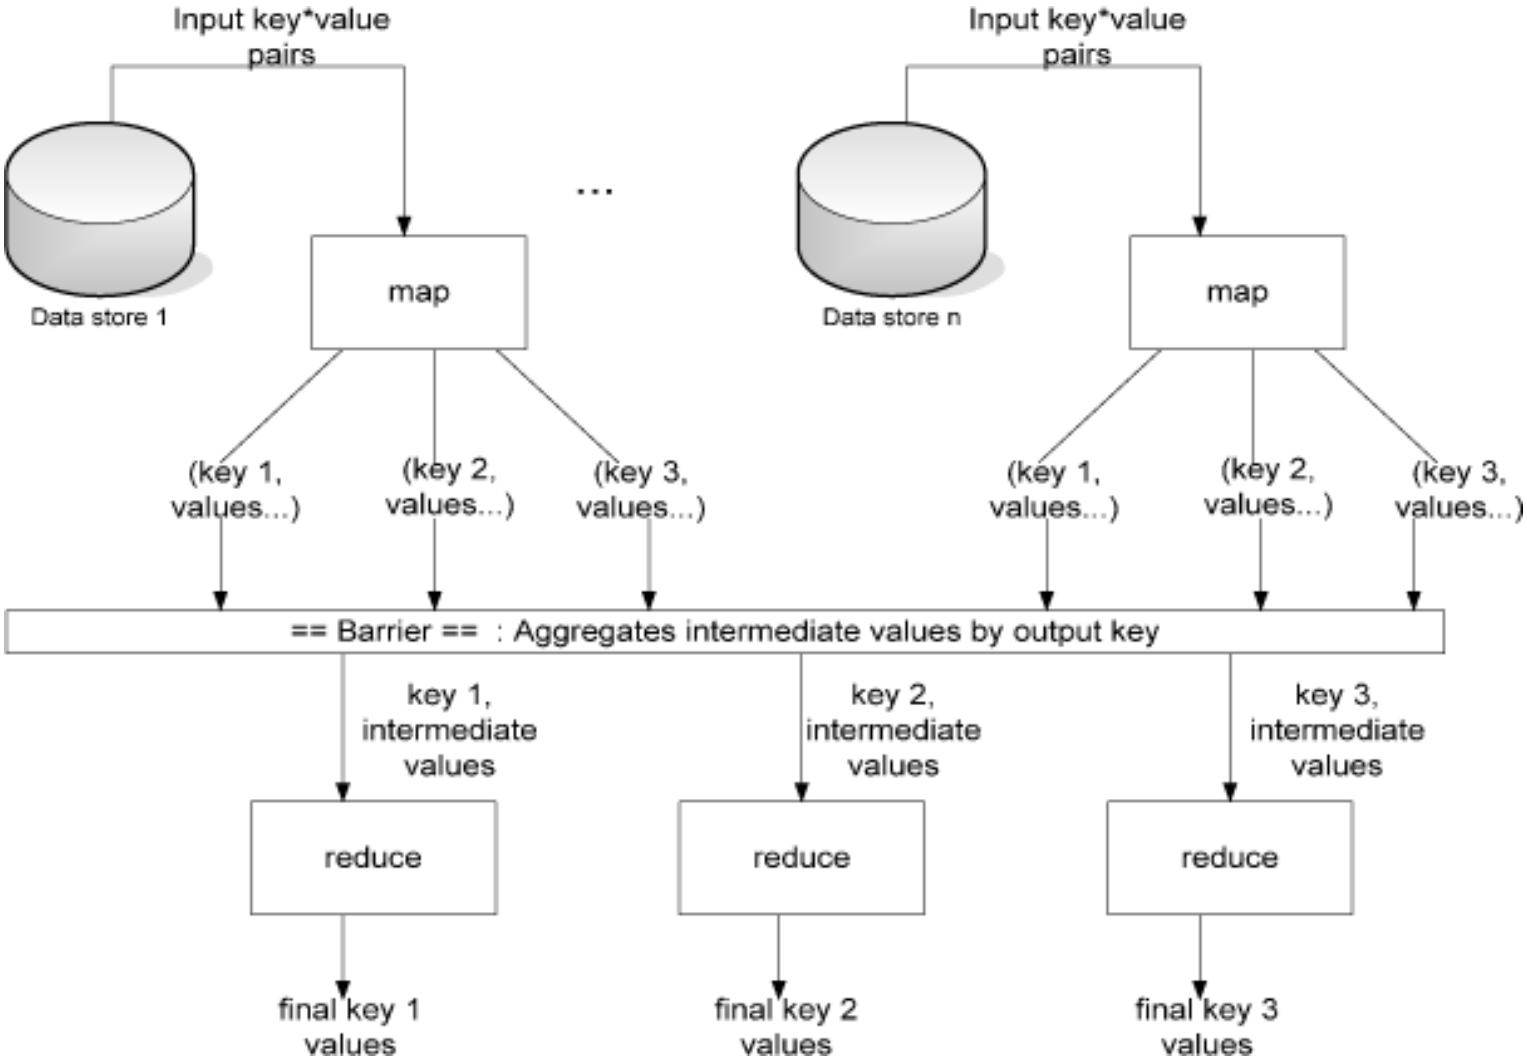

- **Bottleneck: All Map processes have to finish before Reduce starts!**

MapReduce Example

- Counting words in documents

```
map(String input_key, String input_value):
// input_key: document name
// input_value: document contents
    for each word w in input_value:
        EmitIntermediate(w, 1);

reduce(String output_key, Iterator<int> intermediate_values):
// output_key: a word
// output_values: a list of counts
    int result = 0;
    for each v in intermediate_values:
        result += v;
    Emit(result);
```

Wordcount Example

- Hello World Bye World
- Hello Hadoop Goodbye Hadoop

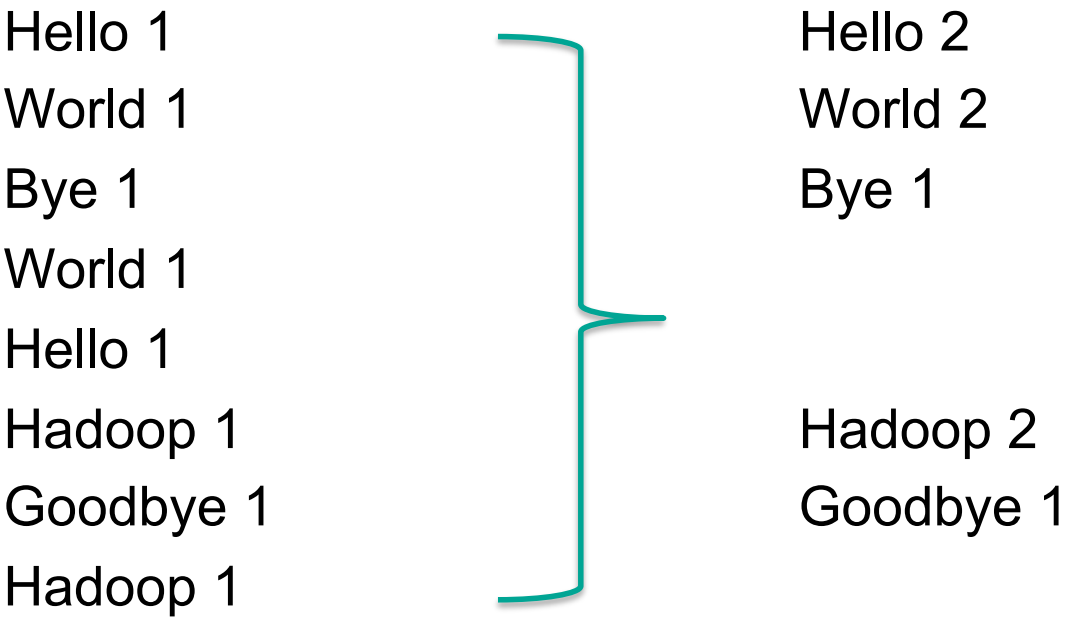

Programming the Cloud: Hadoop

- Reproduce the proprietary software infrastructure developed by Google (Started by Doug Cutting 2004)
- Hadoop implements
  - Parallel programming model (MapReduce)
  - Hadoop Distributed File System (HDFS)
  - Parallel database (HBase)
  - Programming environment (Pig)
  - Data warehouse infrastructure (HIVE)
  - Data collection and analysis (Chukwa)
  - Machine based learning (Mahout)
- Largest Cluster at Yahoo!: 32.000 cores and 16 PetaByte storage (1PB/16h)

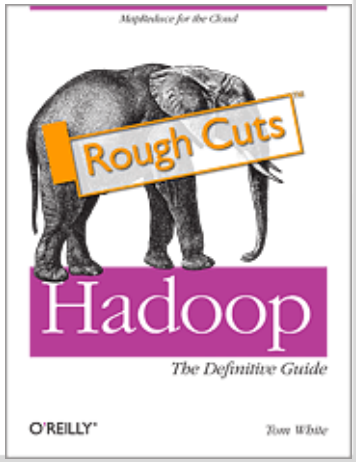

HDFS Architecture

Block Replication

Namenode (Filename, numReplicas, block-ids, …)
/users/sameerp/data/part-0, r:2, {1,3}, …
/users/sameerp/data/part-1, r:3, {2,4,5}, …

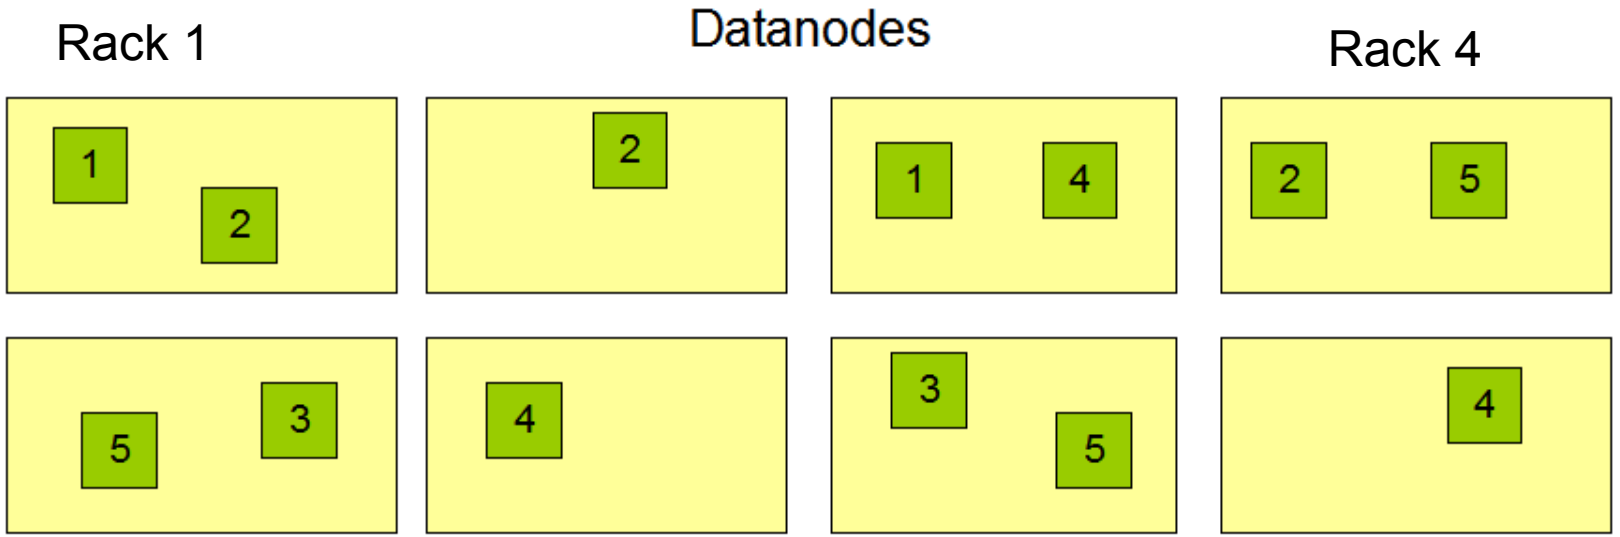

- Namenode holds meta data (File names, block-ids etc.)
- Blocks are replicated over data nodes (Default size: 64 MB)

MapReduce

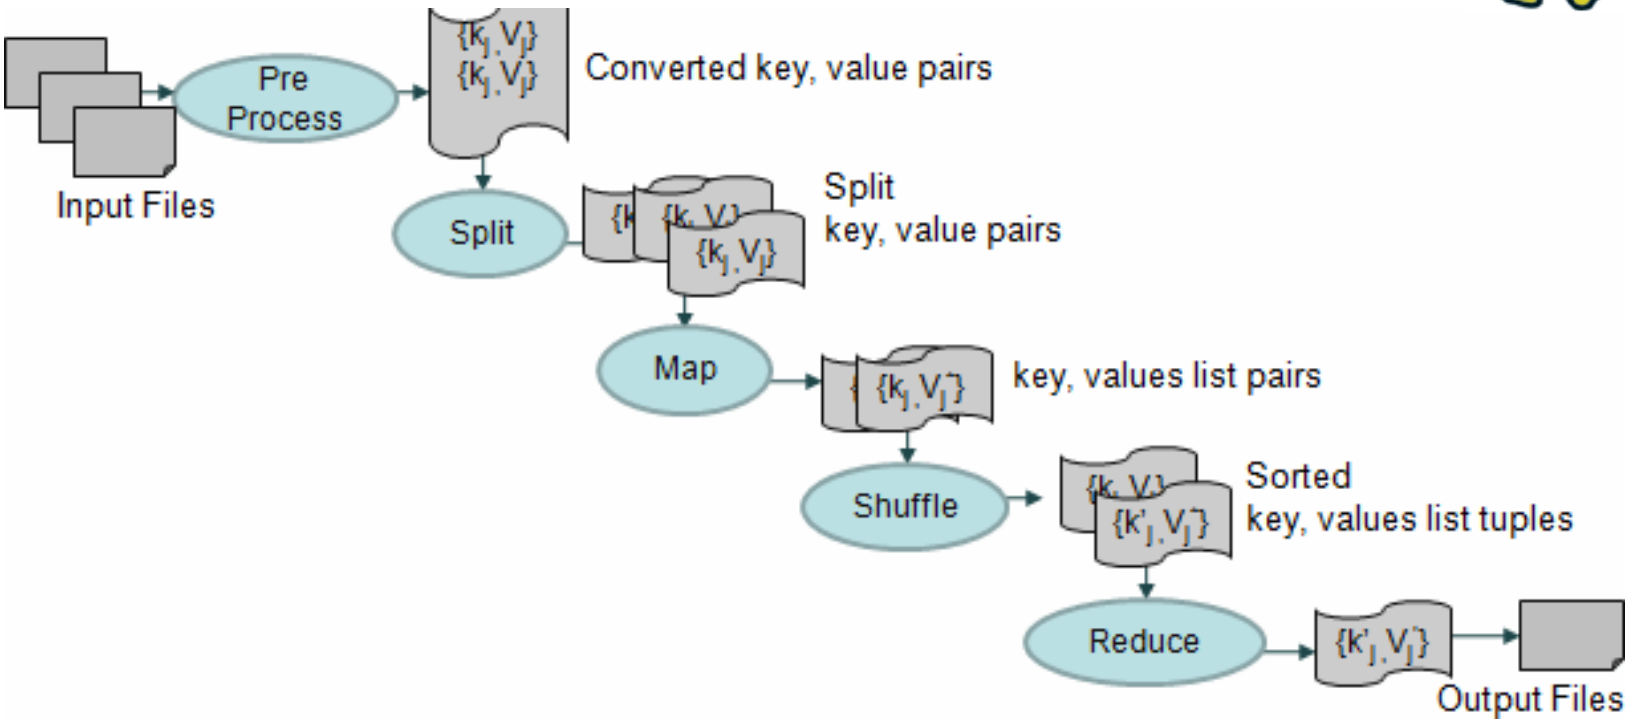

- Map: Generate intermediate Key/Value pairs from input data
  map (k,v) -> list (k1,v1)

- Reduce: Generate output data from intermediate data
  reduce (k1, list(v1)) -> list(v1)

Orchestration of MapReduce Tasks

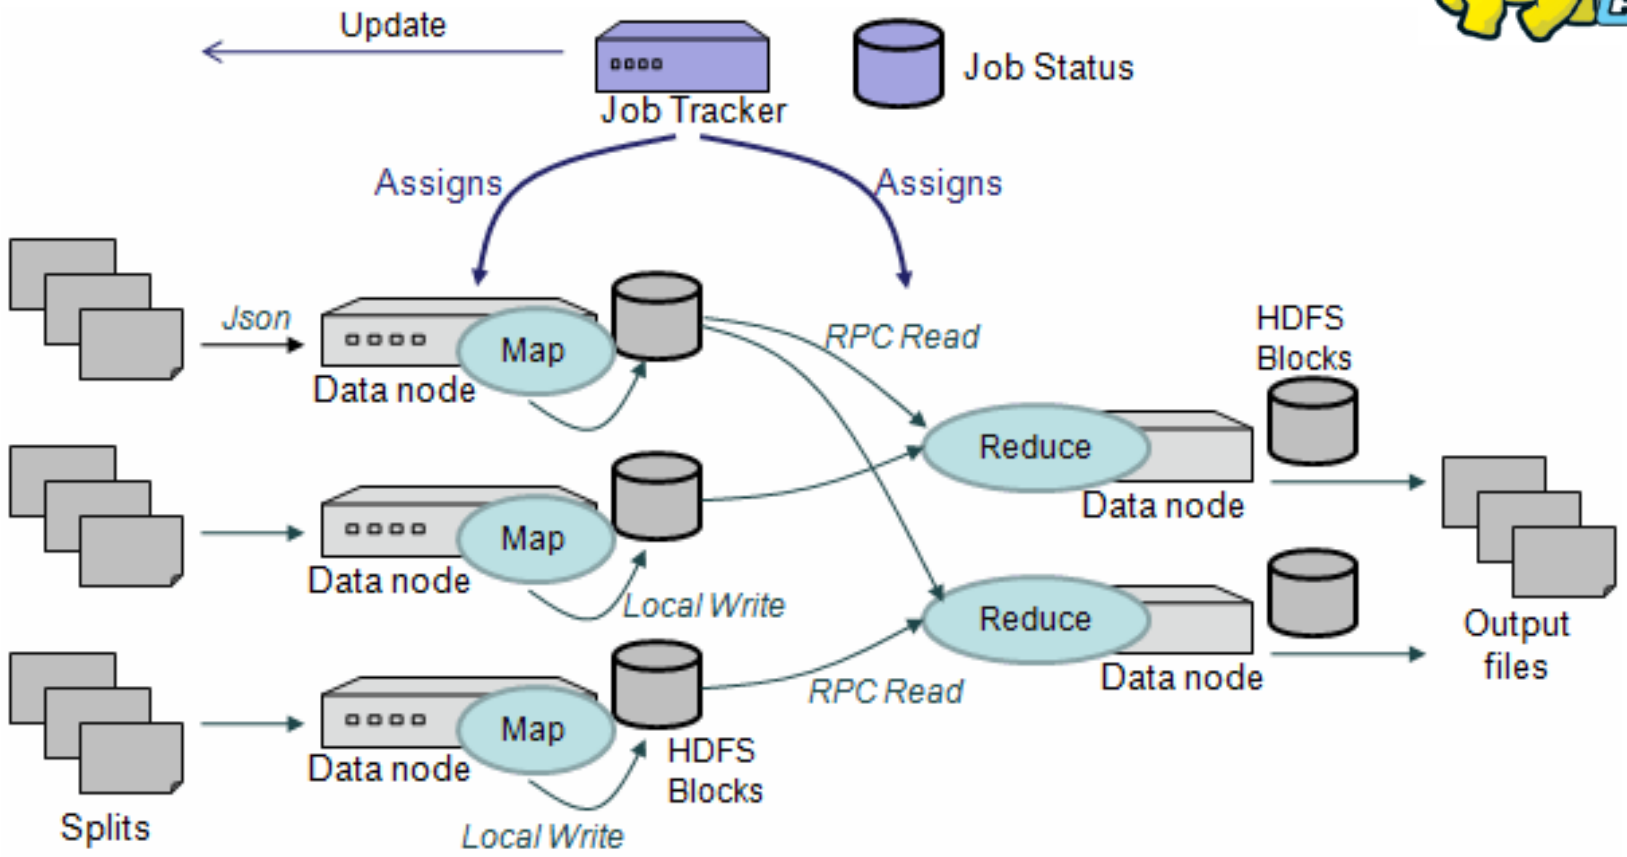

- Store data in HDFS
- Job Tracker starts MapReduce Job
- Assign tasks to task tracker nodes
- Assemble output files in HDFS

Pig

- Platform for analyzing large data sets
  - High level language to express data analysis programs
  - Parallel infrastructure

- Pig Latin
  - Ease of programming. It is trivial to achieve parallel execution of simple, "embarrassingly parallel" data analysis tasks. Complex tasks comprised of multiple interrelated data transformations are explicitly encoded as data flow sequences, making them easy to write, understand, and maintain.
  - Optimization opportunities. The way in which tasks are encoded permits the system to optimize their execution automatically, allowing the user to focus on semantics rather than efficiency.
  - Extensibility. Users can create their own functions to do special-purpose processing.

HAMA

Hama (means a hippopotamus in Korean) is a distributed matrix computation package. It is a library of matrix operations for large-scale processing and development environments as well as a Map/Reduce framework for a large-scale numerical analysis and data mining, that need the intensive computation power of matrix inversion, e.g., linear regression, PCA, SVM and etc. It will be useful for many scientific applications, e.g., physics computations, linear algebra, computational fluid dynamics, statistics, graphic rendering and many more.

- Scientific simulation and modeling
  - Matrix-vector/matrix-matrix multiply
  - Soving linear systems
  - Scientific graphs
- Information retrieval
  - Sorting
  - Finding eigenvalues and eigenvectors
- Computer graphics and computational geometry
  - Matrix multiply
  - Computing matrix determinate

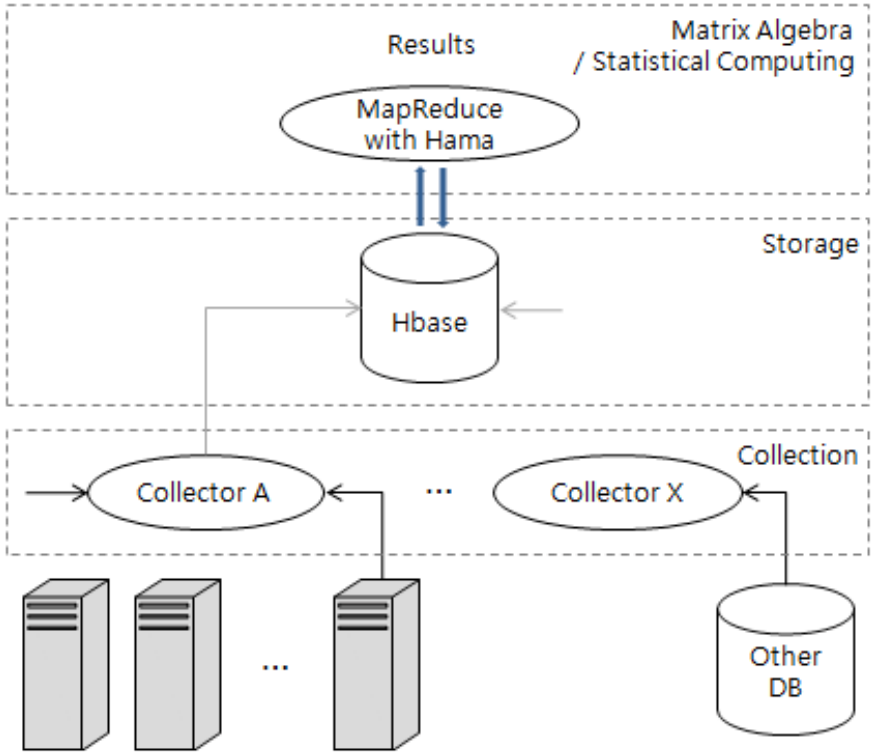

New Hadoop Cluster at KIT

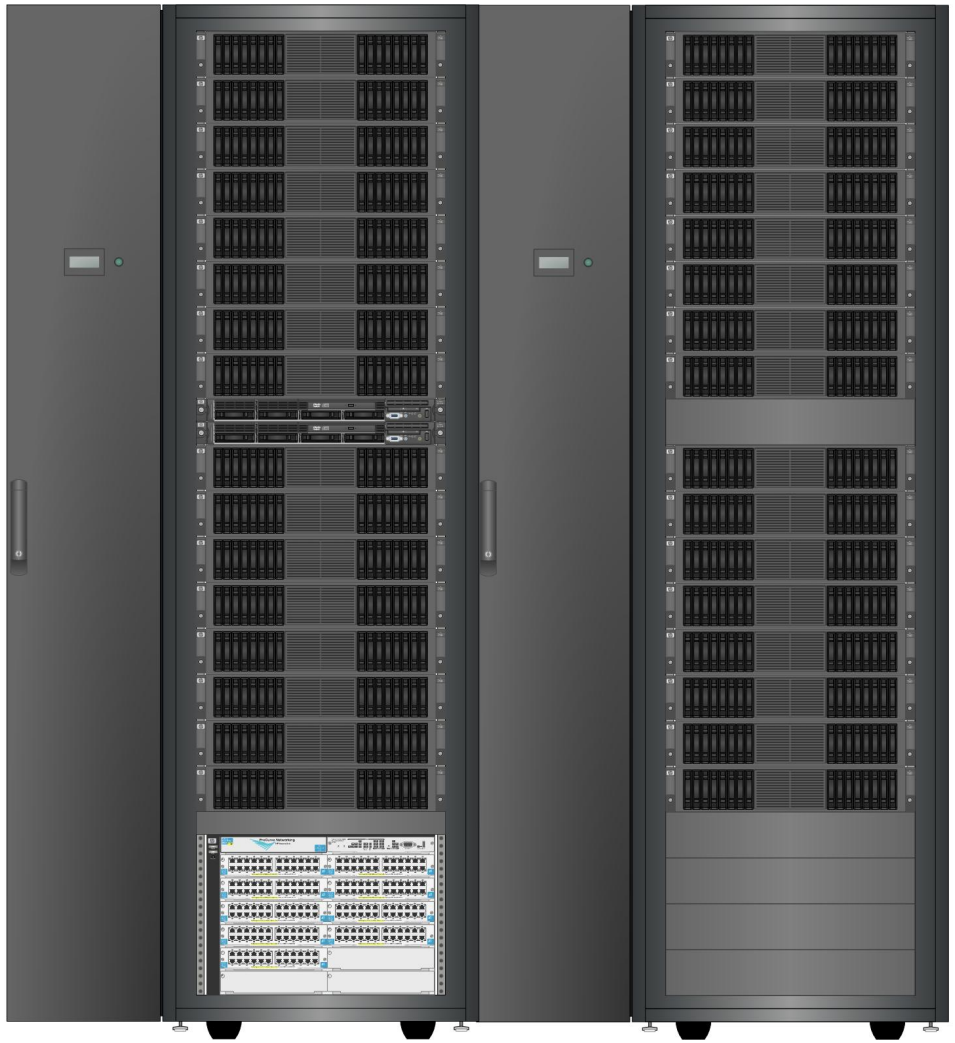

- Hadoop
- 480 Nehalem cores
- 128 TB + 256 TB storage
- OpenCirrus resource
- Application area:
  - Cloud R&D
  - Bioinformatics

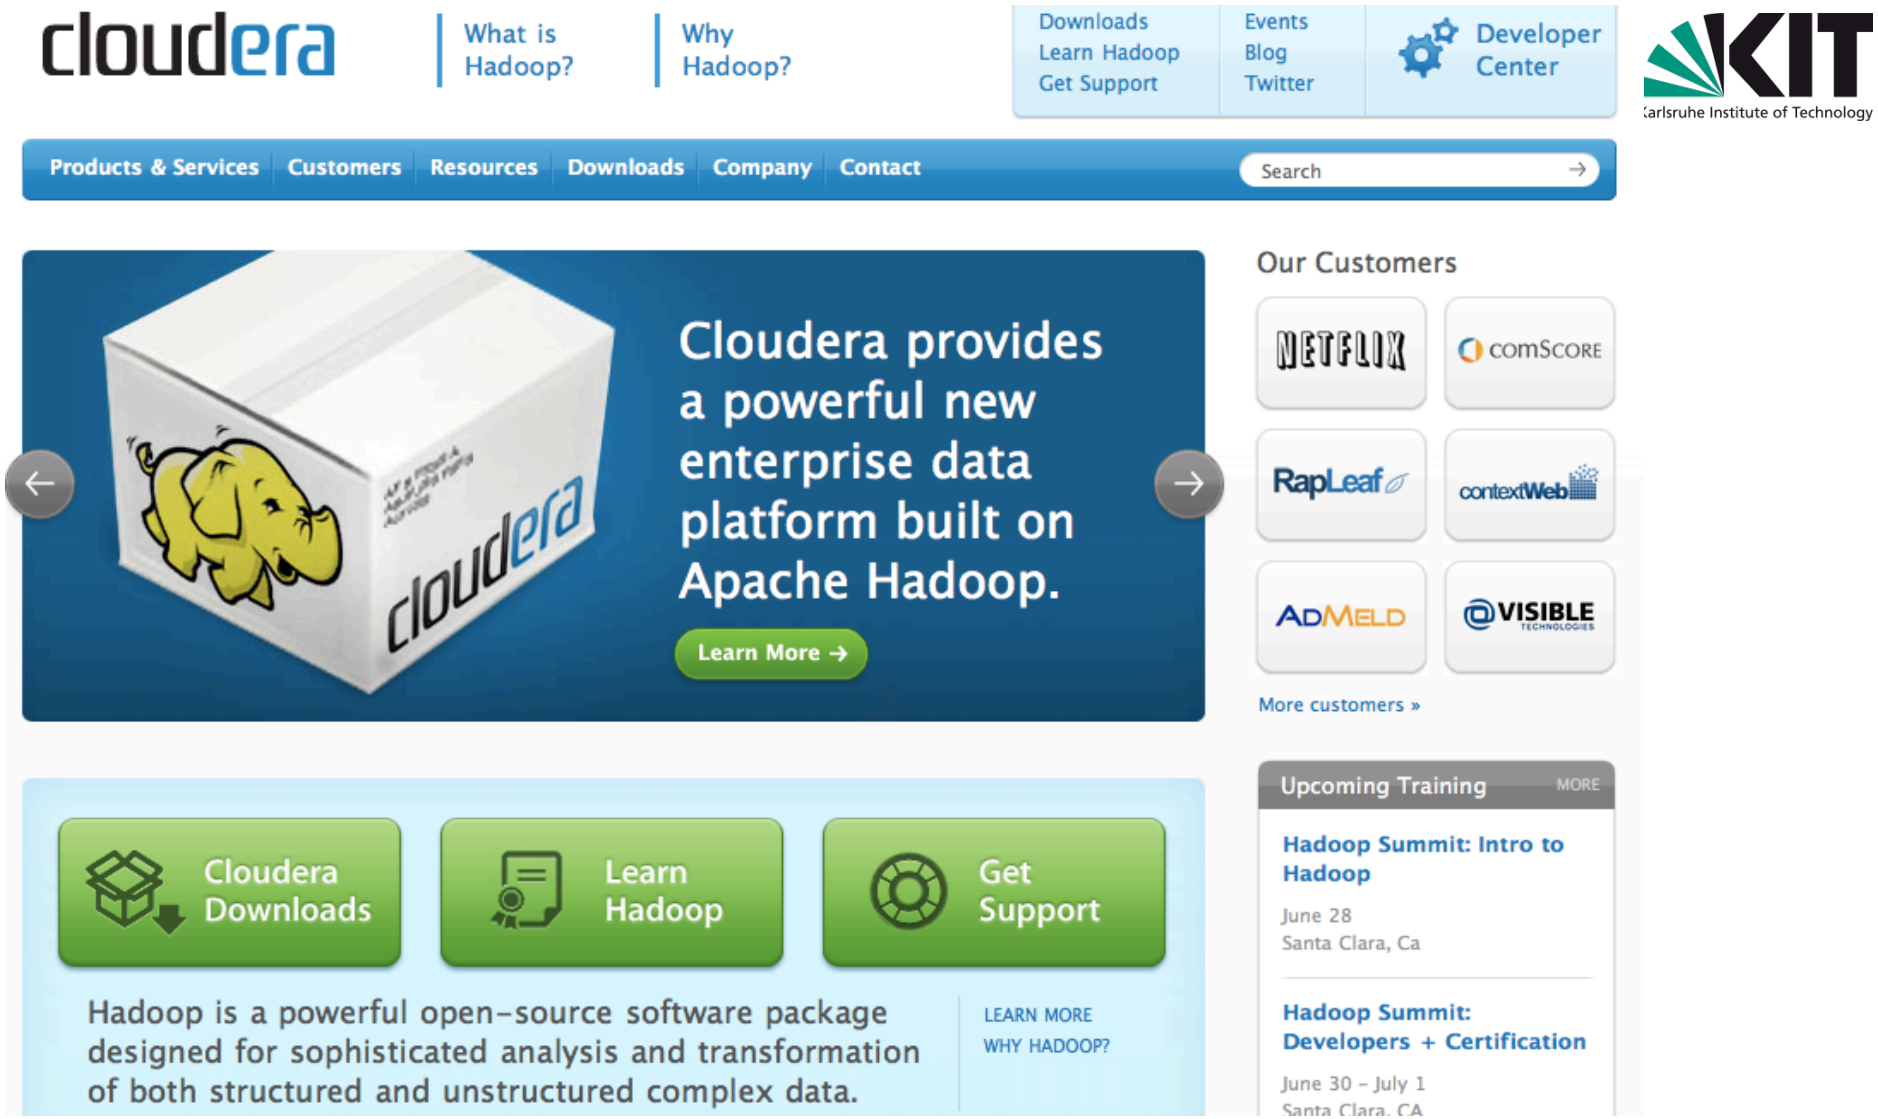

- http://www.cloudera.com/

Cloudera

- Cloudera offers a Hadoop release
  - Free download
  - Enterprise support for Hadoop
- Tutorials and training videos
- Cloudera desktop
- Virtual appliance for VMware Workstation and Fusion based on Ubuntu Linux

Cloudera Desktop

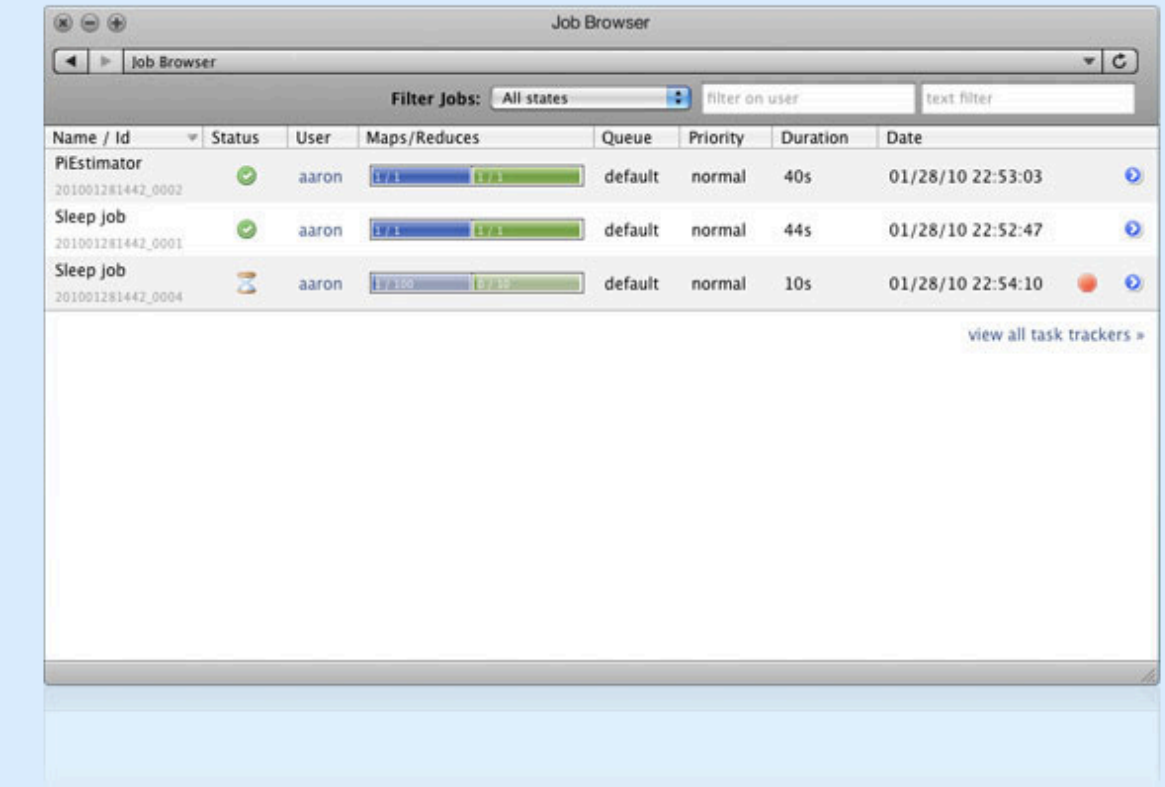

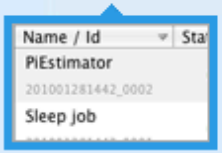

Job Browser

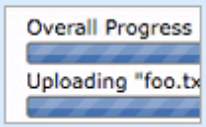

File Browser

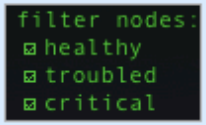

Cluster Health

Job Designer

Demo video

Amazon Web Services

http://aws.amazon.com/

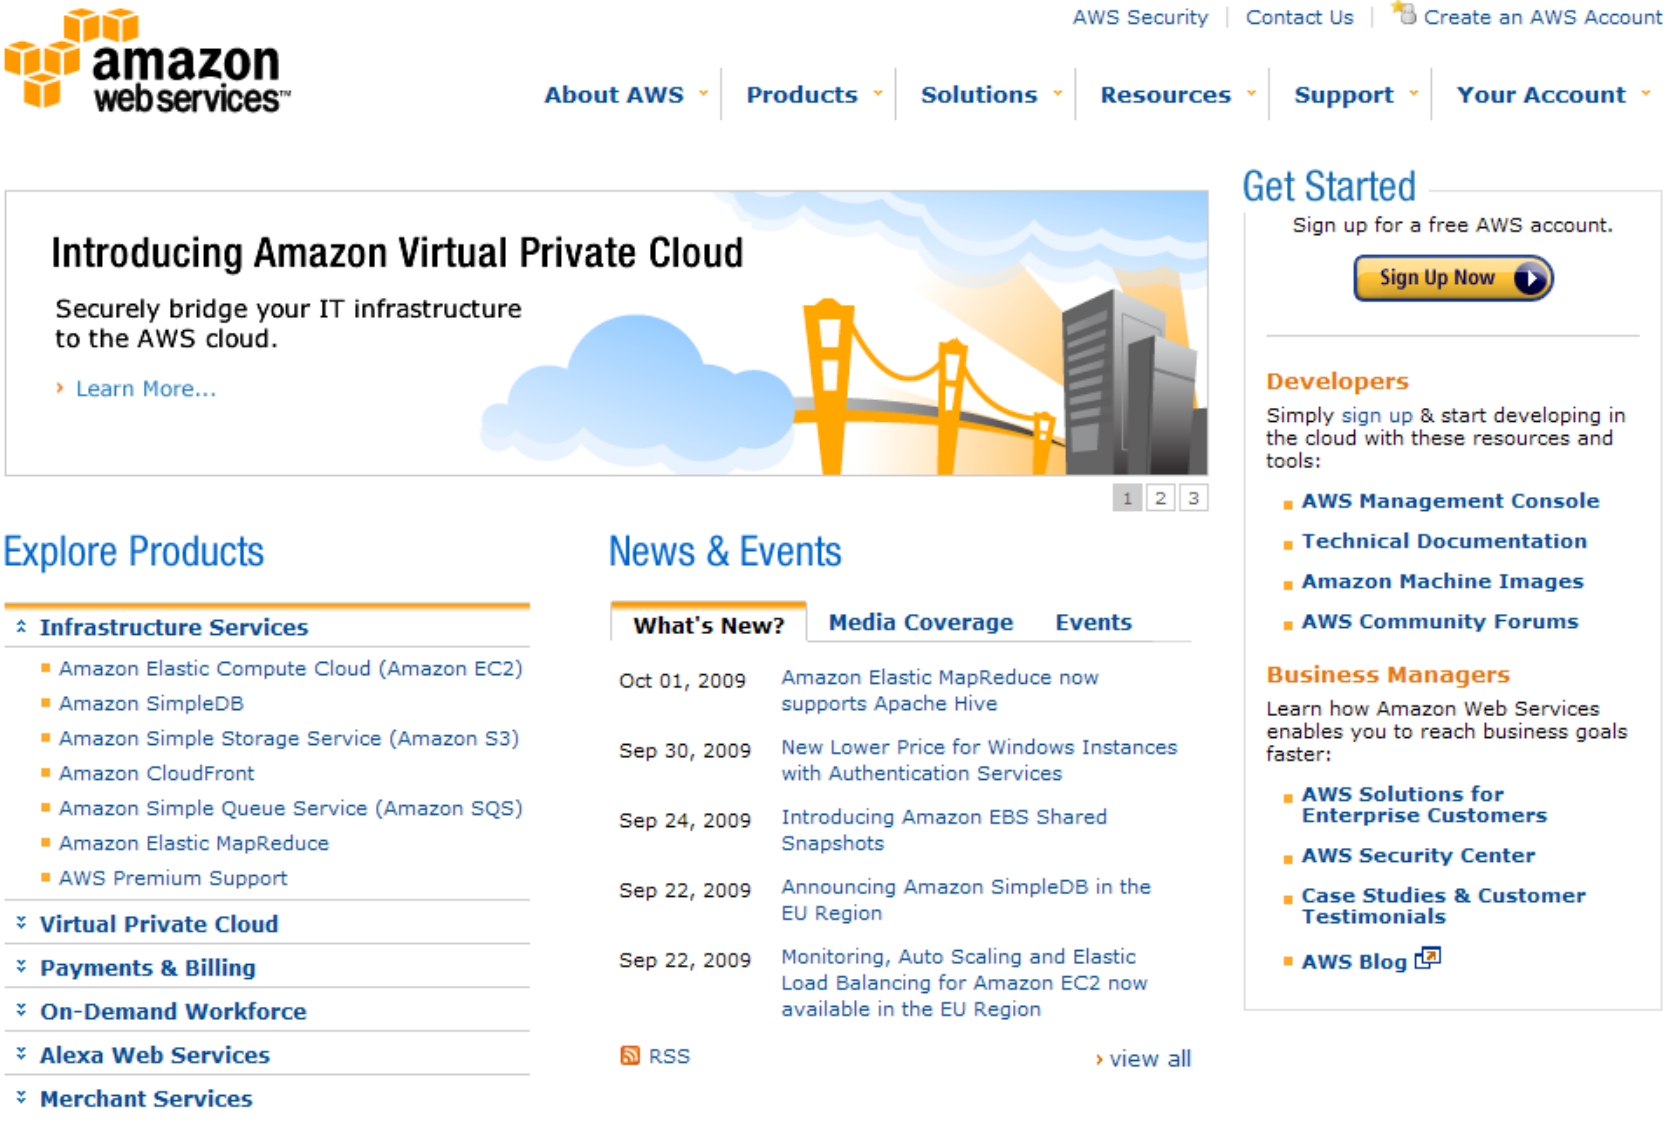

Amazon Elastic MapReduce

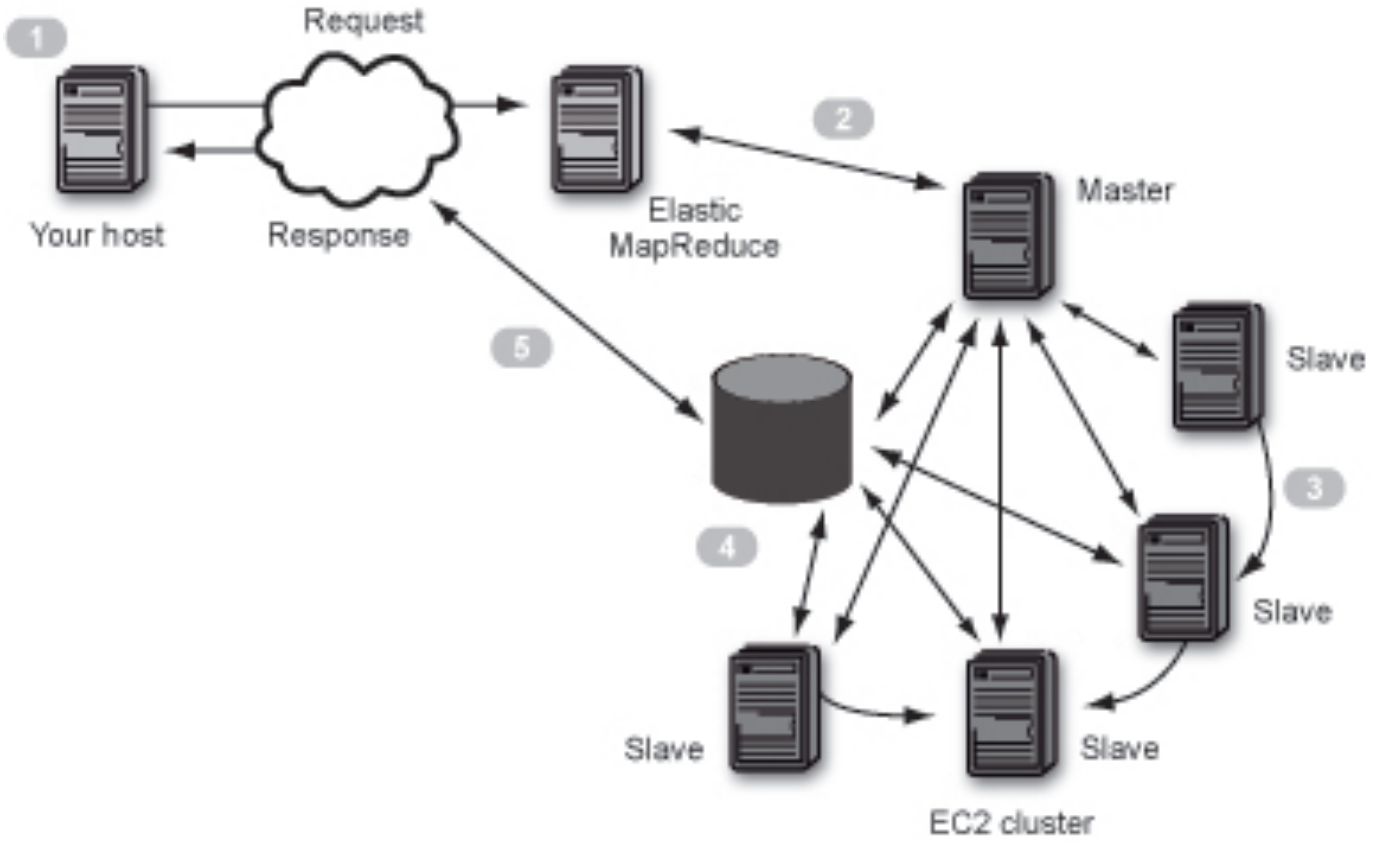

1. Load data, Map and Reduce executables to S3
2. Elastic MapReduce starts EC2 Hadoop-Cluster (Master + Slaves)
3. Hadoop generates Jobflow to distribute S3 data to cluster and to process it
4. Results are copied over to S3
5. Message is sent at the end: Retrieve the results from S3 (Browser, wget,…)

Management of Virtual Machines in EC2

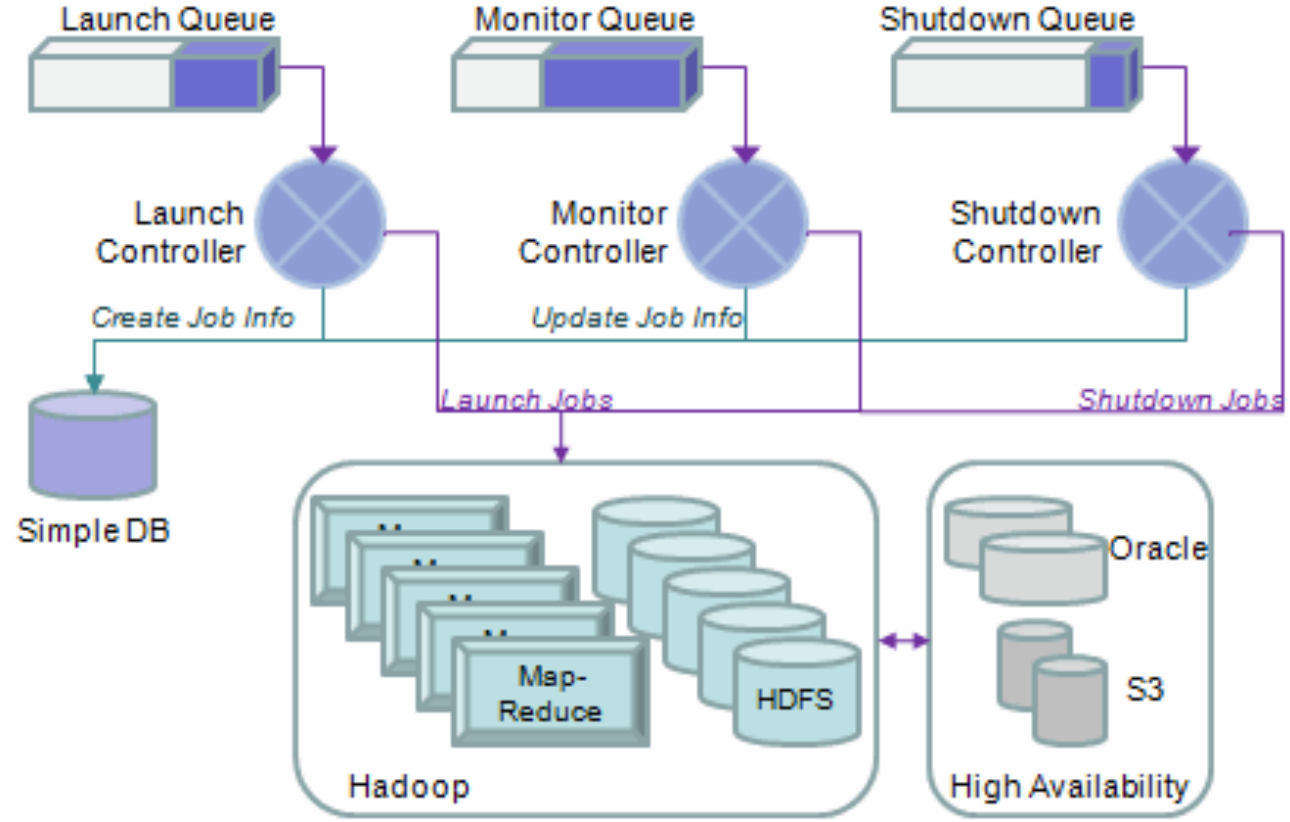

- Request queues: Launch jobs, monitor execution or shut down an existing job.
- Controllers manage and monitor the execution of those requests for jobs.
- Simple DB: Statistics on execution of the controllers and jobs for reporting purpose
- Instances are virtual machines that execute jobs
- HDFS is the default local distributed block-based storage

Amazon Elastic MapReduce API

- RunJobFlow: Creates a job flow request, starts EC2 instances and begins processing.
- DescribeJobFlows: Provides status of your job flow request(s).
- AddJobFlowSteps: Adds additional step to an already running job flow.
- TerminateJobFlows: Terminates running job flow and shutdowns all instances.

Amazon Elastic MapReduce Pricing

Standard Amazon EC2 Instances	Amazon EC2 Price per hour (On-Demand Instances)	Amazon Elastic MapReduce Price per hour
Small (Default)	$0.10 per hour	$0.015 per hour
Large	$0.40 per hour	$0.06 per hour
Extra Large	$0.80 per hour	$0.12 per hour
High CPU Instances	**Amazon EC2 Price per hour (On-Demand Instances)**	**Amazon Elastic MapReduce Price per hour**
Medium	$0.20 per hour	$0.03 per hour
Extra Large	$0.80 per hour	$0.12 per hour

- **Price in addition to EC2**

Amazon MapReduce Management Console (1)

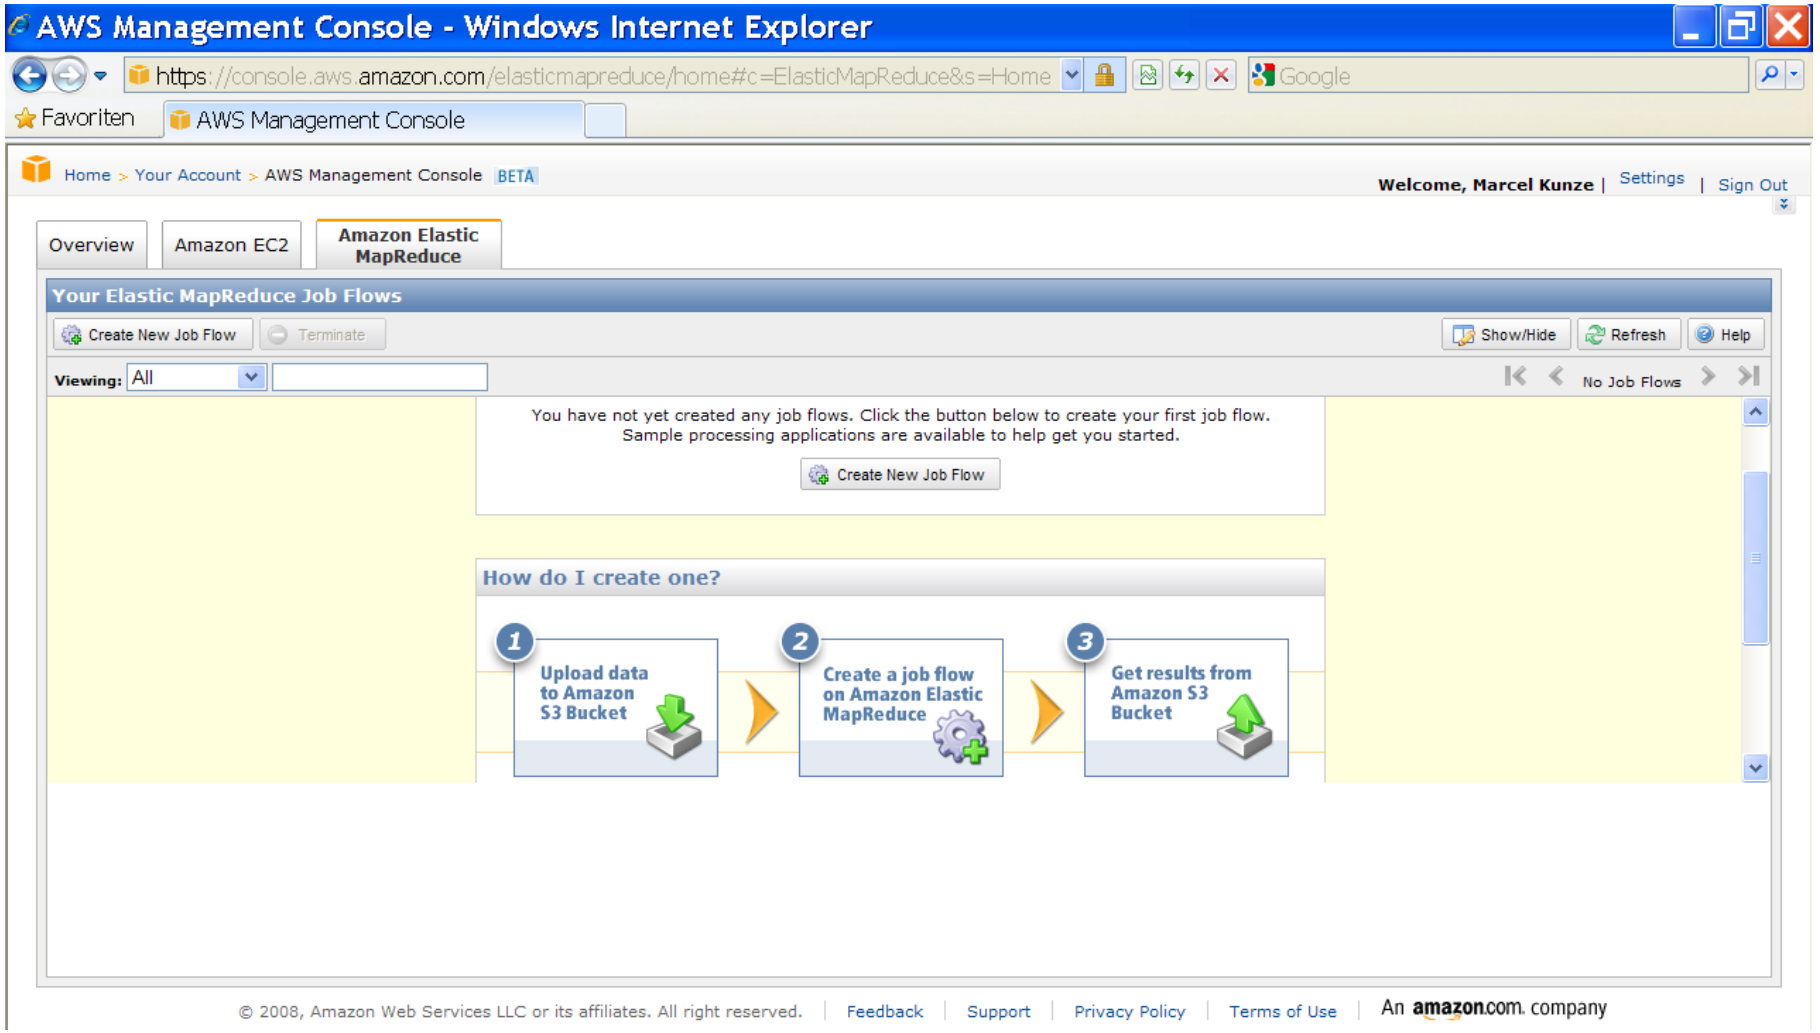

Amazon MapReduce Management Console (2)

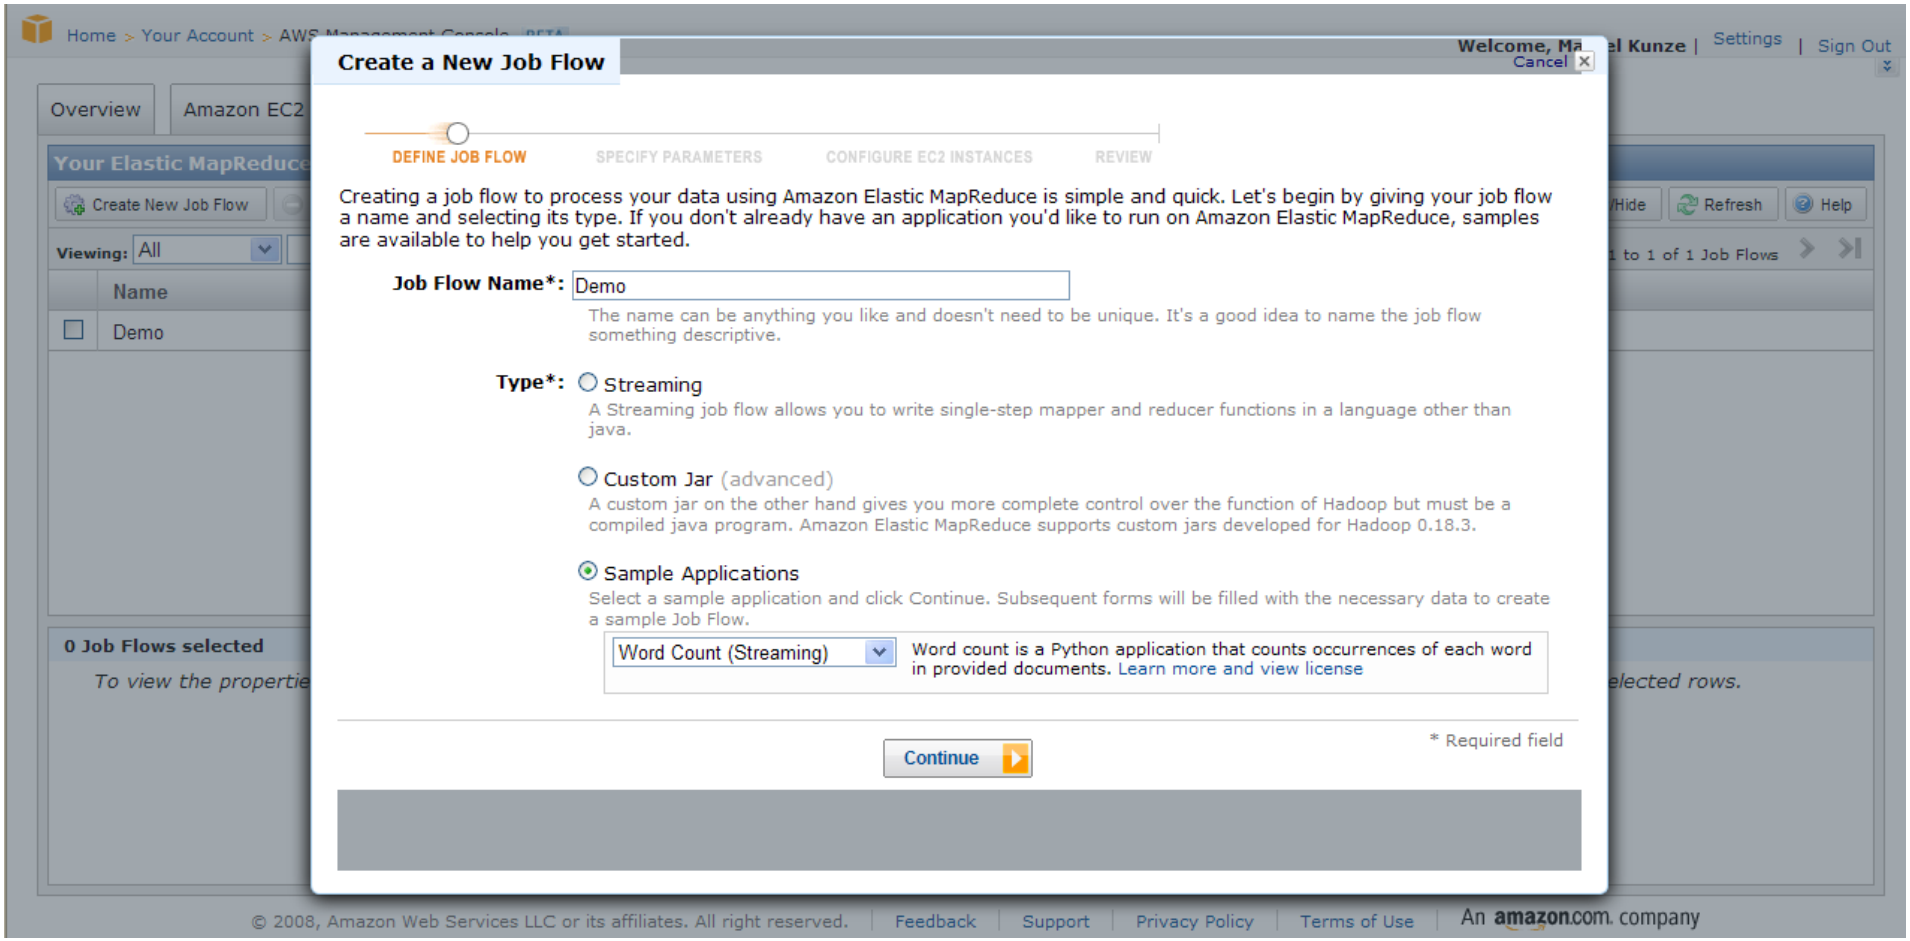

Amazon MapReduce Management Console (3)

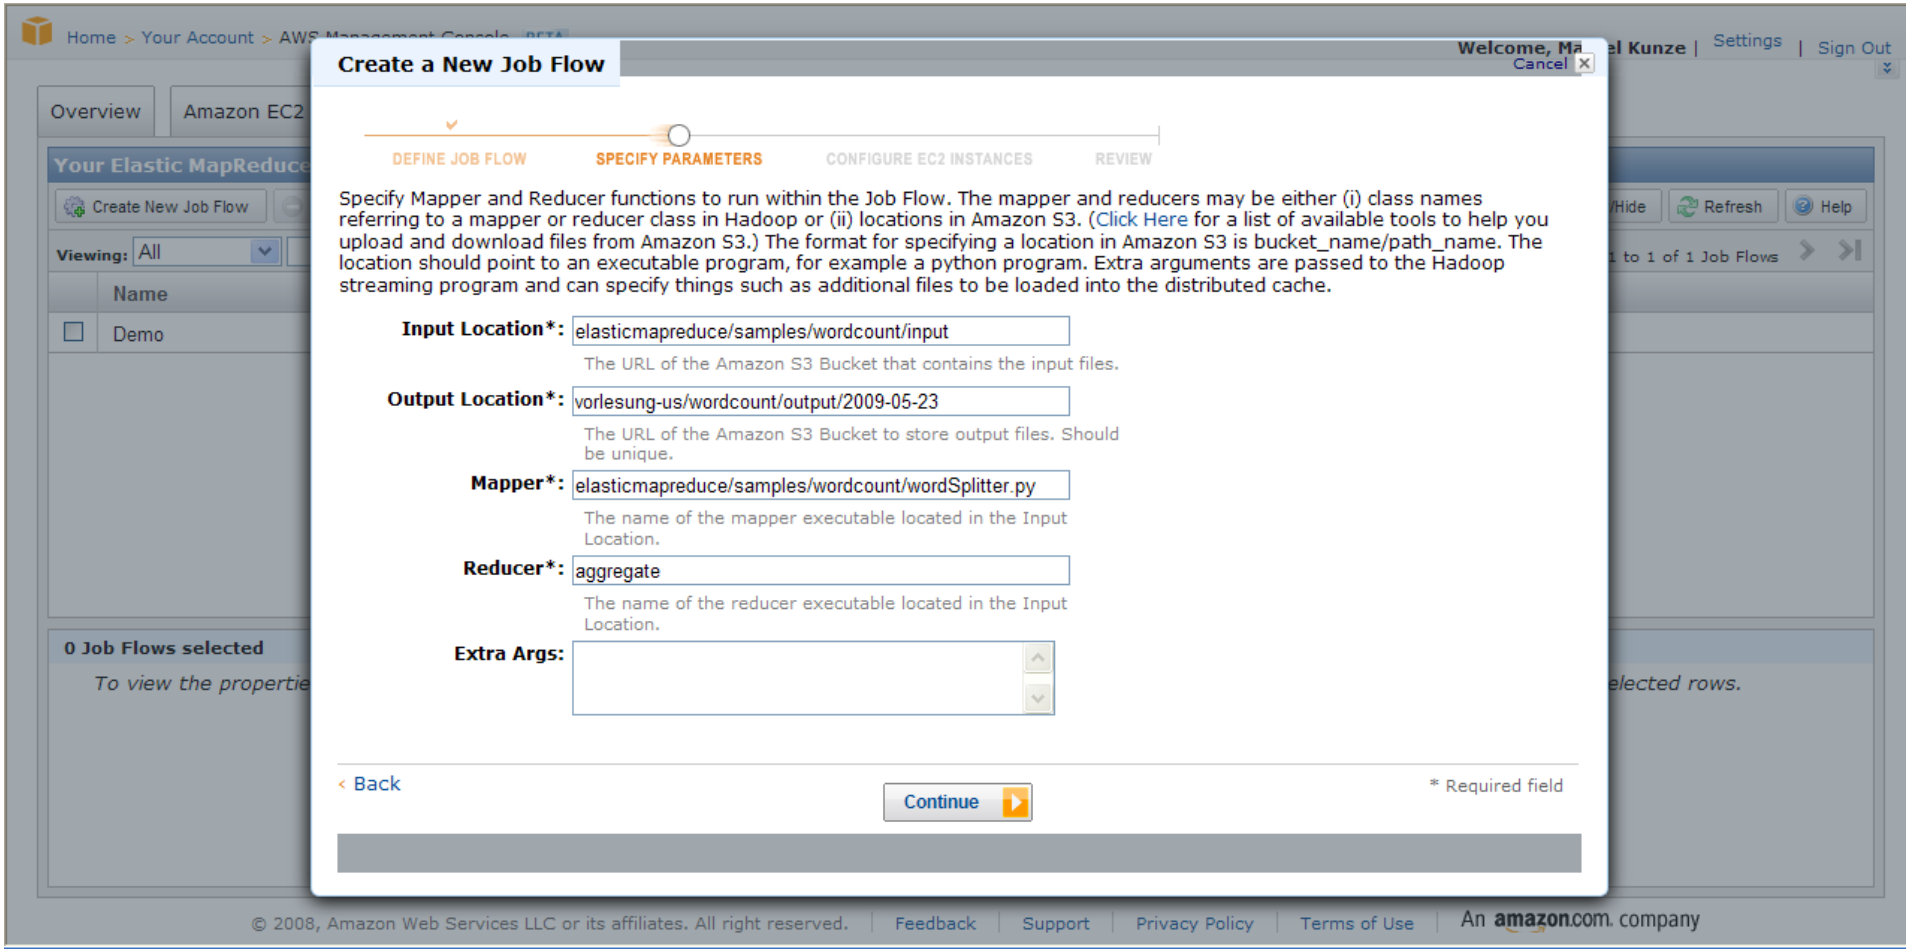

Amazon MapReduce Management Console (4)

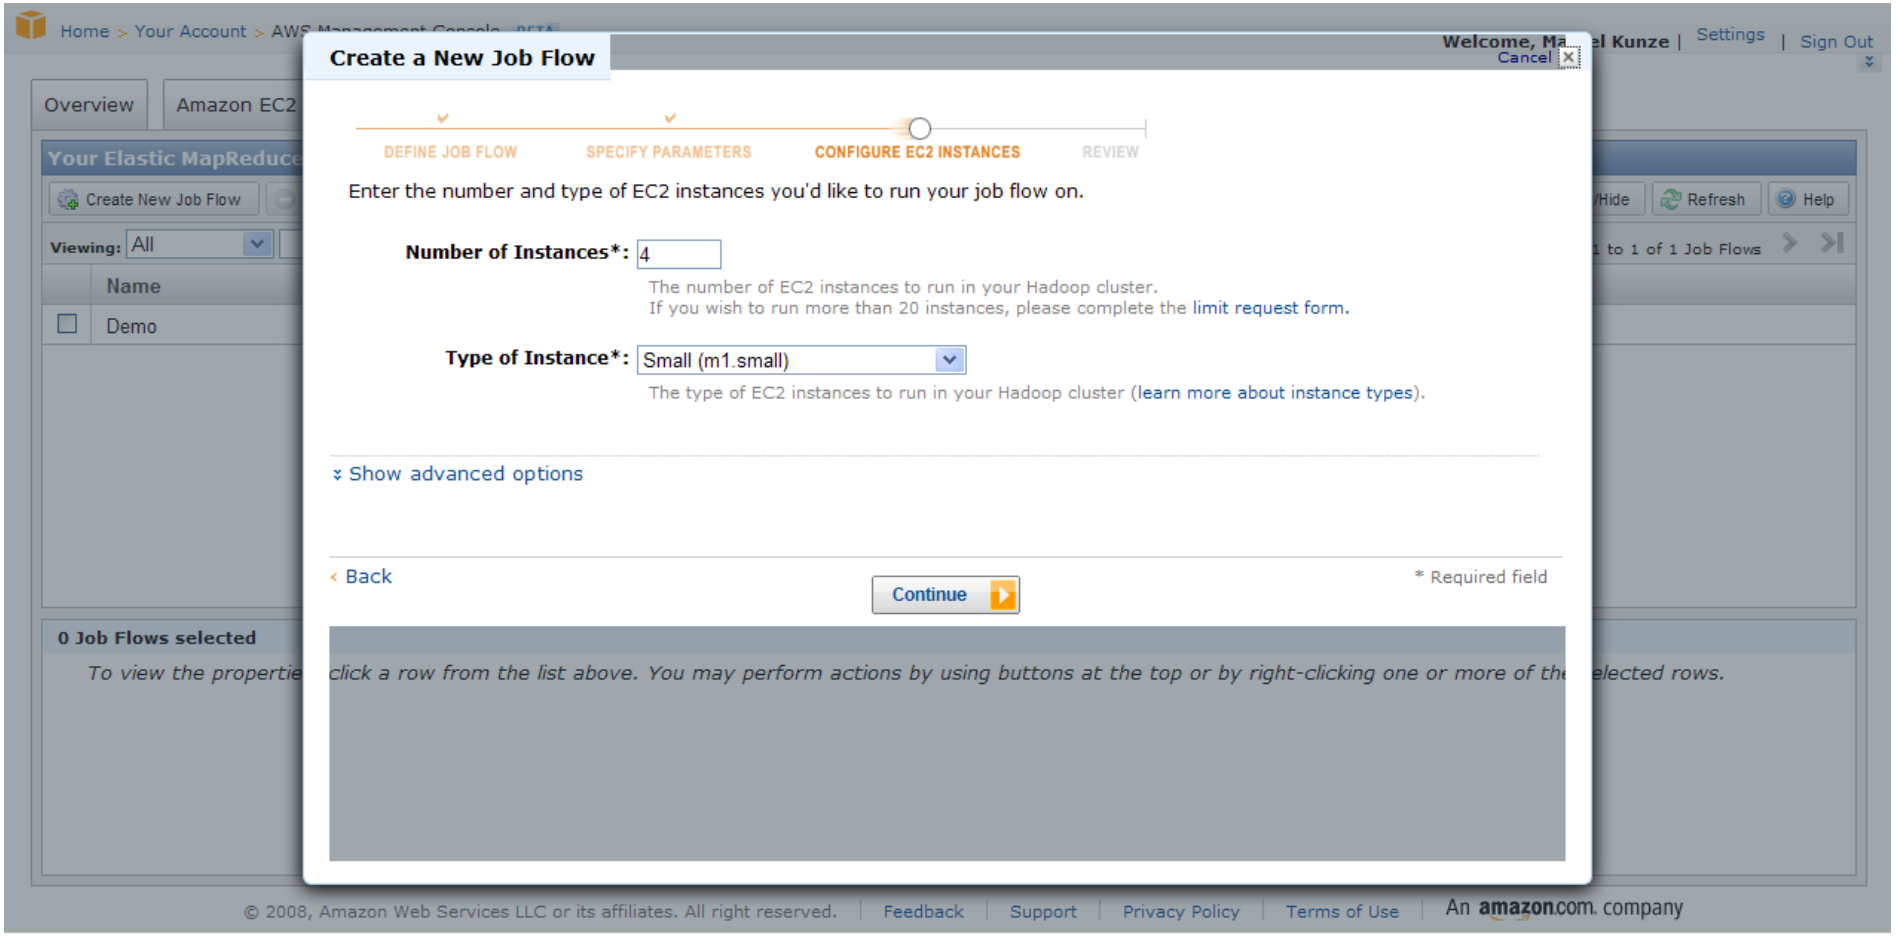

Amazon MapReduce Management Console (5)

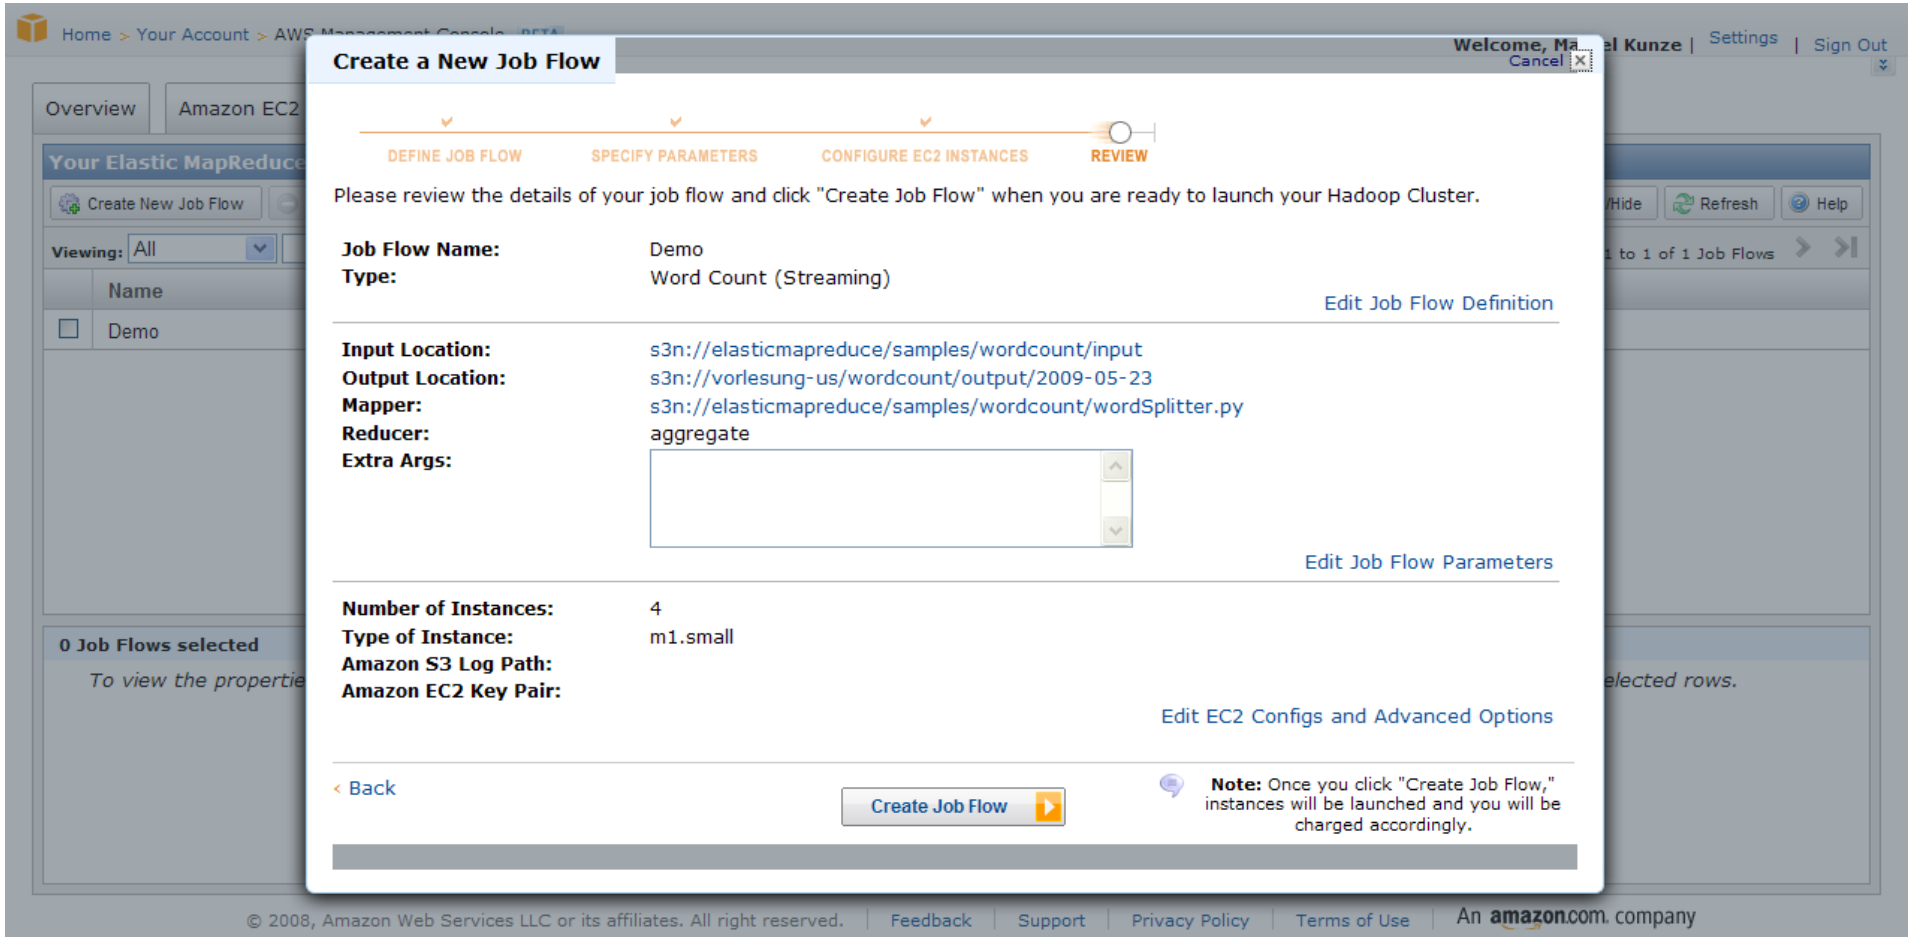

Summary

- MapReduce
  - Useful abstraction to simplify large-scale computations
  - Functional programming paradigm can be applied to large-scale applications
  - Fun to use: focus on problem, let library deal w/ messy details

- Hadoop
  - Re-implements Google software infrastructure as OpenSource
  - Ecosystem of useful services to process large data
  - Cloudera offers Hadoop release with enterprise support
  - Amazon Elastic MapReduce Service


©johnlund.com