

parser@cs.msu.su, kuzloc@ispras.ru

MapReduce

storm

: NoSQL,

; nosql; mapreduce; hadoop; s4;

1.

0.18 , 2011 IDC « » 2006
5 ! 1.8
[1]:

1. -
2. Facebook
3. 500 Internet Archive 2
4. 20 !

2.

« »?

« »? Gartner: ««
» [2].

1.

IT

2.

3.

4.

3.

3.1.

« ?» - , , .

1. : , , — ().
2. — , , .
3. .
4. / — , , .
5. , , — .
6. , « ?» « ?».

1. .

2. .
3. SQL .
4. .
5. .

3.2. NoSQL

— , NoSQL. «NoSQL» 1998 SQL [3]. 2009 , NoSQL [4,5]:

1. ACID
2. NoSQL

3. Hypertable, Zvent, Google Bigtable, Bigtable, 20 [6].

3. NoSQL

4. NoSQL, CAP. (availability), (partition tolerance),

3.2.1 *NoSQL*

NoSQL

(Rick Cattell) [7]:

1.

2.

« ».

147

3. < , , > [6].

4. (,).

-	Redis Scalaris Riak Tokyo Tyrant
	SimpleDB CouchDB MongoDB
	BigTable Hbase Cassandra
	Neo4j

. 1. *NoSQL*

3.3. MapReduce

2003 .

2004 . Google, GFS [8], MapReduce [9]. Apache

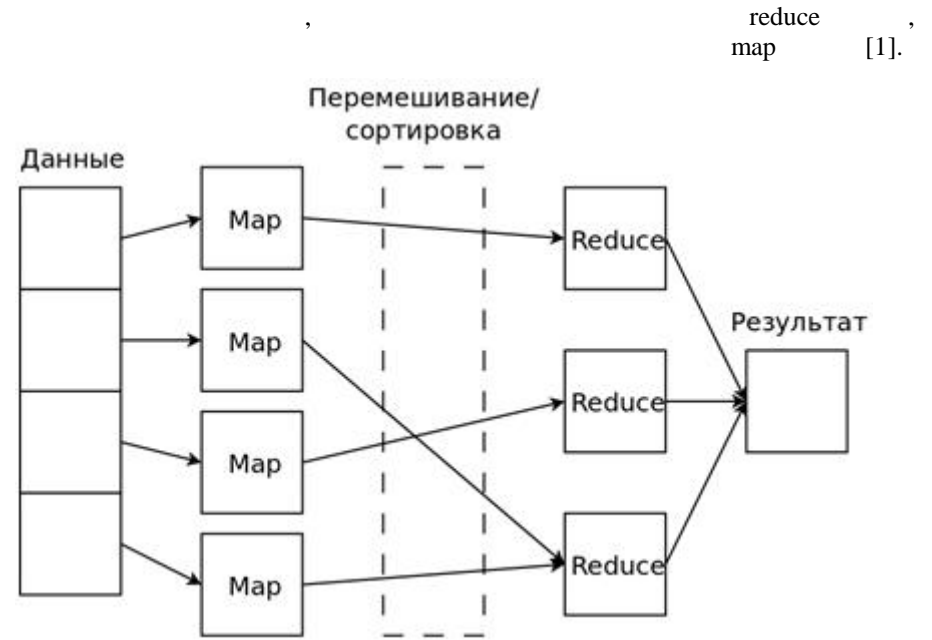
Nutch Hadoop [10],

« ».

148

3.3.1 Hadoop

« Hadoop MapReduce NoSQL [11, 12, 13].
 ~100 / 2.5
 1
 2 NoSQL
 ?
 NoSQL (CouchDB, MongoDB) B- [14].
 Hadoop.
 Hadoop HDFS [15].
 :
 1. HDFS
 2. HDFS
 3. HDFS
 Hadoop map
 reduce
 Map



1. MapReduce

Hadoop 100
 [16], 90
 10
 1
 Yahoo!
 Google — 68
 Hadoop 2009
 62 ! [17]

- Hadoop :
1. Hive — HDFS HiveQL, MapReduce SQL. [18].
 2. Pig — Hadoop. Pig [19].
 3. Hbase — Hbase

4. ZooKeeper —

[20].

Hadoop

[21].

[27].

3.3.2 Disco

1.

. MapReduce

2.

: map reduce.

3.

[22].

4.

MapReduce

[23]

Nokia,

Disco [24].

Dicso

Erlang [25],

Ericsson,

1.

99.999999% [26].

Hadoop, Disco

2.

S4,

2012 .

S4

(PE),

3.4.

S4

PE,

Hadoop,

S4,

«Hadoop

S4 (Simple Scalable Streaming System)».

Yahoo!.

3.4.1 Yahoo! S4

[27].

S4

S4
Hadoop

CTR

CTR

9.7 /c.

[28]
S4.

1.

2.

3.

S4

4.

UDP

3.4.2 Storm

Storm [29]

Twitter

S4.
Storm,
Yahoo!.

Storm

:

Storm

1.

(RPC)

2.

Storm

3.

Storm,

S4,

4.

, Storm


5.

Storm

(. 2).

(spout)

(bolt).



. 2. Storm.

Storm

. Storm

Storm

4.

: NoSQL, MapReduce

[1] Tom White. Hadoop: The Definitive Guide, 3rd Edition. O'Reilly Media, 2012, 688 p.

[2] Mark A. Beyer, Douglas Laney. The Importance of «Big Data»: A Definition. <http://www.gartner.com/DisplayDocument?id=2057415>, 21 June 2012.

[3] Carlo Strozzi. NoSQL: A Relational Database Management System. http://www.strozzi.it/cgi-bin/CSA/tw7/I/en_US/nosql/Home%20Page

[4] Jaroslav Pokorny. NoSQL databases: a step to database scalability in web environment. Proceedings of the 13th International Conference on Information Integration and Web-based Applications and Services, p. 278-283, ACM New York, NY, USA, 2011.

[5] Christof Strauch. NoSQL Databases. <http://www.christof-strauch.de/nosql dbs.pdf>

[6] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, Robert E. Gruber. Bigtable: a distributed storage system for structured data. Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation, vol. 7, p. 15-15, USENIX Association Berkeley, CA, USA, 2006.

[7] Rick Cattell. Scalable SQL and NoSQL data stores. ACM SIGMOD Record, 39(4), p. 12-27, ACM New York, NY, USA, December 2010.

[8] Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung. The Google File System. 19th ACM Symposium on Operating Systems Principles, Lake George, NY, October, 2003.

[9] Jeffrey Dean, Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation, vol. 6, p. 10-10, USENIX Association Berkeley, CA, USA, 2004.

[10] Apache Hadoop. <http://hadoop.apache.org/>

[11] Apache CouchDB. <http://couchdb.apache.org/>

[12] MongoDB. <http://www.mongodb.org/>

[13] Riak. <http://basho.com/products/riak-overview/>

[14] J. Chris Anderson, Jan Lehnardt, Noah Slater. CouchDB: The Definitive Guide. O'Reilly Media, 2010, 272 p.

[15] Konstantin Shvachko, Hairong Kuang, Sanjai Radia, Robert Chansler. The Hadoop Distributed File System. MSST '10 Proceedings of the 2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST), 2010, pp. 1-10.

[16] Sort Benchmark Home Page. <http://sortbenchmark.org/>

[17] Ajay Anand. Hadoop Sorts a Petabyte in 16.25 Hours and a Terabyte in 62 Seconds. http://developer.yahoo.com/blogs/hadoop/posts/2009/05/hadoop_sorts_a_petabyte_in_162/, 2009.

[18] Apache Hive. <http://hive.apache.org/>

[19] Apache Pig. <http://pig.apache.org/>

[20] Apache Hbase. <http://hbase.apache.org/>

[21] P. Hunt, M. Konar, F. P. Junqueira, and B. Reed. ZooKeeper: wait-free coordination for internet-scale systems. USENIXATC'10: Proceedings of the 2010 USENIX conference on USENIX annual technical conference. Berkeley, CA, USA: USENIX Association, 2010, pp. 11-11.

[22] . « » ,
02, 2012.

[23] G. Agha. Actors: A Model of Concurrent Computation in Distributed Systems. Cambridge, MA, USA: MIT Press, 1986.

[24] The Disco Project. <http://discoproject.org/>

[25] Erlang Programming Language. <http://www.erlang.org/>

[26] Joe Armstrong. Concurrency Oriented Programming in Erlang. <http://ll2.ai.mit.edu/talks/armstrong.pdf>, November 2002.

[27] Leonardo Neumeyer, Bruce Robbins, Anish Nair, Anand Kesari. S4: Distributed Stream Computing Platform. Data Mining Workshops (ICDMW), 2010 IEEE International Conference, 2010.

[28] Jagmohan Chauhan, Shaiful Chowdhury and Dwight Makaroff, Performance Evaluation of Yahoo! S4: A First Look, IEEE Seventh International Conference on P2P, Parallel, GRID, Cloud and Internet computing, 2012.

[29] Storm: Distributed and Fault-tolerant realtime computation. <http://storm-project.net/>

Big data: modern approaches to storage and analysis

Pavel Klemenkov
MSU, Moscow, Russia
parser@cs.msu.su
Sergey Kuznetsov
ISP RAS, Moscow, Russia
kuzloc@ispras.ru

Abstract: Big data challenged traditional storage and analysis systems in several new ways. In this paper we try to figure out how to overcome this challenges, why it's not possible to make it efficiently and describe three modern approaches to big data handling: NoSQL, MapReduce and real-time stream processing. The first section of the paper is the introduction. The second section discuss main issues of Big Data: volume, diversity, velocity, and value. The third section describes different approaches to solving the problem of Big Data. Traditionally one might use a relational DBMS. The paper propose some steps that allow to continue RDBMS using when it's capacity becomes not enough. Another way is to use a NoSQL approach. The basic ideas of the NoSQL approach are: simplification, high throughput, and unlimited scaling out. Different kinds of NoSQL stores allow to use such systems in different applications of Big Data. MapReduce and it's free implementation Hadoop may be used to provide scaling out Big Data analytics. Finally, several data management products support real time stream processing under Big Data. The paper briefly overviews these products. The final section of the paper is the conclusion.

Keywords: big data; relational model; nosql; mapreduce; hadoop; s4; storm

References

- [1]. Storm: Distributed and Fault-tolerant realtime computation. <http://storm-project.net/>
- [2]. Tom White. Hadoop: The Definitive Guide, 3rd Edition. O'Reilly Media, 2012, 688 p.
- [3]. Mark A. Beyer, Douglas Laney. The Importance of «Big Data»: A Definition. <http://www.gartner.com/DisplayDocument?id=2057415>, 21 June 2012.
- [4]. Carlo Strozzi. NoSQL: A Relational Database Management System. http://www.strozzi.it/cgi-bin/CSA/tw7/1/en_US/nosql/Home%20Page
- [5]. Jaroslav Pokorny. NoSQL databases: a step to database scalability in web environment. Proceedings of the 13th International Conference on Information Integration and Web-based Applications and Services, p. 278-283, ACM New York, NY, USA, 2011.
- [6]. Christof Strauch. NoSQL Databases. <http://www.christof-strauch.de/nosql dbs.pdf>
- [7]. Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, Robert E. Gruber. Bigtable: a distributed storage system for structured data. Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation, vol. 7, p. 15-15, USENIX Association Berkeley, CA, USA, 2006.
- [8]. Rick Cattell. Scalable SQL and NoSQL data stores. ACM SIGMOD Record, 39(4), p. 12-27, ACM New York, NY, USA, December 2010.
- [9]. Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung. The Google File System. 19th ACM Symposium on Operating Systems Principles, Lake George, NY, October, 2003.
- [10]. Jeffrey Dean, Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation, vol. 6, p. 10-10, USENIX Association Berkeley, CA, USA, 2004.
- [11]. Apache Hadoop. <http://hadoop.apache.org/>
- [12]. Apache CouchDB. <http://couchdb.apache.org/>
- [13]. MongoDB. <http://www.mongodb.org/>
- [14]. Riak. <http://basho.com/products/riak-overview/>
- [15]. J. Chris Anderson, Jan Lehnardt, Noah Slater. CouchDB: The Definitive Guide. O'Reilly Media, 2010, 272 p.
- [16]. Konstantin Shvachko, Hairong Kuang, Sanjai Radia, Robert Chansler. The Hadoop Distributed File System. MSST '10 Proceedings of the 2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST), 2010, pp. 1-10.
- [17]. Sort Benchmark Home Page. <http://sortbenchmark.org/>
- [18]. Ajay Anand. Hadoop Sorts a Petabyte in 16.25 Hours and a Terabyte in 62 Seconds. http://developer.yahoo.com/blogs/hadoop/posts/2009/05/hadoop_sorts_a_petabyte_in_1625_hours_and_a_terabyte_in_62_seconds/, 2009.
- [19]. Apache Hive. <http://hive.apache.org/>
- [20]. Apache Pig. <http://pig.apache.org/>
- [21]. Apache Hbase. <http://hbase.apache.org/>
- [22]. P. Hunt, M. Konar, F. P. Junqueira, and B. Reed. ZooKeeper: wait-free coordination for internet-scale systems. USENIXATC'10: Proceedings of the 2010 USENIX conference on USENIX annual technical conference. Berkeley, CA, USA: USENIX Association, 2010, pp. 11–11.
- [23]. Sergej Kuznetsov. K svobode ot problemy Bol'shih Dannyh [Toward the freedom from the Big Data problem]. «Otkrytye sistemy» [Open Systems], 02, 2012 (in Russian).
- [24]. G. Agha. Actors: A Model of Concurrent Computation in Distributed Systems. Cambridge, MA, USA: MIT Press, 1986.
- [25]. The Disco Project. <http://discoproject.org/>
- [26]. Erlang Programming Language. <http://www.erlang.org/>
- [27]. Joe Armstrong. Concurrency Oriented Programming in Erlang. <http://il2.ai.mit.edu/talks/armstrong.pdf>, November 2002.
- [28]. Leonardo Neumeyer, Bruce Robbins, Anish Nair, Anand Kesari. S4: Distributed Stream Computing Platform. Data Mining Workshops (ICDMW), 2010 IEEE International Conference, 2010.
- [29]. Jagmohan Chauhan, Shaiful Chowdhury and Dwight Makaroff, Performance Evaluation of Yahoo! S4: A First Look, IEEE Seventh International Conference on P2P, Parallel, GRID, Cloud and Internet computing, 2012.
- [30]. Storm: Distributed and Fault-tolerant realtime computation. <http://storm-project.net/>