



## **MapReduce & Apache Hadoop**

Enis Söztutar  
enis@apache.org

Ekinoks Yazılım

# Hakkında

- Enis Söztutar
- Apache Hadoop committer & Project Committee Member
- Ekinoks Yazılım
- [mailto: enis @ apache . org](mailto:enis@apache.org)



# İçerik

- Problem
- Hadoop Nedir ?
- Dağıtık Dosya Sistemi (HDFS)
- MapReduce
- İlgili Projeler




# Problem

- Büyük miktarda veriyi saklamak ve işlemek
  - Ölçeklenebilirlik
    - TB'larca veri
  - Performans
    - GB lar büyüklüğünde dosyalar
  - Güvenilirlik
    - 1000 lerce makinada hata toleransı



# Hadoop Nedir

-  Apache Software Organization (ASF) Top Level Project
- Open Source software for reliable, scalable, distributed computing
- Core, Dağıtık Dosya Sistemi(HDFS) ve MapReduce
- Commodity Hardware



# Hadoop Tarihçesi

- Feb 2003 – First MapReduce library written at Google
- Oct 2003 – Google File System paper published
- Dec 2004 – Google MapReduce paper published
- Jul 2005 – Doug Cutting reports that Nutch now uses new MapReduce implementation
- Nov 2006 – Google Bigtable paper published
- Feb 2006 – Hadoop code moves out of Nutch into new Lucene sub-project
- Feb 2007 – First HBase code drop from Mike Cafarella
- Apr 2007 – Yahoo! running Hadoop on 1000-node cluster
- Jan 2008 – Hadoop made an Apache Top Level Project
- Apr 2008 – Wins Terabyte sort benchmark
- Aug 2008 – Yahoo! ran 4000 node Hadoop cluster



# Hadoop Kullanıcıları

- <http://wiki.apache.org/hadoop/PoweredBy>
- Yahoo (20K node, biggest cluster: 2K nodes)
- Facebook (1K nodes)
- IBM/Google Cloud Computing Initiative(100 nodes)
- CMU/Yahoo supercomputing cluster
- A9, AOL, Joost, Last.fm, NYT, Ning, Baidu, ...



# Yahoo Clusters



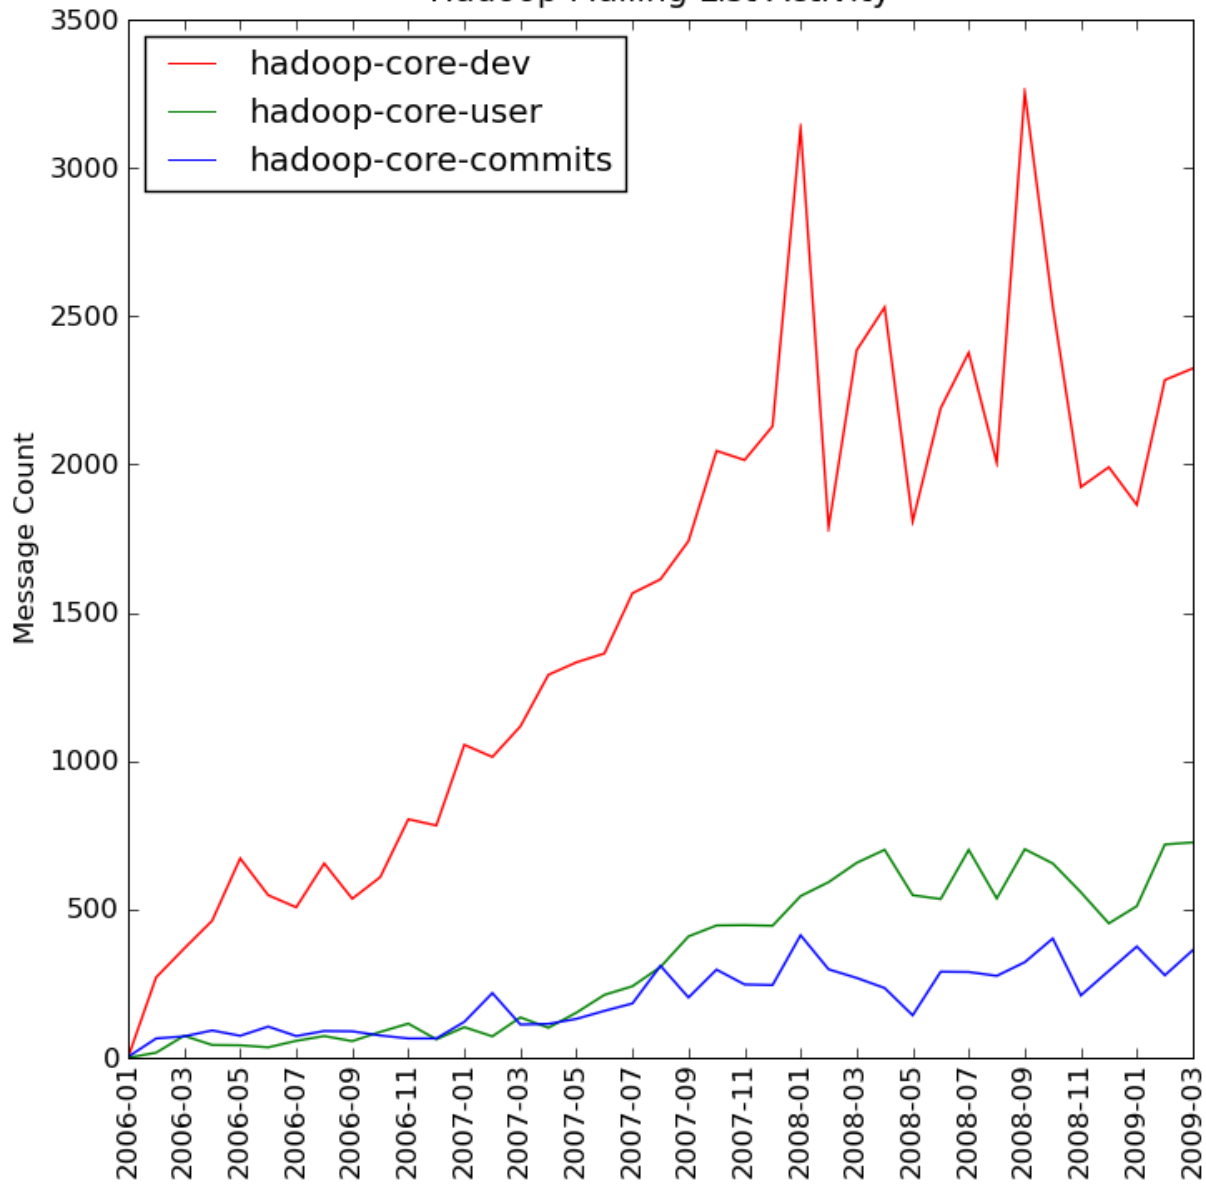
- 20.000 Node
- En büyük 4000 node
- Bir kaç PetaByte
- Aylık yüzbinlerce Hadoop görevi



# Çözüm

- Google File System (GFS) 2003 Paper
- Google MapReduce 2004 Paper
- Hadoop
  - HDFS
  - MapReduce

Hadoop Mailing List Activity



[www.cloudera.com](http://www.cloudera.com) 'dan alınmıştır



# Hadoop Dağıtık Dosya Sistemi

- Amaçlar:
  - Hata toleransı
  - Batch processing
  - Büyük veri kümeleri
  - Write once veri erişimi (append yakında)
  - Veri transferi yerine hesaplama transferi




# Hadoop Dağıtık Dosya Sistemi

- Master / Worker yapısı
- Dosya isimleri için tek isim-uzayı
- Write, append, read, delete, rename
- Non - Posix
- Dosya blokları
- Blok replikasyonu
- Java, C, Thrift, Shell, http, FUSE API



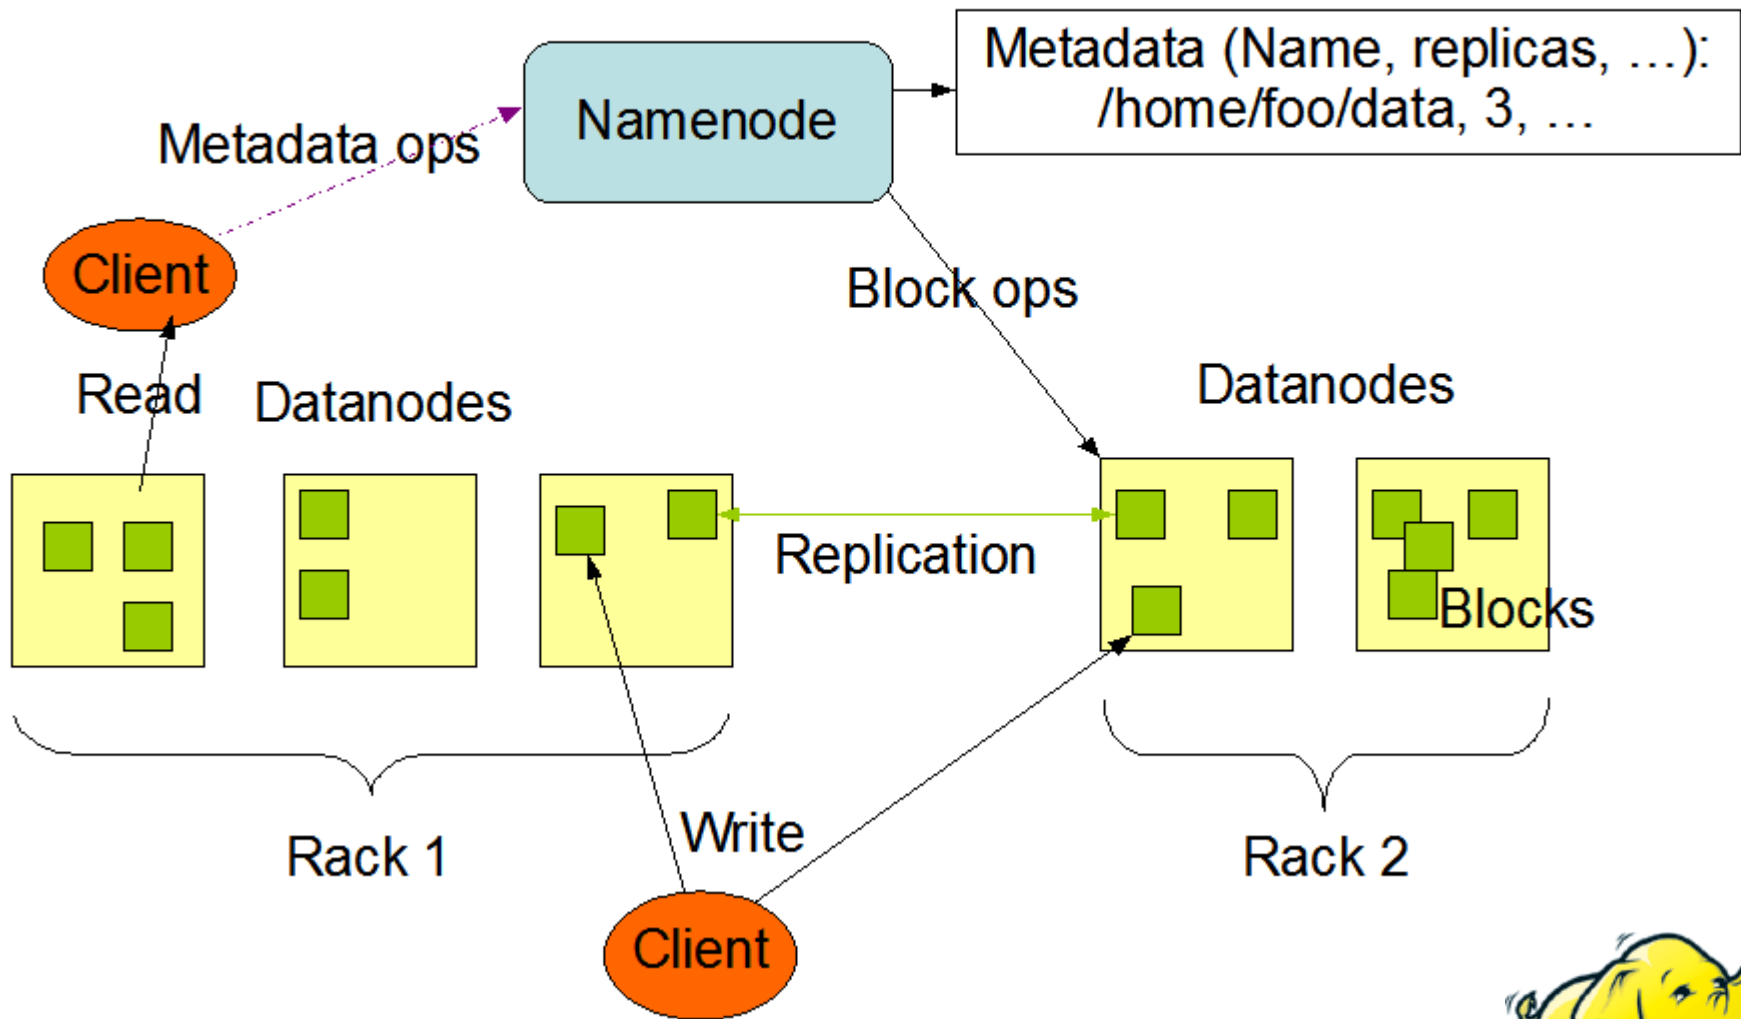
# Konfigurasyon



- 2 - 4 CPU, 16 GB Ram, 1 TB HD PC
- 40 Node rack
- Ağaç Topolojisi




# 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}, ...

### Datanodes



# Blok Yerleşimi

- Bloklar öntanımlı olarak 64 MB
- Öntanımlı 3 kopya
- Topolojiye göre blok yerleşimi
  - Aynı node/rack
- İstemci veriyi datanode'lara direk iletiyor
- Balancer
- Checksum ile veri doğruluğu





# Hadoop Dosya Sistemleri

- Hadoop birden fazla dosya sistemi ile çalışabilir
  - Local
  - Memory
  - Kosmos FS
  - HDFS
  - Amazon S3



# MapReduce

## map in lisp :

- The map function takes a function and a list and returns the result of applying the function to successive elements.

```
> (map (fn (x) (+ x 10)) '(1 2 3))  
(11 12 13)
```

## reduce in lisp:

- The reduce function combines all the elements of a sequence using a binary operation; for example, using + one can add up all the elements.

```
> (reduce '+ '(1 2 3 4))  
10
```



# Hadoop MapReduce

- Map

$(K1, V1) \rightarrow \text{list}(K2, V2)$

- Reduce

$(K2, \text{list}(K2)) \rightarrow \text{list}(K3, V3)$

- MapReduce programı :

$\text{list}(K1, V1) \rightarrow \text{list}(K3, V3)$



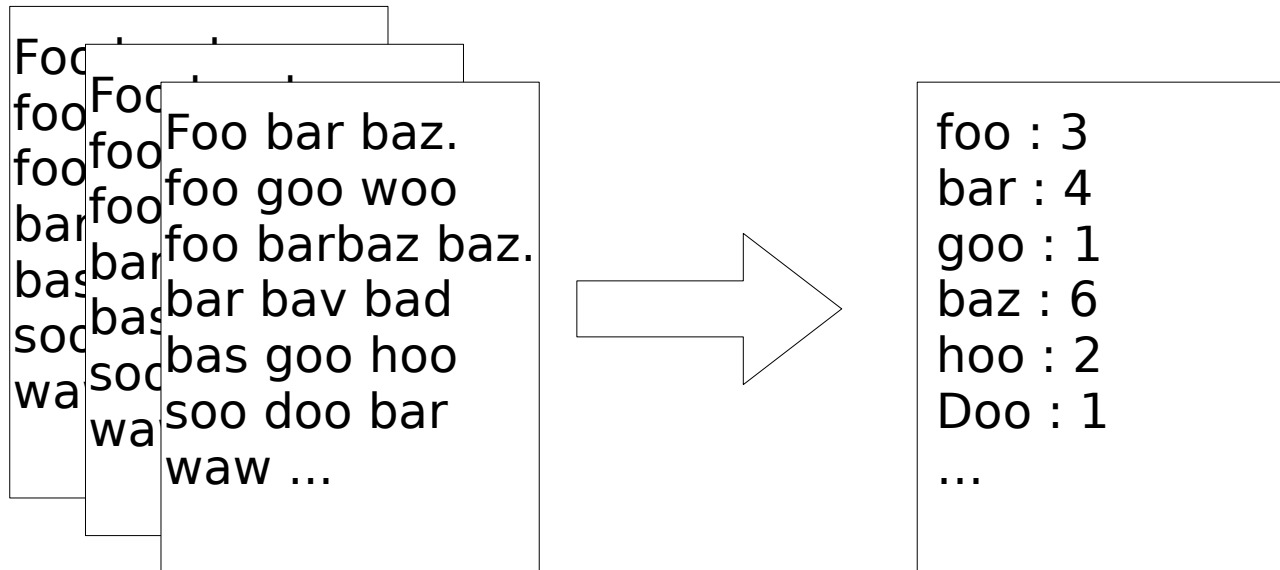
# Hadoop MapReduce Özellikler

- Büyük miktarda veriyi parçalar halinde, paralel işler
- Basitleştirilmiş key/value arayüzü ile herkesin dağıtık alg. yazabilmesi
- Batch streaming I/O
- Hata toleranslı
- Yüksek data locality
- Key/value tabanlı Map/Reduce modeli data mining / machine learning / Web crawl gibi pek çok alana uygun



# Word Count

- Bir Dosya Kümesinde geçen tüm kelimeleri saymak.



# Word Count Pseudo Code

```
mapper (filename, file-contents):  
    for each word in file-contents:  
        emit (word, 1)
```

```
reducer (word, values):  
    sum = 0  
    for each value in values:  
        sum = sum + value  
    emit (word, sum)
```



# Word Count Mapper

```
public static class TokenizerMapper
    extends Mapper<Object, Text, Text, IntWritable>{

    private final static IntWritable one = new IntWritable(1);
    private Text word = new Text();

    public void map(Object key, Text value, Context context
        ) throws IOException, InterruptedException {
        StringTokenizer itr = new StringTokenizer(value.toString());
        while (itr.hasMoreTokens()) {
            word.set(itr.nextToken());
            context.write(word, one);
        }
    }
}
```



# Word Count Reducer

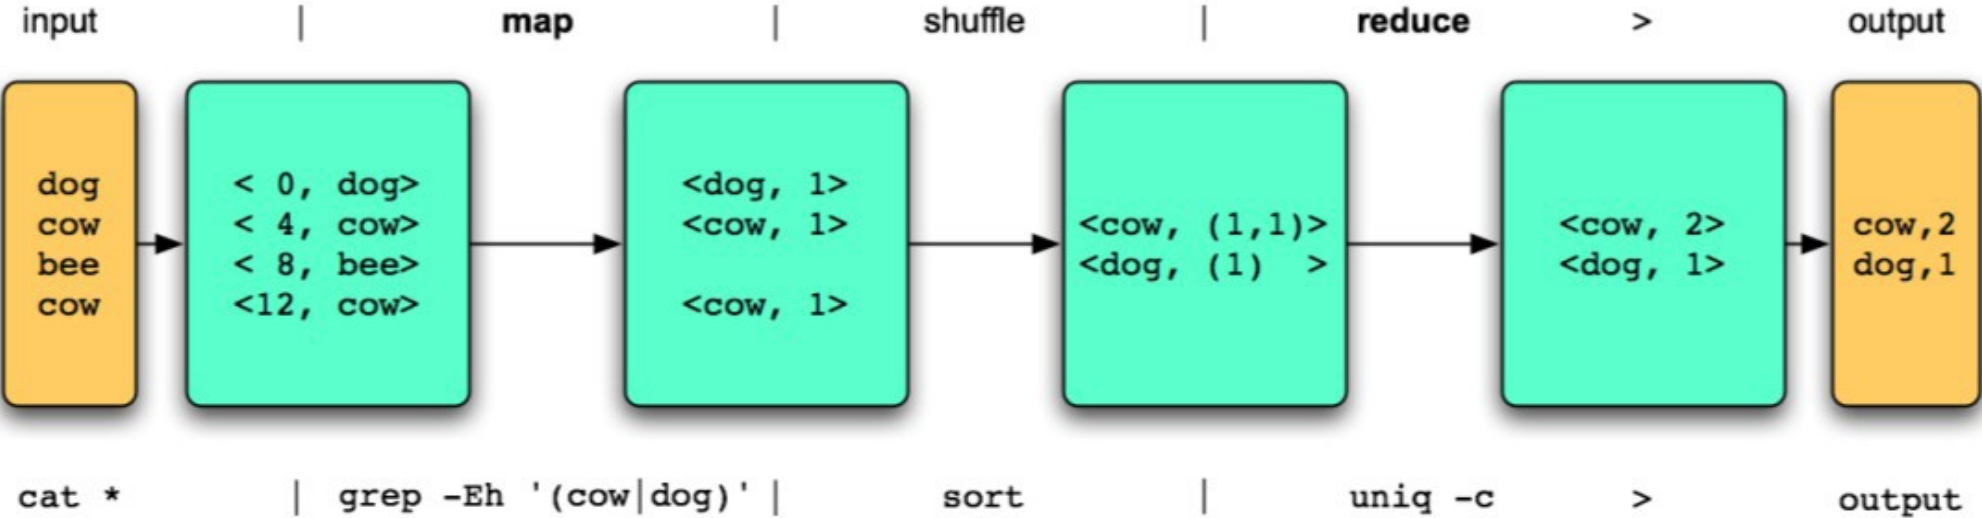
```
public static class IntSumReducer
    extends Reducer<Text,IntWritable,Text,IntWritable> {
    private IntWritable result = new IntWritable();

    public void reduce(Text key, Iterable<IntWritable> values,
        Context context
        ) throws IOException, InterruptedException {
        int sum = 0;
        for (IntWritable val : values) {
            sum += val.get();
        }
        result.set(sum);
        context.write(key, result);
    }
}
```

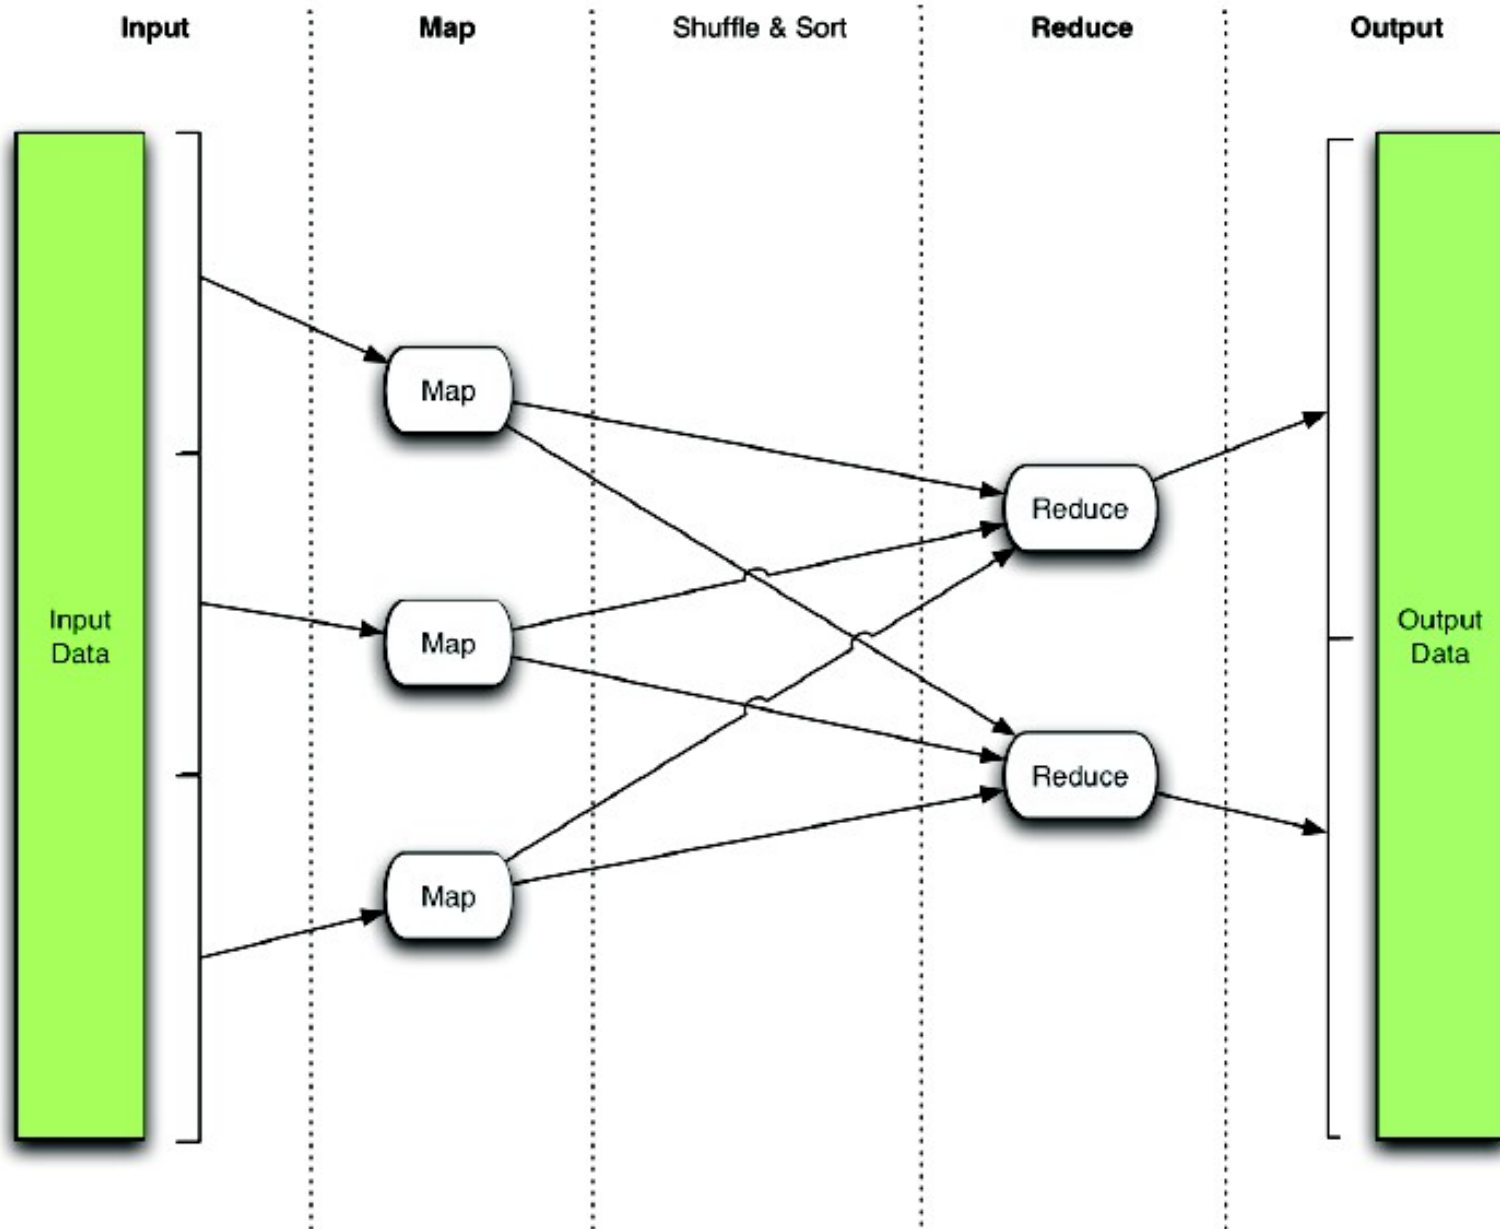


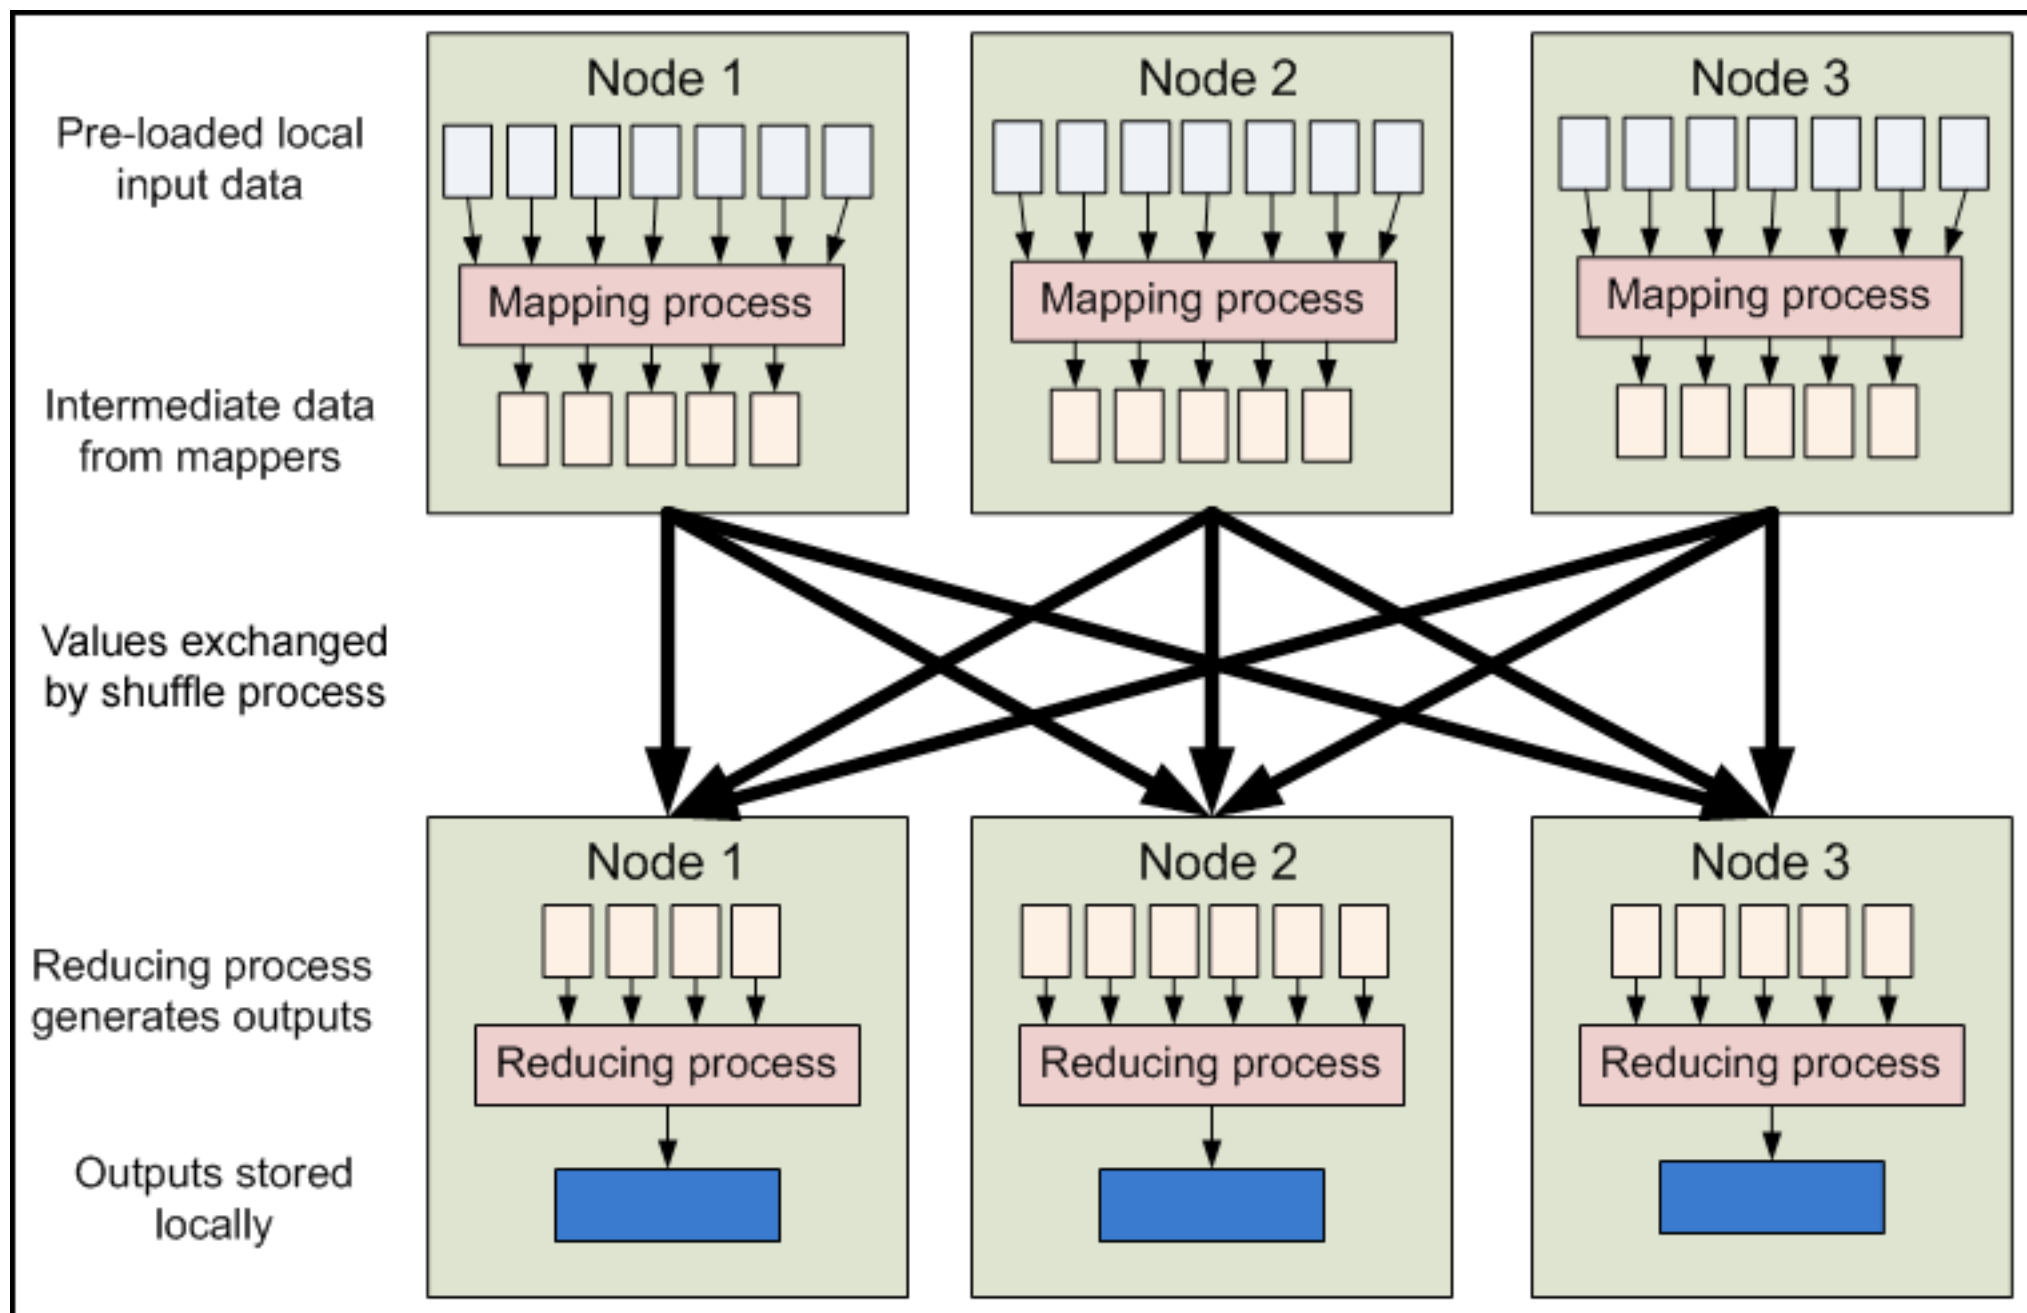


# Mantıksal Akış



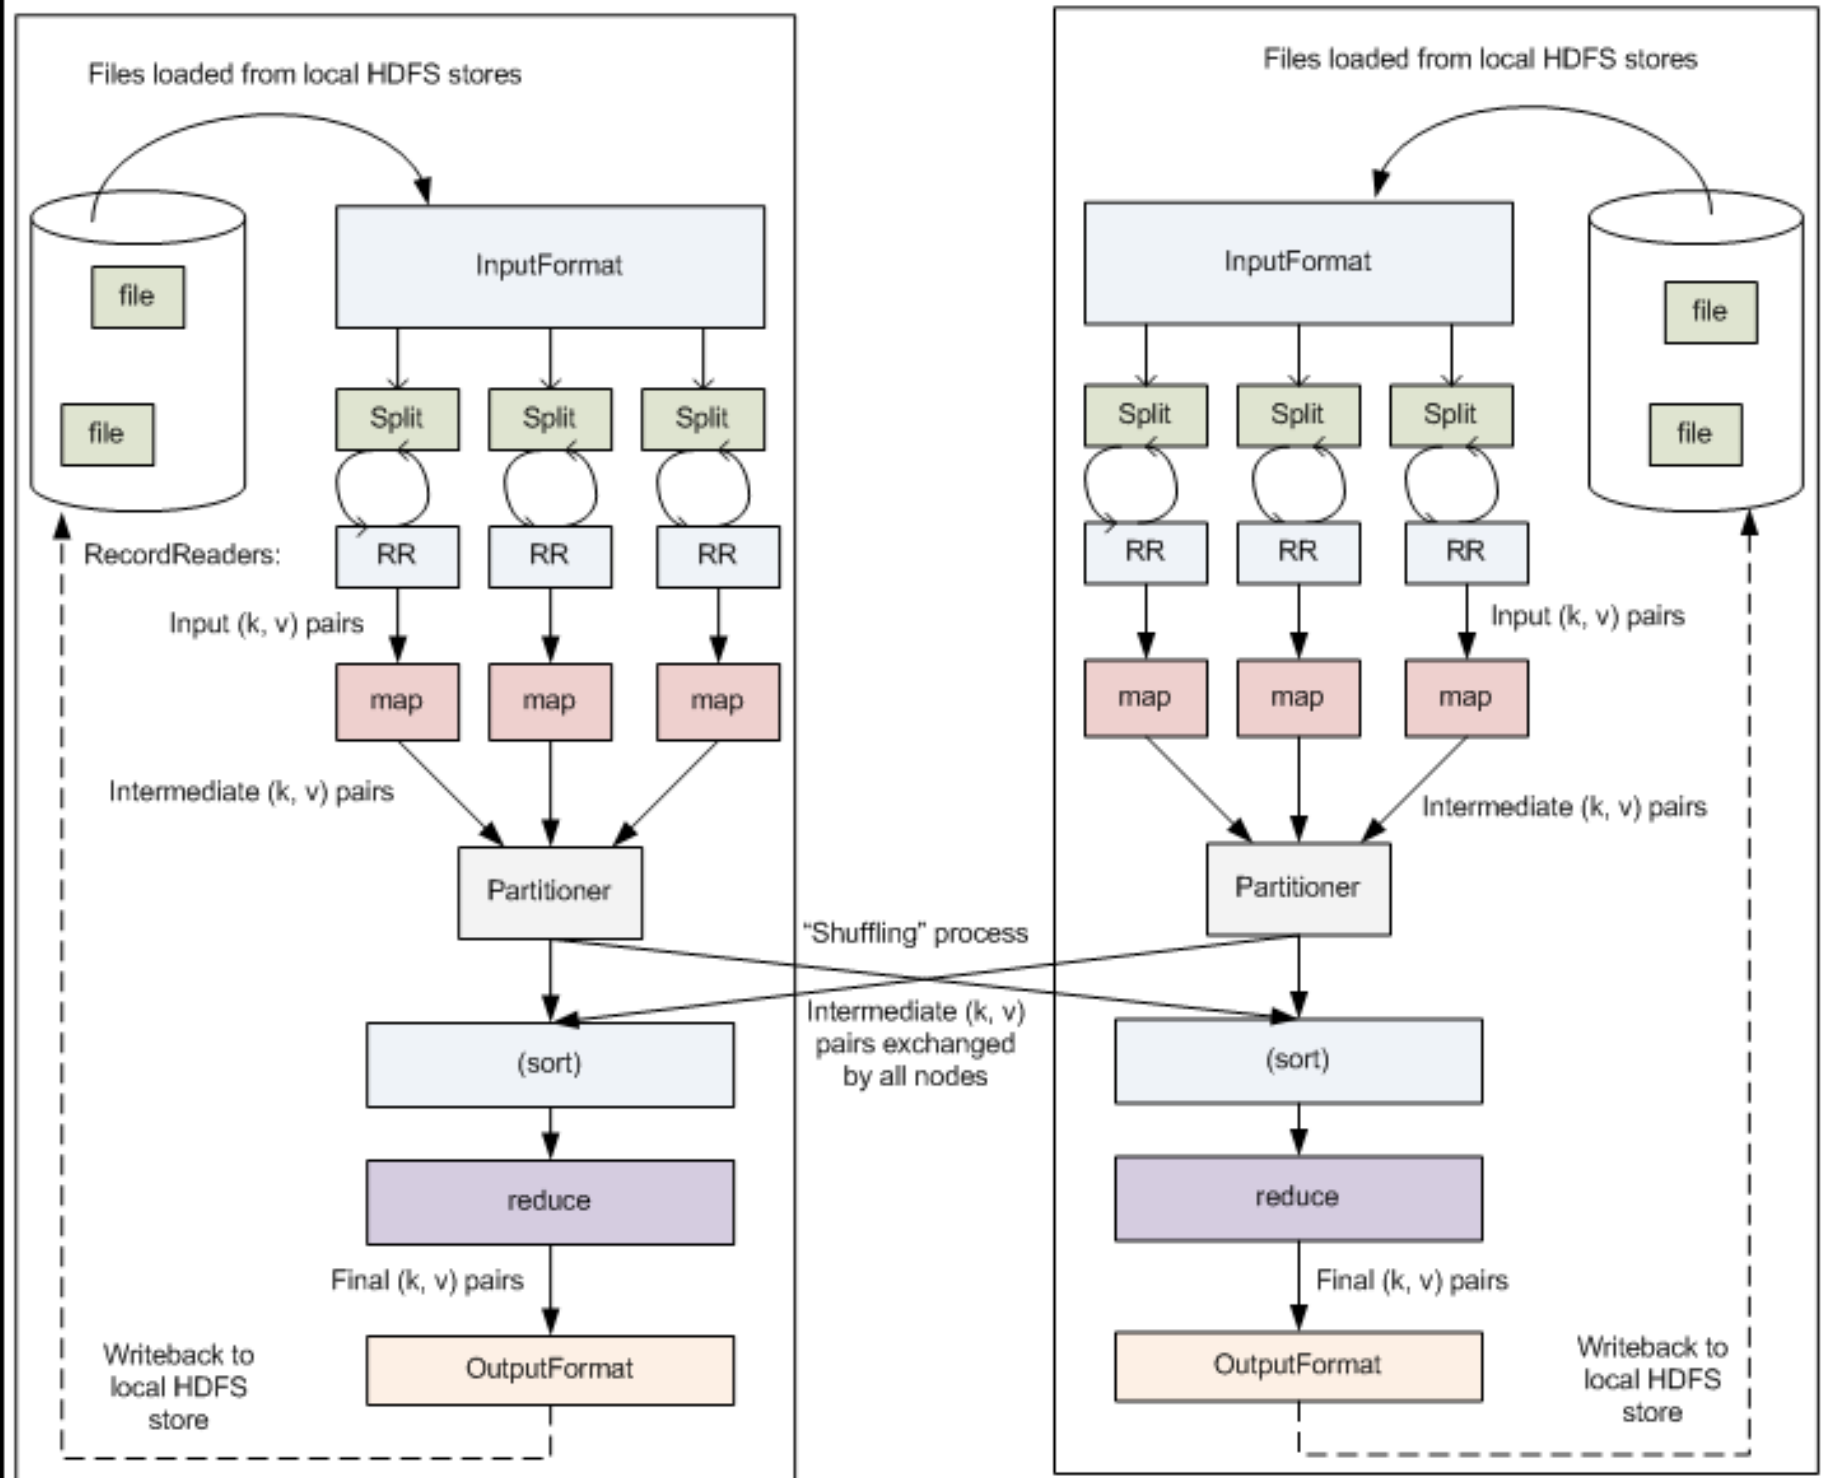
# Veri Akışı





# Node 1

# Node 2



# MapReduce Örnek Algoritmalar

- Distributed Grep
- Web access statistics
- Reverse web link graph
- Compute inverted index
- Find near-duplicate documents
- Sort - Terasort
- Machine learning ( collaborative filtering, EM, recommendation, )



# MapReduce Uygulama Alanları

- Yahoo WebMap
- Nutch – Internet crawl / parse / index / find duplicates
- Facebook – Datawarehouse with SQL-like interface
- Amazon Elastic MapReduce
- Log Processing / Reporting / Analysis



# İlgili Projeler



Hbase



Pig



Zookeeper



Hive



Mahout

Cascading



(Note: Hive logo is currently under discussion)

# Sorular

- Daha fazla bilgi için
  - Hadoop Web sitesi : <http://hadoop.apache.org>
  - Hadoop mailing listeleri
    - [core-users@hadoop.apache.org](mailto:core-users@hadoop.apache.org)
    - [core-dev@hadop.apache.org](mailto:core-dev@hadop.apache.org)
  - Hadoop wiki : <http://wiki.apache.org/hadoop>

