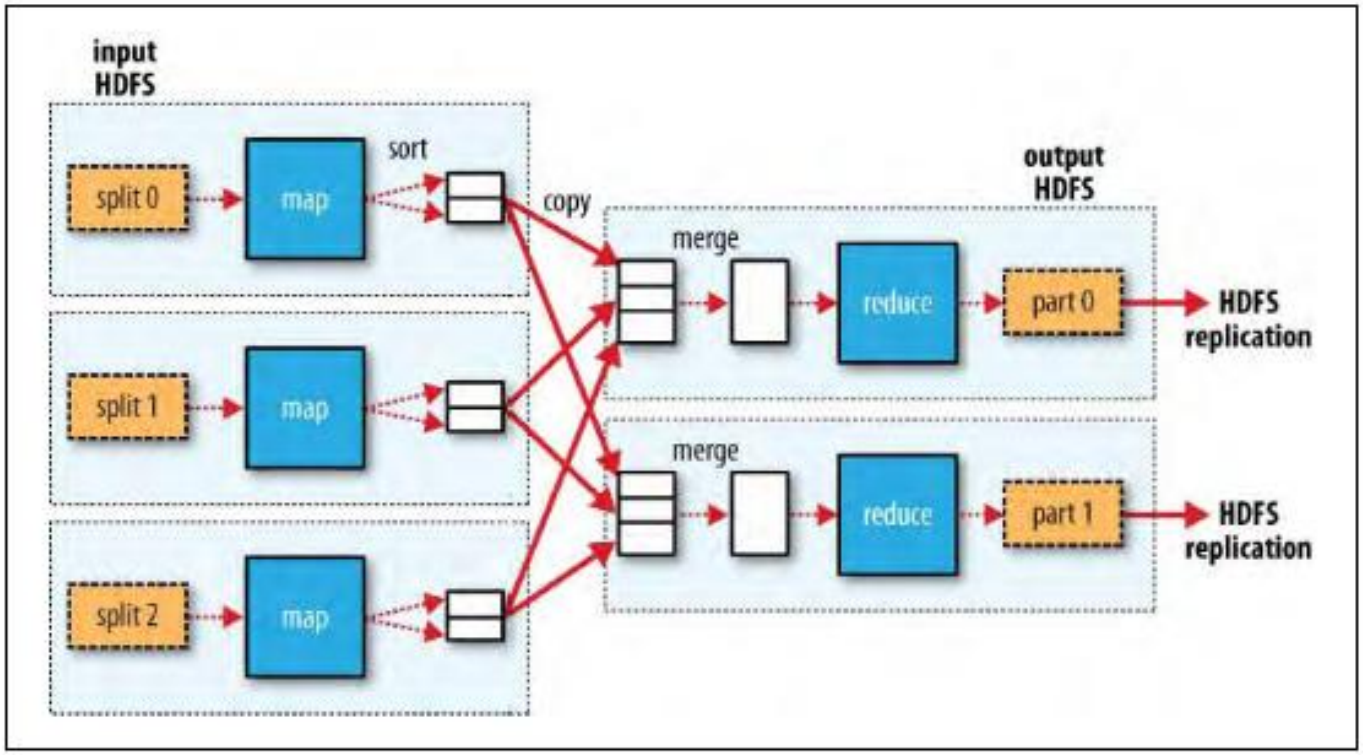

Hadoop数据分析平台 第3周

2012.9.3

Hello, World!

- 对刚安装好的hadoop集群做个测试

```
[grid@h1 ~]$
[grid@h1 ~]$ mkdir input
[grid@h1 ~]$ cd input
[grid@h1 input]$ echo "hello world" >test1.txt
[grid@h1 input]$ echo "hello hadoop" > test2.txt
[grid@h1 input]$ cat test1.txt
hello world
[grid@h1 input]$ cat test2.txt
hello hadoop
[grid@h1 input]$ cd ../hadoop-0.20.2
```

```
[grid@h1 hadoop-0.20.2]$ bin/hadoop dfs -put ../input in
[grid@h1 hadoop-0.20.2]$ bin/hadoop dfs -ls ./in/*\
>
-rw-r--r--   3 grid supergroup         12 2012-06-22 17:00 /user/grid/in/test1.txt
-rw-r--r--   3 grid supergroup         13 2012-06-22 17:00 /user/grid/in/test2.txt
```

测试

```
[grid@h1 hadoop-0.20.2]$ bin/hadoop jar hadoop-0.20.2-examples.jar wordcount in out
12/06/22 17:04:25 INFO input.FileInputFormat: Total input paths to process : 2
12/06/22 17:04:25 INFO mapred.JobClient: Running job: job_201206221659_0001
12/06/22 17:04:26 INFO mapred.JobClient:  map 0% reduce 0%
12/06/22 17:04:38 INFO mapred.JobClient:  map 100% reduce 0%
12/06/22 17:04:51 INFO mapred.JobClient:  map 100% reduce 100%
12/06/22 17:04:53 INFO mapred.JobClient: Job complete: job_201206221659_0001
12/06/22 17:04:53 INFO mapred.JobClient: Counters: 17
12/06/22 17:04:53 INFO mapred.JobClient:   Job Counters
12/06/22 17:04:53 INFO mapred.JobClient:     Launched reduce tasks=1
12/06/22 17:04:53 INFO mapred.JobClient:     Launched map tasks=2
12/06/22 17:04:53 INFO mapred.JobClient:     Data-local map tasks=2
12/06/22 17:04:53 INFO mapred.JobClient:   FileSystemCounters
12/06/22 17:04:53 INFO mapred.JobClient:     FILE_BYTES_READ=55
12/06/22 17:04:53 INFO mapred.JobClient:     HDFS_BYTES_READ=25
12/06/22 17:04:53 INFO mapred.JobClient:     FILE_BYTES_WRITTEN=180
12/06/22 17:04:53 INFO mapred.JobClient:     HDFS_BYTES_WRITTEN=25
12/06/22 17:04:53 INFO mapred.JobClient:   Map-Reduce Framework
12/06/22 17:04:53 INFO mapred.JobClient:     Reduce input groups=3
12/06/22 17:04:53 INFO mapred.JobClient:     Combine output records=4
12/06/22 17:04:53 INFO mapred.JobClient:     Map input records=2
12/06/22 17:04:53 INFO mapred.JobClient:     Reduce shuffle bytes=61
12/06/22 17:04:53 INFO mapred.JobClient:     Reduce output records=3
12/06/22 17:04:53 INFO mapred.JobClient:     Spilled Records=8
12/06/22 17:04:53 INFO mapred.JobClient:     Map output bytes=41
12/06/22 17:04:53 INFO mapred.JobClient:     Combine input records=4
12/06/22 17:04:53 INFO mapred.JobClient:     Map output records=4
12/06/22 17:04:53 INFO mapred.JobClient:     Reduce input records=4
[grid@h1 hadoop-0.20.2]$
```

测试结果

```
[grid@h1 hadoop-0.20.2]$ bin/hadoop dfs -ls
Found 2 items
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:00 /user/grid/in
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:04 /user/grid/out
[grid@h1 hadoop-0.20.2]$ bin/hadoop dfs -ls ./out
Found 2 items
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:04 /user/grid/out/_logs
-rw-r--r--   3 grid supergroup         25 2012-06-22 17:04 /user/grid/out/part-r-00000
[grid@h1 hadoop-0.20.2]$ bin/hadoop dfs -cat ./out/*
hadoop  1
hello   2
world   1
cat: Source must be a file.
[grid@h1 hadoop-0.20.2]$
```

通过web了解Hadoop的活动

- 通过用浏览器和http访问jobtracker所在节点的50030端口监控jobtracker
- 通过用浏览器和http访问namenode所在节点的50070端口监控集群

http://192.168.1.102:50030/jobtracker.jsp

h1 Hadoop Map/Reduce Administration

State: RUNNING
Started: Fri Jun 22 16:59:02 EDT 2012
Version: 0.20.2, r911707
Compiled: Fri Feb 19 08:07:34 UTC 2010 by chrisdo
Identifier: 201206221659

Cluster Summary (Heap Size is 15.31 MB/966.69 MB)

Maps	Reduces	Total Submissions	Nodes	Map Task Capacity	Reduce Task Capacity	Avg. Tasks/Node	Blacklisted Nodes
0	0	1	1	2	2	4.00	0

Scheduling Information

Queue Name	Scheduling Information
default	N/A

Filter (Jobid, Priority, User, Name)
Example: 'user:smith 3200' will filter by 'smith' only in the user field and '3200' in all fields

Jobtracker监控

Filter (Jobid, Priority, User, Name)
Example: 'user:smith 3200' will filter by 'smith' only in the user field and '3200' in all fields

Running Jobs

none

Completed Jobs

Jobid	Priority	User	Name	Map % Complete	Map Total	Maps Completed	Reduce % Complete	Reduce Total	Reduces Completed	Job Scheduling Information
job_201206221659_0001	NORMAL	grid	word count	100.00%	2	2	100.00%	1	1	NA

Failed Jobs

none

Local Logs

Log directory, Job Tracker History

Hadoop, 2012.

Jobtracker监控

Hadoop job_201206221659_0001 on h1

User: grid
Job Name: word count
Job File: hdfs://h1:9000/tmp/hadoop-grid/mapred/system/job_201206221659_0001/job.xml
Job Setup: Successful
Status: Succeeded
Started at: Fri Jun 22 17:04:25 EDT 2012
Finished at: Fri Jun 22 17:04:52 EDT 2012
Finished in: 26sec
Job Cleanup: Successful

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	100.00%	2	0	0	2	0	0 / 0
reduce	100.00%	1	0	0	1	0	0 / 0

	Counter	Map	Reduce	Total
Job Counters	Launched reduce tasks	0	0	1
	Launched map tasks	0	0	2
	Data-local map tasks	0	0	2
FileSystemCounters	FILE_BYTES_READ	0	55	55
	HDFS_BYTES_READ	25	0	25
	FILE_BYTES_WRITTEN	125	55	180

2012.9.3

http://192.168.1.102:50070/dfshealth.jsp

NameNode 'h1:9000'

Started: Fri Jun 22 16:58:58 EDT 2012
Version: 0.20.2, r911707
Compiled: Fri Feb 19 08:07:34 UTC 2010 by chrisdo
Upgrades: There are no upgrades in progress.

Browse the filesystem
Namenode Logs

Cluster Summary

17 files and directories, 11 blocks = 28 total. Heap Size is 15.31 MB / 966.69 MB (1%)

Configured Capacity	:	94.81 GB
DFS Used	:	100 KB
Non DFS Used	:	8.79 GB
DFS Remaining	:	86.02 GB
DFS Used%	:	0 %
DFS Remaining%	:	90.73 %
Live Nodes	:	1
Dead Nodes	:	0

NameNode Storage:

Storage Directory	Type	State
/home/grid/name	IMAGE_AND_EDITS	Active

观看日志

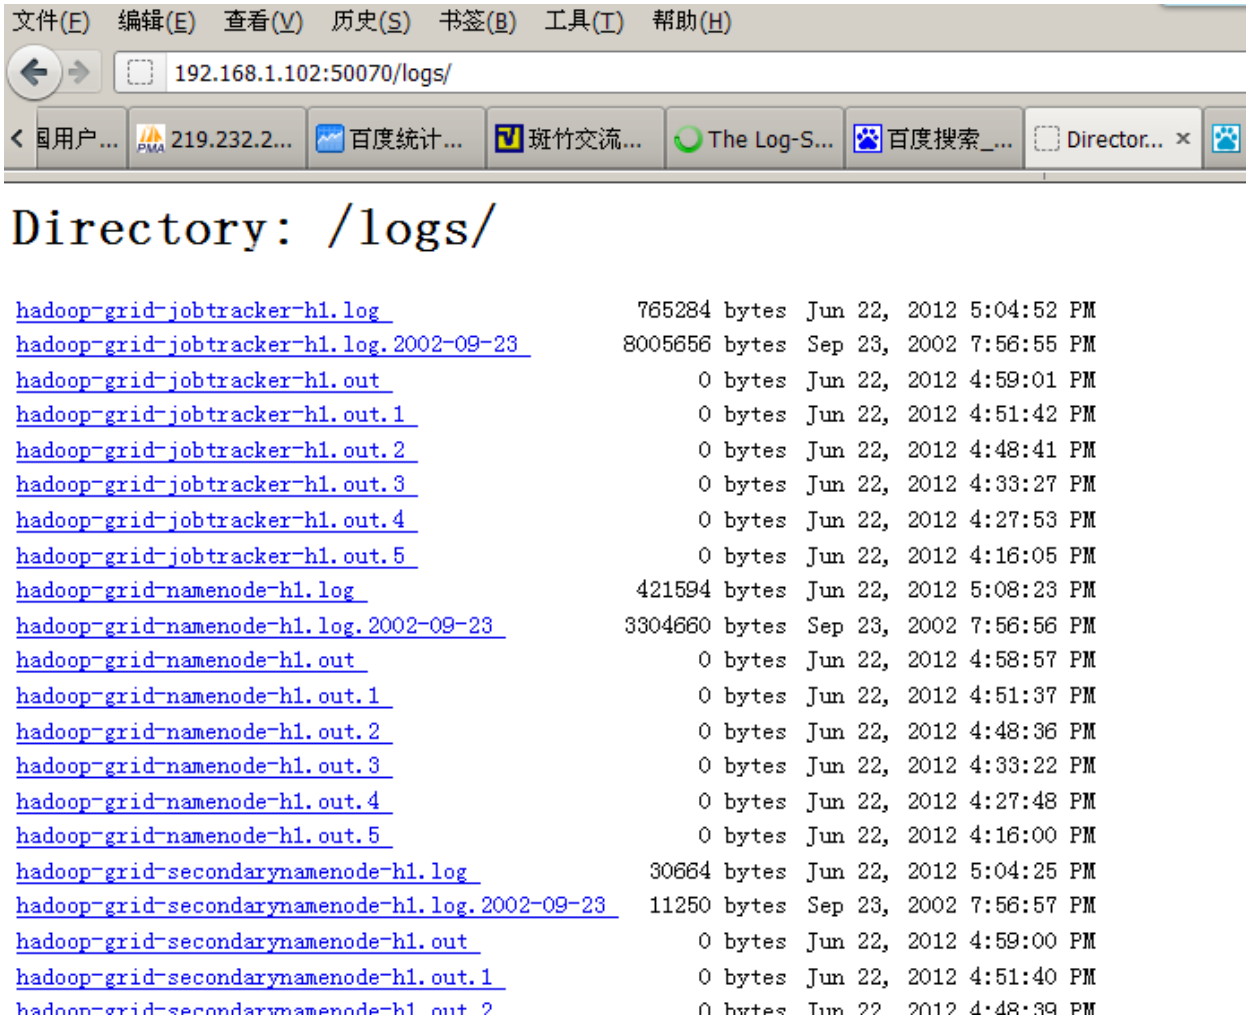
文件(F) 编辑(E) 查看(V) 历史(S) 书签(B) 工具(T) 帮助(H)
192.168.1.102:50070/logs/
219.232.2...
百度统计...
斑竹交流...
The Log-S...
百度搜索_...
Director...
Directory: /logs/
hadoop-grid-jobtracker-h1.log 765284 bytes Jun 22, 2012 5:04:52 PM
hadoop-grid-jobtracker-h1.log.2002-09-23 8005656 bytes Sep 23, 2002 7:56:55 PM
hadoop-grid-jobtracker-h1.out 0 bytes Jun 22, 2012 4:59:01 PM
hadoop-grid-jobtracker-h1.out.1 0 bytes Jun 22, 2012 4:51:42 PM
hadoop-grid-jobtracker-h1.out.2 0 bytes Jun 22, 2012 4:48:41 PM
hadoop-grid-jobtracker-h1.out.3 0 bytes Jun 22, 2012 4:33:27 PM
hadoop-grid-jobtracker-h1.out.4 0 bytes Jun 22, 2012 4:27:53 PM
hadoop-grid-jobtracker-h1.out.5 0 bytes Jun 22, 2012 4:16:05 PM
hadoop-grid-namenode-h1.log 421594 bytes Jun 22, 2012 5:08:23 PM
hadoop-grid-namenode-h1.log.2002-09-23 3304660 bytes Sep 23, 2002 7:56:56 PM
hadoop-grid-namenode-h1.out 0 bytes Jun 22, 2012 4:58:57 PM
hadoop-grid-namenode-h1.out.1 0 bytes Jun 22, 2012 4:51:37 PM
hadoop-grid-namenode-h1.out.2 0 bytes Jun 22, 2012 4:48:36 PM
hadoop-grid-namenode-h1.out.3 0 bytes Jun 22, 2012 4:33:22 PM
hadoop-grid-namenode-h1.out.4 0 bytes Jun 22, 2012 4:27:48 PM
hadoop-grid-namenode-h1.out.5 0 bytes Jun 22, 2012 4:16:00 PM
hadoop-grid-secondarynamenode-h1.log 30664 bytes Jun 22, 2012 5:04:25 PM
hadoop-grid-secondarynamenode-h1.log.2002-09-23 11250 bytes Sep 23, 2002 7:56:57 PM
hadoop-grid-secondarynamenode-h1.out 0 bytes Jun 22, 2012 4:59:00 PM
hadoop-grid-secondarynamenode-h1.out.1 0 bytes Jun 22, 2012 4:51:40 PM

2012.9.3

浏览文件系统

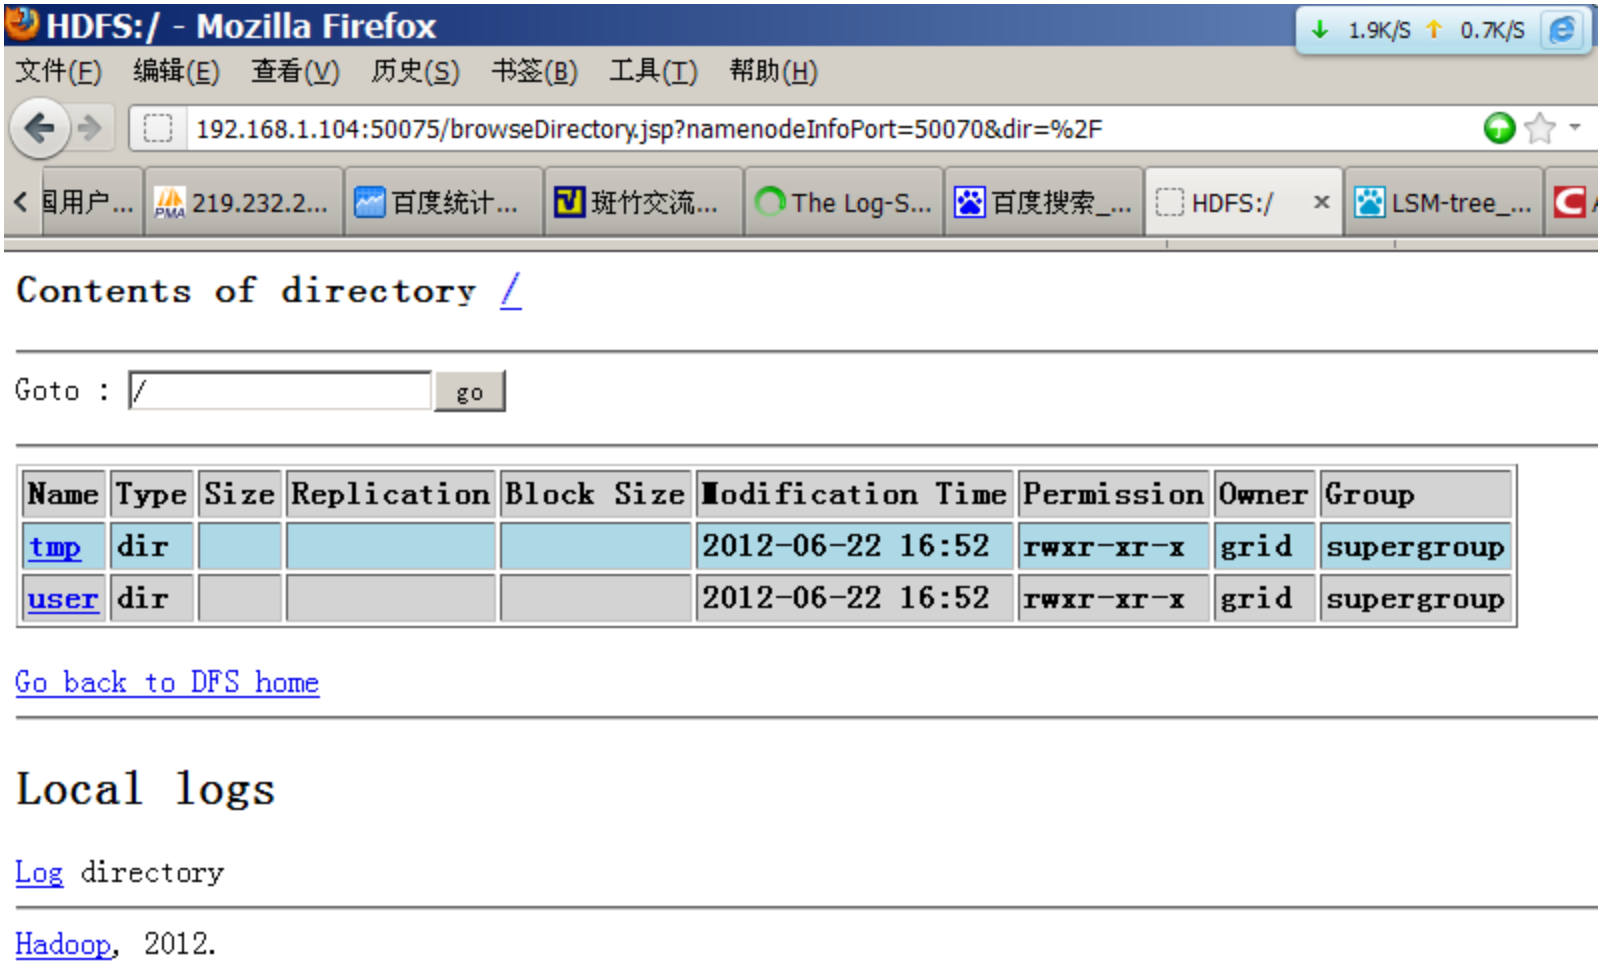
HDFS:/ - Mozilla Firefox
1.9K/S 0.7K/S
文件(F) 编辑(E) 查看(V) 历史(S) 书签(B) 工具(T) 帮助(H)
192.168.1.104:50075/browseDirectory.jsp?namenodeInfoPort=50070&dir=%2F
219.232.2...
百度统计...
斑竹交流...
The Log-S...
百度搜索_...
HDFS:/
LSM-tree_...
Contents of directory /
Goto : /
go
Name Type Size Replication Block Size Modification Time Permission Owner Group
tmp dir 2012-06-22 16:52 rwxr-xr-x grid supergroup
user dir 2012-06-22 16:52 rwxr-xr-x grid supergroup
Go back to DFS home
Local logs
Log directory
Hadoop, 2012.

数据写在了哪儿（从OS看）

```
[grid@h3 data]$ ls -lR
.:
total 16
drwxrwxr-x 2 grid grid 4096 Jun 22 17:04 current
drwxrwxr-x 2 grid grid 4096 Jun 22 16:59 detach
-rw-rw-r-- 1 grid grid    0 Jun 22 16:59 in_use.lock
-rw-rw-r-- 1 grid grid  157 Jun 22 16:59 storage
drwxrwxr-x 2 grid grid 4096 Jun 22 17:04 tmp

./current:
total 80
-rw-rw-r-- 1 grid grid  8668 Jun 22 17:04 blk_1864853225322117619
-rw-rw-r-- 1 grid grid    75 Jun 22 17:04 blk_1864853225322117619_1037.meta
-rw-rw-r-- 1 grid grid 16666 Jun 22 17:04 blk_-2152022802549252155
-rw-rw-r-- 1 grid grid   139 Jun 22 17:04 blk_-2152022802549252155_1036.meta
-rw-rw-r-- 1 grid grid    25 Jun 22 17:04 blk_3251284937612942908
-rw-rw-r-- 1 grid grid    11 Jun 22 17:04 blk_3251284937612942908_1037.meta
-rw-rw-r-- 1 grid grid    13 Jun 22 17:00 blk_7499053175856941845
-rw-rw-r-- 1 grid grid    11 Jun 22 17:00 blk_7499053175856941845_1031.meta
-rw-rw-r-- 1 grid grid    12 Jun 22 17:00 blk_-7677903991784965860
-rw-rw-r-- 1 grid grid    11 Jun 22 17:00 blk_-7677903991784965860_1030.meta
-rw-rw-r-- 1 grid grid     4 Jun 22 16:59 blk_8963744517366784034
-rw-rw-r-- 1 grid grid    11 Jun 22 16:59 blk_8963744517366784034_1029.meta
-rw-rw-r-- 1 grid grid   770 Jun 22 17:13 dncp_block_verification.log.curr
-rw-rw-r-- 1 grid grid   157 Jun 22 16:59 VERSION
```

HDFS设计基础与目标

- 硬件错误是常态。因此需要冗余
- 流式数据访问。即数据批量读取而非随机读写，Hadoop擅长做的是数据分析而不是事务处理
- 大规模数据集
- 简单一致性模型。为了降低系统复杂度，对文件采用一次性写多次读的逻辑设计，即是文件一经写入，关闭，就再也不能修改
- 程序采用“数据就近”原则分配节点执行

HDFS体系结构

- NameNode
- DataNode
- 事务日志
- 映像文件
- SecondaryNameNode

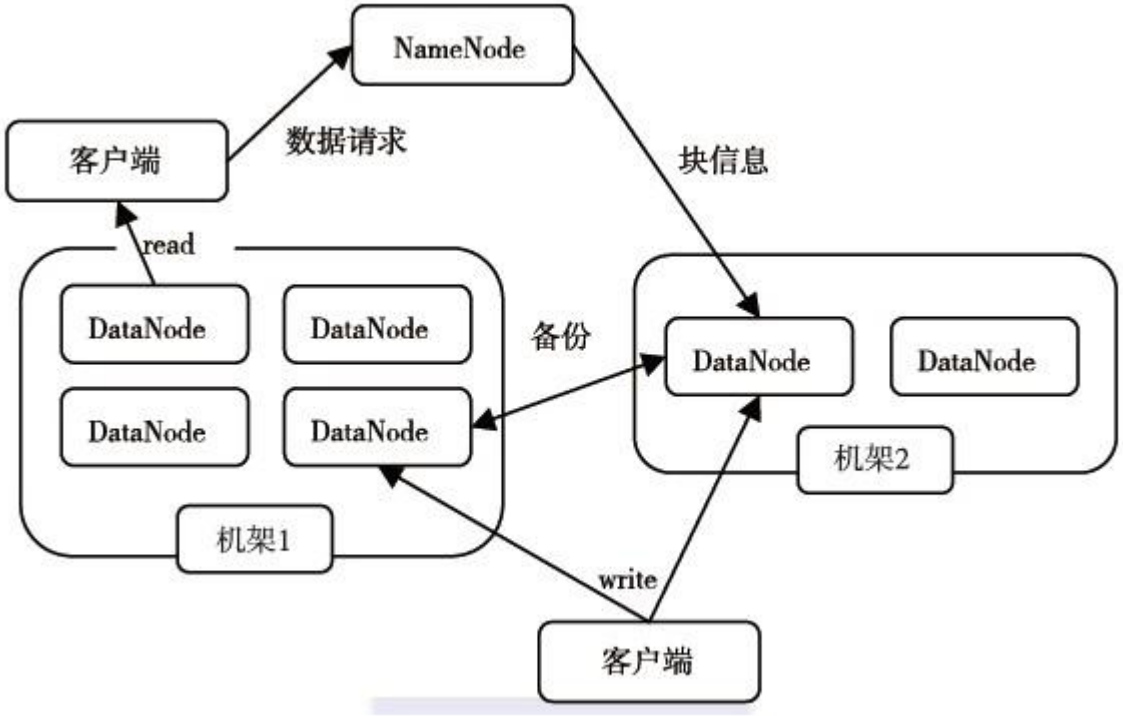

Namenode

- 管理文件系统的命名空间
- 记录每个文件数据块在各个Datanode上的位置和副本信息
- 协调客户端对文件的访问
- 记录命名空间内的改动或空间本身属性的改动
- Namenode使用事务日志记录HDFS元数据的变化。使用映像文件存储文件系统的命名空间，包括文件映射，文件属性等

```
[grid@h1 name]$ ls -Rl
.:
total 8
drwxrwxr-x 2 grid grid 4096 Jun 22 17:04 current
drwxrwxr-x 2 grid grid 4096 Jun 22 16:51 image
-rw-rw-r-- 1 grid grid    0 Jun 22 16:58 in_use.lock

./current:
total 24
-rw-rw-r-- 1 grid grid 1049092 Jun 22 17:04 edits
-rw-rw-r-- 1 grid grid     971 Jun 22 17:04 fsimage
-rw-rw-r-- 1 grid grid       8 Jun 22 17:04 fstime
-rw-rw-r-- 1 grid grid      99 Jun 22 17:04 VERSION

./image:
total 4
-rw-rw-r-- 1 grid grid 157 Jun 22 17:04 fsimage
[grid@h1 name]$
```

Datanode

- 负责所在物理节点的存储管理
- 一次写入，多次读取（不修改）
- 文件由数据块组成，典型的块大小是64MB
- 数据块尽量散布道各个节点

```
[grid@h3 data]$ ls -lR
.:
total 16
drwxrwxr-x 2 grid grid 4096 Jun 22 17:04 current
drwxrwxr-x 2 grid grid 4096 Jun 22 16:59 detach
-rw-rw-r-- 1 grid grid    0 Jun 22 16:59 in_use.lock
-rw-rw-r-- 1 grid grid  157 Jun 22 16:59 storage
drwxrwxr-x 2 grid grid 4096 Jun 22 17:04 tmp

./current:
total 80
-rw-rw-r-- 1 grid grid  8668 Jun 22 17:04 blk_1864853225322117619
-rw-rw-r-- 1 grid grid    75 Jun 22 17:04 blk_1864853225322117619_1037.meta
-rw-rw-r-- 1 grid grid 16666 Jun 22 17:04 blk_-2152022802549252155
-rw-rw-r-- 1 grid grid   139 Jun 22 17:04 blk_-2152022802549252155_1036.meta
-rw-rw-r-- 1 grid grid    25 Jun 22 17:04 blk_3251284937612942908
-rw-rw-r-- 1 grid grid    11 Jun 22 17:04 blk_3251284937612942908_1037.meta
-rw-rw-r-- 1 grid grid    13 Jun 22 17:00 blk_7499053175856941845
-rw-rw-r-- 1 grid grid    11 Jun 22 17:00 blk_7499053175856941845_1031.meta
-rw-rw-r-- 1 grid grid    12 Jun 22 17:00 blk_-7677903991784965860
-rw-rw-r-- 1 grid grid    11 Jun 22 17:00 blk_-7677903991784965860_1030.meta
-rw-rw-r-- 1 grid grid     4 Jun 22 16:59 blk_8963744517366784034
-rw-rw-r-- 1 grid grid    11 Jun 22 16:59 blk_8963744517366784034_1029.meta
-rw-rw-r-- 1 grid grid   770 Jun 22 17:13 dncp_block_verification.log.curr
-rw-rw-r-- 1 grid grid   157 Jun 22 16:59 VERSION
```

读取数据流程

- 客户端要访问HDFS中的一个文件
- 首先从namenode获得组成这个文件的数据块位置列表
- 根据列表知道存储数据块的datanode
- 访问datanode获取数据
- Namenode并不参与数据实际传输

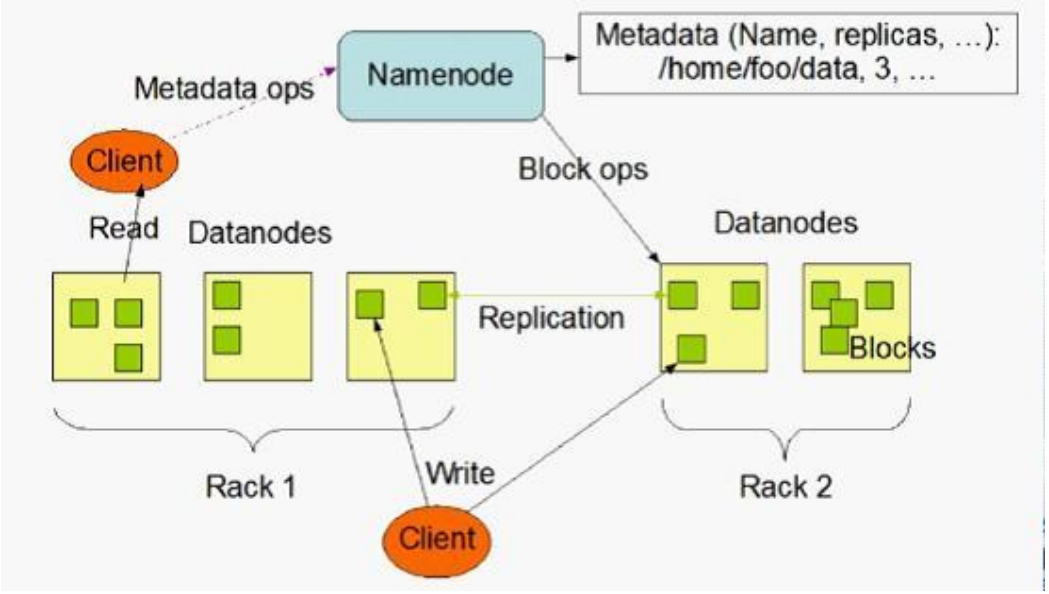

HDFS的可靠性

- 冗余副本策略
- 机架策略
- 心跳机制
- 安全模式
- 校验和
- 回收站
- 元数据保护
- 快照机制

冗余副本策略

- 可以在hdfs-site.xml中设置复制因子指定副本数量
- 所有数据块都有副本
- Datanode启动时，遍历本地文件系统，产生一份hdfs数据块和本地文件的对应关系列表（blockreport）汇报给namenode

机架策略

- 集群一般放在不同机架上，机架间带宽要比机架内带宽要小
- HDFS的“机架感知”
- 一般在本机架存放一个副本，在其它机架再存放别的副本，这样可以防止机架失效时丢失数据，也可以提高带宽利用率

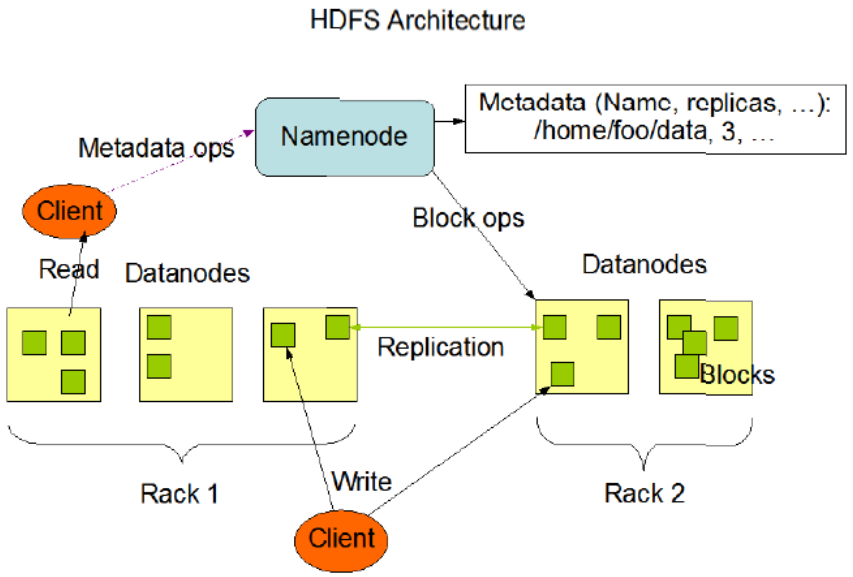

心跳机制

- Namenode周期性从datanode接收心跳信号和块报告
- Namenode根据块报告验证元数据
- 没有按时发送心跳的datanode会被标记为宕机，不会再给它任何I/O请求
- 如果datanode失效造成副本数量下降，并且低于预先设置的阈值，namenode会检测出这些数据块，并在合适的时机进行重新复制
- 引发重新复制的原因还包括数据副本本身损坏、磁盘错误，复制因子被增大等

安全模式

- Namenode启动时会先经过一个“安全模式”阶段
- 安全模式阶段不会产生数据写
- 在此阶段Namenode收集各个datanode的报告，当数据块达到最小副本数以上时，会被认为是“安全”的
- 在一定比例（可设置）的数据块被确定为“安全”后，再过若干时间，安全模式结束
- 当检测到副本数不足的数据块时，该块会被复制直到达到最小副本数

校验和

- 在文件创立时，每个数据块都产生校验和
- 校验和会作为单独一个隐藏文件保存在命名空间下
- 客户端获取数据时可以检查校验和是否相同，从而发现数据块是否损坏
- 如果正在读取的数据块损坏，则可以继续读取其它副本

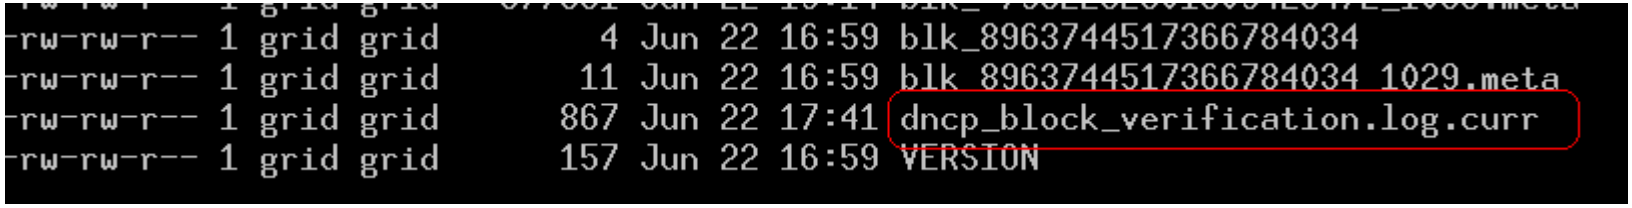

回收站

- 删除文件时，其实是放入回收站/trash
- 回收站里的文件可以快速恢复
- 可以设置一个时间阈值，当回收站里文件的存放时间超过这个阈值，就被彻底删除，并且释放占用的数据块

元数据保护

- 映像文件刚和事务日志是Namenode的核心数据。可以配置为拥有多个副本
- 副本会降低Namenode的处理速度，但增加安全性
- Namenode依然是单点，如果发生故障要手工切换

快照

- 支持存储某个时间点的映像，需要时可以使数据重返这个时间点的状态
- Hadoop目前还不支持快照，已经列入开发计划

HDFS文件操作

- 命令行方式
- API方式

列出HDFS下的文件

- 注意，hadoop没有当前目录的概念，也没有cd命令

```
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -ls
Found 3 items
-rw-r--r--   3 grid supergroup  182543327 2012-06-22 19:14 /user/grid/abc
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:00 /user/grid/in
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:04 /user/grid/out
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -ls ./in
Found 2 items
-rw-r--r--   3 grid supergroup         12 2012-06-22 17:00 /user/grid/in/test1.txt
-rw-r--r--   3 grid supergroup         13 2012-06-22 17:00 /user/grid/in/test2.txt
```

上传文件到HDFS

```
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -put ../abc abc
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -ls
Found 3 items
-rw-r--r--   3 grid supergroup  182543327 2012-06-22 19:14 /user/grid/abc
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:00 /user/grid/in
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:04 /user/grid/out
[grid@h3 hadoop-0.20.2]$ cd ../data
[grid@h3 data]$ ls
current  detach  in_use.lock  storage  tmp
[grid@h3 data]$ ls -lR
.:
total 16
drwxrwxr-x 2 grid grid 4096 Jun 22 19:14 current
drwxrwxr-x 2 grid grid 4096 Jun 22 16:59 detach
-rw-rw-r-- 1 grid grid    0 Jun 22 16:59 in_use.lock
-rw-rw-r-- 1 grid grid  157 Jun 22 16:59 storage
drwxrwxr-x 2 grid grid 4096 Jun 22 19:14 tmp

./current:
total 179952
-rw-rw-r-- 1 grid grid     8668 Jun 22 17:04 blk_1864853225322117619
-rw-rw-r-- 1 grid grid       75 Jun 22 17:04 blk_1864853225322117619_1037.meta
-rw-rw-r-- 1 grid grid    16666 Jun 22 17:04 blk_-2152022802549252155
-rw-rw-r-- 1 grid grid      139 Jun 22 17:04 blk_-2152022802549252155_1036.meta
-rw-rw-r-- 1 grid grid       25 Jun 22 17:04 blk_3251284937612942908
-rw-rw-r-- 1 grid grid       11 Jun 22 17:04 blk_3251284937612942908_1037.meta
-rw-rw-r-- 1 grid grid 67108864 Jun 22 19:14 blk_-4764241716670312355
-rw-rw-r-- 1 grid grid   524295 Jun 22 19:14 blk_-4764241716670312355_1038.meta
-rw-rw-r-- 1 grid grid 67108864 Jun 22 19:14 blk_-6923484112699544961
```

2012.9.3

将HDFS的文件复制到本地

```
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -get abc ./xyz
[grid@h3 hadoop-0.20.2]$ ls -l
total 183348
drwxr-xr-x  2 grid grid      4096 Sep 23  2002 bin
-rw-rw-r--  1 grid grid     74035 Sep 23  2002 build.xml
drwxr-xr-x  4 grid grid      4096 Feb 19  2010 c++
-rw-rw-r--  1 grid grid    348624 Sep 23  2002 CHANGES.txt
drwxr-xr-x  2 grid grid      4096 Jun 22 16:58 conf
drwxr-xr-x 13 grid grid      4096 Feb 19  2010 contrib
drwxr-xr-x  7 grid grid      4096 Sep 23  2002 docs
-rw-rw-r--  1 grid grid      6839 Sep 23  2002 hadoop-0.20.2-ant.jar
-rw-rw-r--  1 grid grid   2689741 Sep 23  2002 hadoop-0.20.2-core.jar
-rw-rw-r--  1 grid grid    142466 Sep 23  2002 hadoop-0.20.2-examples.jar
-rw-rw-r--  1 grid grid   1563859 Sep 23  2002 hadoop-0.20.2-test.jar
-rw-rw-r--  1 grid grid     69940 Sep 23  2002 hadoop-0.20.2-tools.jar
drwxr-xr-x  2 grid grid      4096 Sep 23  2002 ivy
-rw-rw-r--  1 grid grid      8852 Sep 23  2002 ivy.xml
drwxr-xr-x  5 grid grid      4096 Sep 23  2002 lib
drwxr-xr-x  2 grid grid      4096 Sep 23  2002 librecordio
-rw-rw-r--  1 grid grid     13366 Sep 23  2002 LICENSE.txt
drwxrwxr-x  3 grid grid      4096 Jun 22 17:04 logs
-rw-rw-r--  1 grid grid       101 Sep 23  2002 NOTICE.txt
-rw-rw-r--  1 grid grid      1366 Sep 23  2002 README.txt
drwxr-xr-x 15 grid grid      4096 Sep 23  2002 src
drwxr-xr-x  8 grid grid      4096 Feb 19  2010 webapps
-rw-rw-r--  1 grid grid 182543327 Jun 22 20:40 xyz
```

删除HDFS下的文档

```
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -rmr abc
Deleted hdfs://h1:9000/user/grid/abc
[grid@h3 hadoop-0.20.2]$ cd ../data
[grid@h3 data]$ ls -lR
.:
total 16
drwxrwxr-x 2 grid grid 4096 Jun 22 20:41 current
drwxrwxr-x 2 grid grid 4096 Jun 22 16:59 detach
-rw-rw-r-- 1 grid grid    0 Jun 22 16:59 in_use.lock
-rw-rw-r-- 1 grid grid  157 Jun 22 16:59 storage
drwxrwxr-x 2 grid grid 4096 Jun 22 19:14 tmp

./current:
total 80
-rw-rw-r-- 1 grid grid  8668 Jun 22 17:04 blk_1864853225322117619
-rw-rw-r-- 1 grid grid    75 Jun 22 17:04 blk_1864853225322117619_1037.meta
-rw-rw-r-- 1 grid grid 16666 Jun 22 17:04 blk_-2152022802549252155
-rw-rw-r-- 1 grid grid   139 Jun 22 17:04 blk_-2152022802549252155_1036.meta
-rw-rw-r-- 1 grid grid    25 Jun 22 17:04 blk_3251284937612942908
-rw-rw-r-- 1 grid grid    11 Jun 22 17:04 blk_3251284937612942908_1037.meta
-rw-rw-r-- 1 grid grid    13 Jun 22 17:00 blk_7499053175856941845
-rw-rw-r-- 1 grid grid    11 Jun 22 17:00 blk_7499053175856941845_1031.meta
-rw-rw-r-- 1 grid grid    12 Jun 22 17:00 blk_-7677903991784965860
-rw-rw-r-- 1 grid grid    11 Jun 22 17:00 blk_-7677903991784965860_1030.meta
-rw-rw-r-- 1 grid grid     4 Jun 22 16:59 blk_8963744517366784034
-rw-rw-r-- 1 grid grid    11 Jun 22 16:59 blk_8963744517366784034_1029.meta
-rw-rw-r-- 1 grid grid  1158 Jun 22 20:22 dncp_block_verification.log.curr
-rw-rw-r-- 1 grid grid   157 Jun 22 16:59 VERSION
```

查看HDFS下某个文件的内容

```
[grid@h3 hadoop-0.20.2]$
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -ls
Found 2 items
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:00 /user/grid/in
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:04 /user/grid/out
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -ls ./out
Found 2 items
drwxr-xr-x   - grid supergroup          0 2012-06-22 17:04 /user/grid/out/_logs
-rw-r--r--   3 grid supergroup         25 2012-06-22 17:04 /user/grid/out/part-r-00000
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfs -cat ./out/part-r-00000
hadoop  1
hello   2
world   1
[grid@h3 hadoop-0.20.2]$
```

查看HDFS基本统计信息

```
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfsadmin -report
Configured Capacity: 101801435136 (94.81 GB)
Present Capacity: 91998890308 (85.68 GB)
DFS Remaining: 91998568448 (85.68 GB)
DFS Used: 321860 (314.32 KB)
DFS Used%: 0%
Under replicated blocks: 6
Blocks with corrupt replicas: 0
Missing blocks: 0

-------------------------------------------------
Datanodes available: 1 (1 total, 0 dead)

Name: 192.168.1.104:50010
Decommission Status : Normal
Configured Capacity: 101801435136 (94.81 GB)
DFS Used: 321860 (314.32 KB)
Non DFS Used: 9802544828 (9.13 GB)
DFS Remaining: 91998568448(85.68 GB)
DFS Used%: 0%
DFS Remaining%: 90.37%
Last contact: Fri Jun 22 20:44:36 EDT 2012
```

进入和退出安全模式

```
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfsadmin -safemode enter
Safe mode is ON
[grid@h3 hadoop-0.20.2]$
[grid@h3 hadoop-0.20.2]$ bin/hadoop dfsadmin -safemode leave
Safe mode is OFF
[grid@h3 hadoop-0.20.2]$
```

怎样添加节点？

- 在新节点安装好hadoop
- 把namenode的有关配置文件复制到该节点点
- 修改masters和slaves文件，增加该节点点
- 设置ssh免密码进出该节点点
- 单独启动该节点点上的datanode和tasktracker（hadoop-daemon.sh start datanode/tasktracker）
- 运行start-balancer.sh进行数据负载均衡

启动某些特定后台进程而非所有后台进程

- Start-all.sh的内容

```
bin=`dirname "$0"`
bin=`cd "$bin"; pwd`

. "$bin"/hadoop-config.sh

# start dfs daemons
"$bin"/start-dfs.sh --config $HADOOP_CONF_DIR

# start mapred daemons
"$bin"/start-mapred.sh --config $HADOOP_CONF_DIR
[grid@h3 bin]$
```

负载均衡

- 作用：当节点出现故障，或新增加节点时，数据块分布可能不均匀，负载均衡可以重新平衡各个datanode上数据块的分布

```
[grid@h3 bin]$ ./start-balancer.sh
starting balancer, logging to /home/grid/hadoop-0.20.2/bin/../logs/hadoop-grid-balancer-h3.out
[grid@h3 bin]$
```

HDFS API

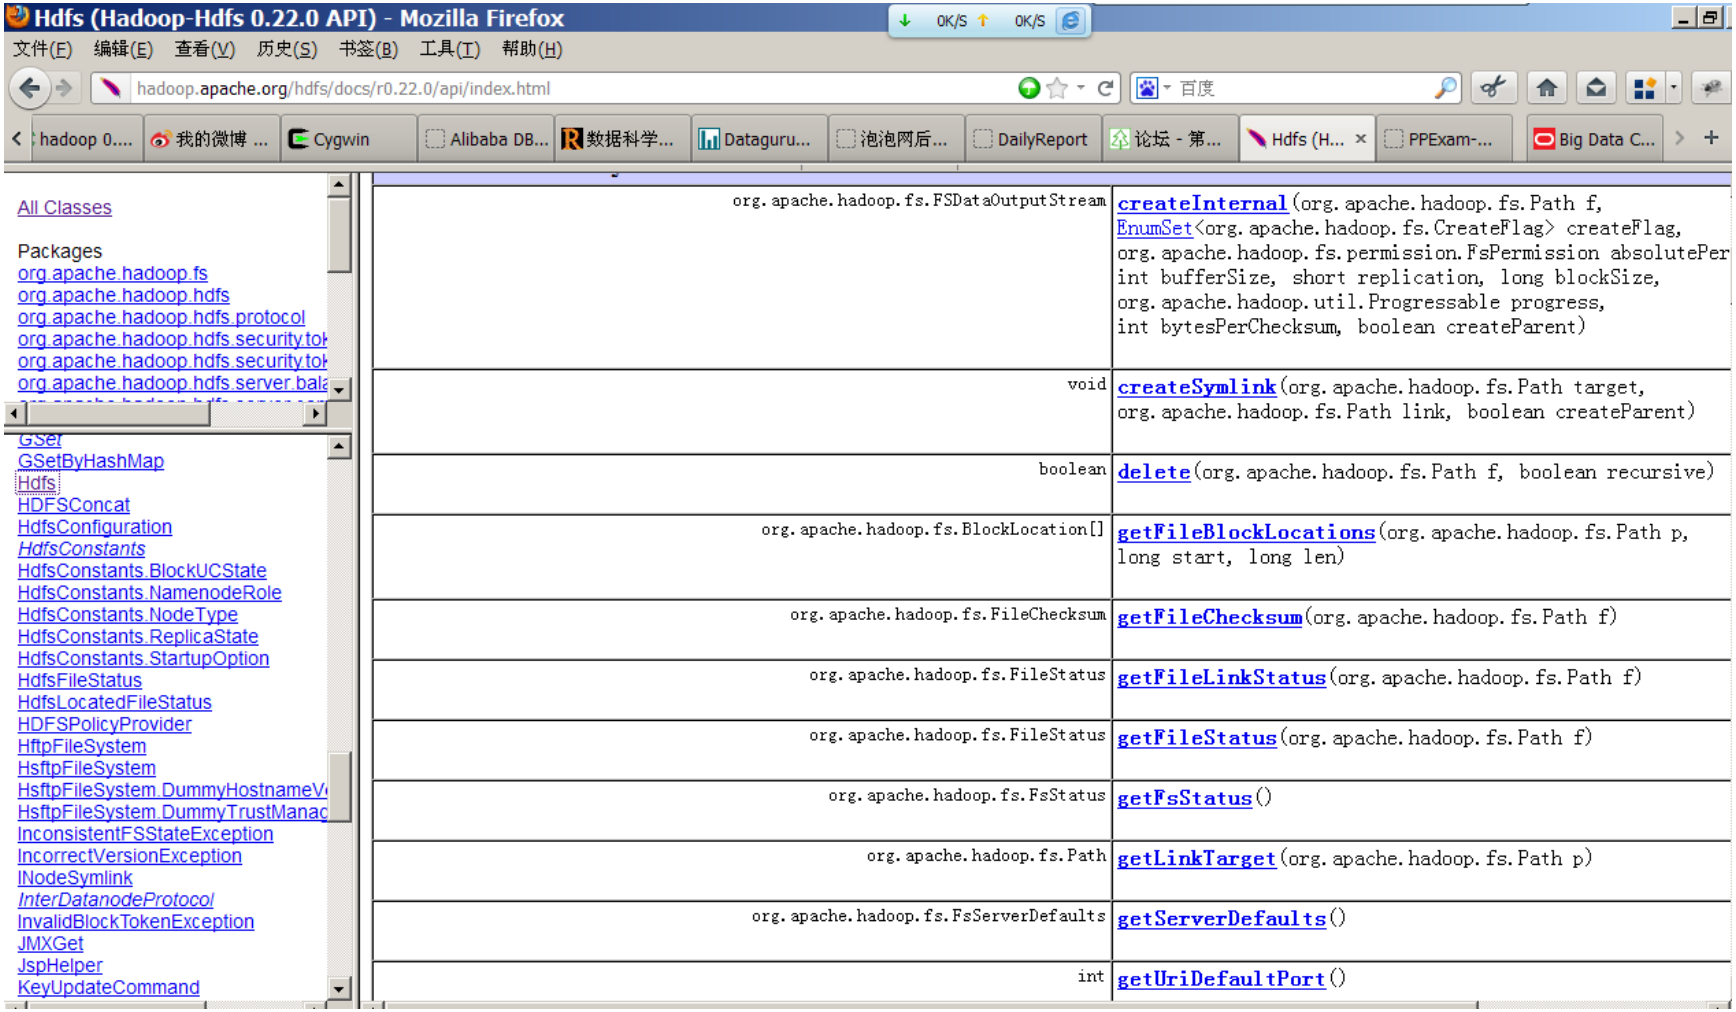

2012.9.3

学习Java的参考书

- 《Java就业培训教程》，张孝祥，清华大学出版社
- 《Java程序设计教程》，雍俊海，清华大学出版社

上传本地文件到HDFS

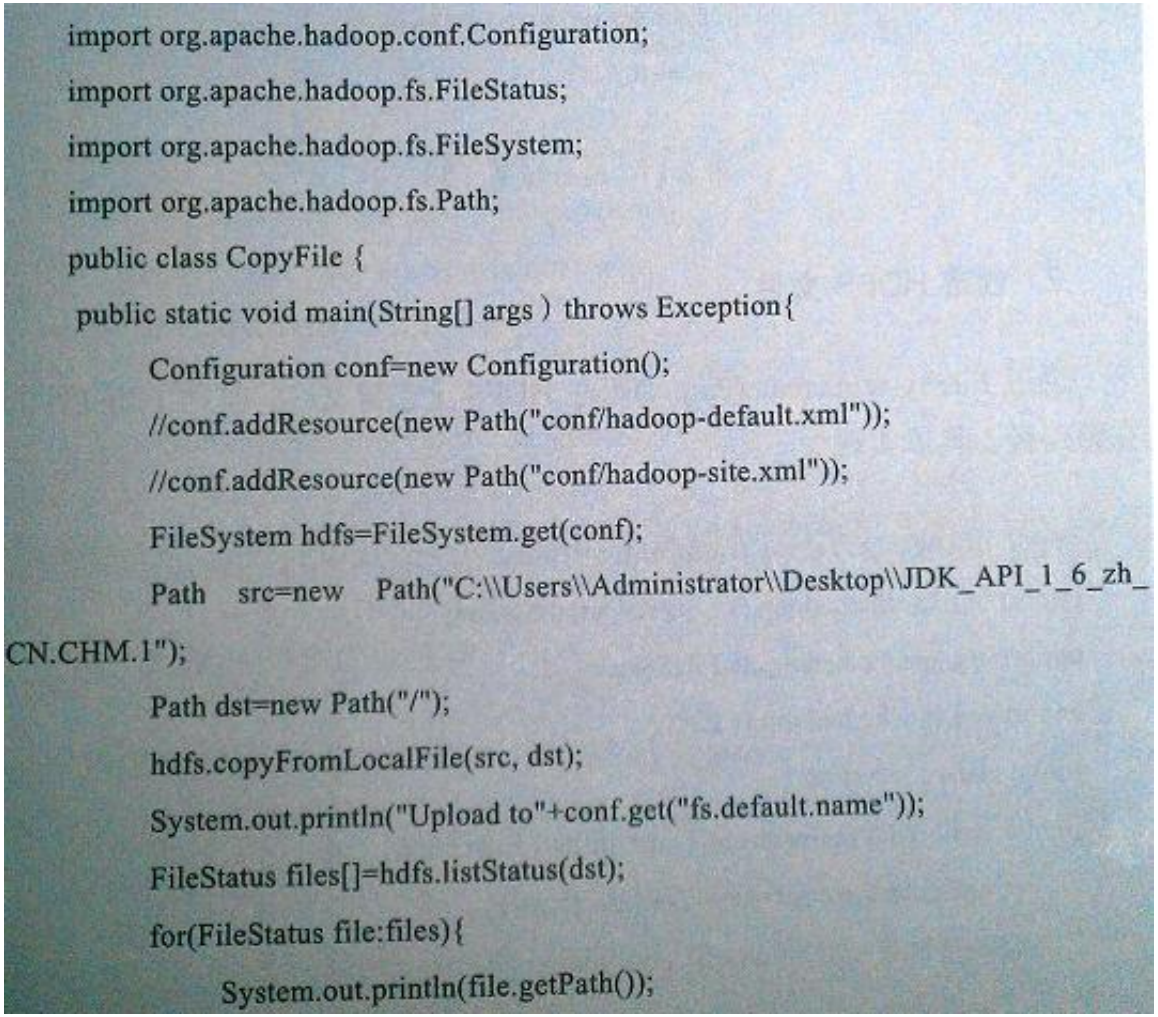

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.fs.FileStatus;
import org.apache.hadoop.fs.FileSystem;
import org.apache.hadoop.fs.Path;
public class CopyFile {
 public static void main(String[] args ) throws Exception{
        Configuration conf=new Configuration();
        //conf.addResource(new Path("conf/hadoop-default.xml"));
        //conf.addResource(new Path("conf/hadoop-site.xml"));
        FileSystem hdfs=FileSystem.get(conf);
        Path   src=new   Path("C:\\Users\\Administrator\\Desktop\\JDK_API_1_6_zh_
CN.CHM.1");
        Path dst=new Path("/");
        hdfs.copyFromLocalFile(src, dst);
        System.out.println("Upload to"+conf.get("fs.default.name"));
        FileStatus files[]=hdfs.listStatus(dst);
        for(FileStatus file:files){
              System.out.println(file.getPath());
```

2012.9.3

创建HDFS文件

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.fs.FSDataOutputStream;
import org.apache.hadoop.fs.FileSystem;
import org.apache.hadoop.fs.Path;
public class CreateFile {
 public static void main(String[] args)throws Exception{
        Configuration conf=new Configuration();
        byte[] buff ="hello word!".getBytes();
        FileSystem hdfs = FileSystem.get(conf);
        Path dfs = new Path("/test");
        FSDataOutputStream outputStream = hdfs.create(dfs);
        outputStream.write(buff, 0, buff.length);
        }
```

重命名HDFS文件

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.fs.FileSystem;
import org.apache.hadoop.fs.Path;
public class Rename {
 public static void main(String[] args)throws Exception{
        Configuration conf=new Configuration();
        FileSystem hdfs=FileSystem.get(conf);
        Path frpath=new Path("/test");
        Path topath=new Path("/test1");
        boolean isRename=hdfs.rename(frpath, topath);
 }
```

查看HDFS文件的最后修改时间

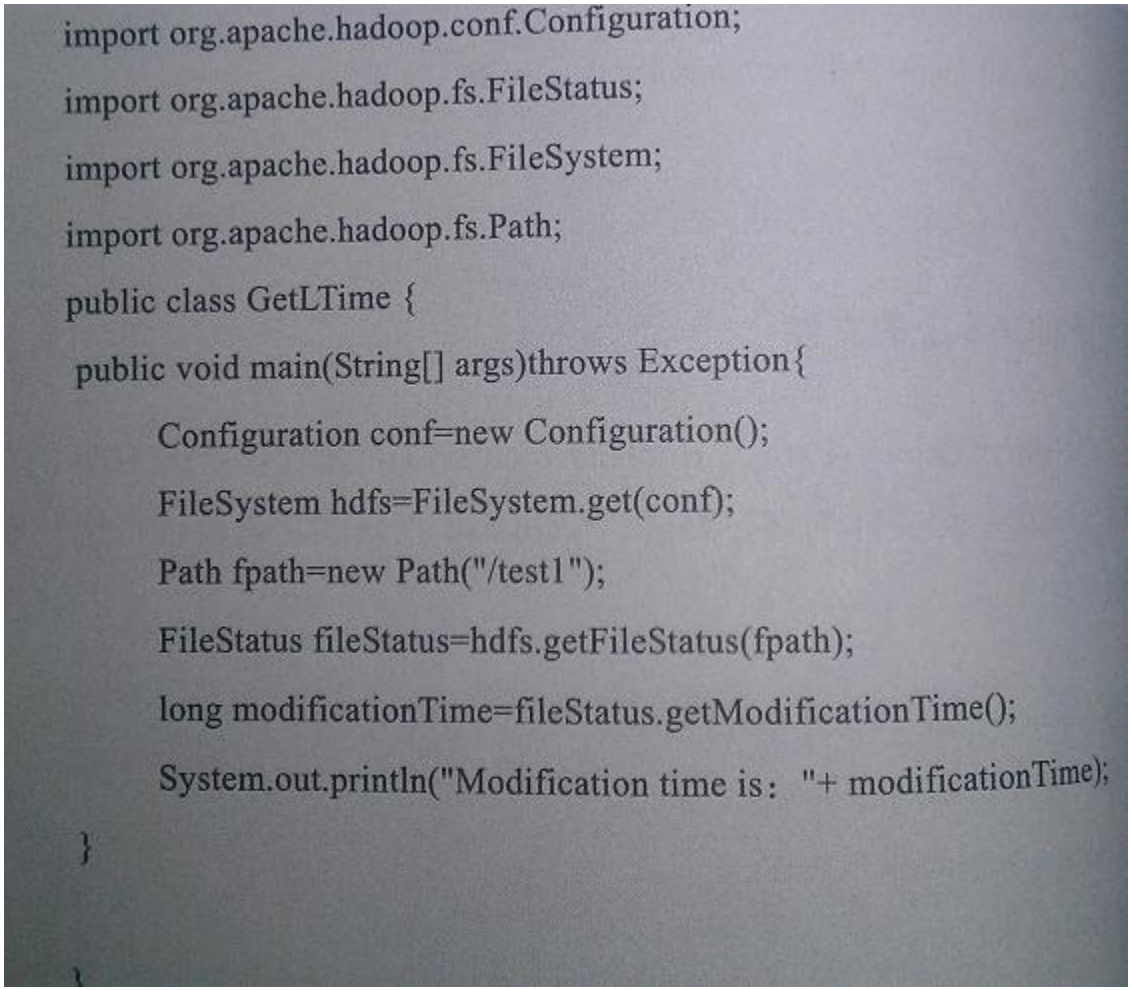

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.fs.FileStatus;
import org.apache.hadoop.fs.FileSystem;
import org.apache.hadoop.fs.Path;
public class GetLTime {
 public void main(String[] args)throws Exception{
        Configuration conf=new Configuration();
        FileSystem hdfs=FileSystem.get(conf);
        Path fpath=new Path("/test1");
        FileStatus fileStatus=hdfs.getFileStatus(fpath);
        long modificationTime=fileStatus.getModificationTime();
        System.out.println("Modification time is：  "+ modificationTime);
 }
}
```

Thanks

FAQ时间