



链滴

# CentOS 7.9 系统搭建 Hadoop 集群

作者: [lingyundu](#)

原文链接: <https://ld246.com/article/1647588066643>

来源网站: [链滴](#)

许可协议: [署名-相同方式共享 4.0 国际 \(CC BY-SA 4.0\)](#)

# 安装环境

虚拟软件: VMware® Workstation 16 Pro

虚拟机操作系统: CentOS 7.9-Minimal

虚拟机 IP: 192.168.153.11、192.168.153.12、192.168.153.13

# 前期规划

Hadoop 集群包含两个集群: HDFS 集群、YARN 集群, 两个集群在逻辑上分离, 但通常会共用主机。

两个集群都是标准的主从架构集群。

HDFS 集群包含的角色 (守护进程):

- 主角色: NameNode
- 从角色: DataNode
- 主角色辅助角色: SecondaryNameNode

YARN 集群包含的角色 (守护进程):

- 主角色: ResourceManager
- 从角色: NodeManager

集群规划:

服务器 程)	IP 地址	运行角色 (守护
node1.hadoop.com ameNode DataNode ResourceManager NodeManager	192.168.153.11	
node2.hadoop.com econdaryNameNode DataNode NodeManager	192.168.153.12	
node3.hadoop.com ataNode NodeManager	192.168.153.13	

# 环境配置

每台虚拟机都要配置, 使用 root 用户。

1、关闭防火墙

```
systemctl stop firewalld  
systemctl disable firewalld
```

2、同步时间

```
yum -y install ntpdate  
ntpdate ntp5.aliyun.com
```

### 3、配置主机名

```
vi /etc/hostname
```

按照规划，将三台虚拟机的主机名分别设置为：`node1.hadoop.com`、`node2.hadoop.com` 和 `node3.hadoop.com`。

### 4、配置 hosts 文件

```
vi /etc/hosts
```

添加下面的内容：

```
192.168.153.11 node1 node1.hadoop.com
192.168.153.12 node2 node1.hadoop.com
192.168.153.13 node3 node1.hadoop.com
```

### 5、安装 JDK

```
yum -y install java-1.8.0-openjdk java-1.8.0-openjdk-devel
```

配置 JAVA\_HOME

```
cat <<EOF | tee /etc/profile.d/hadoop_java.sh
export JAVA_HOME=\$(dirname \$(dirname \$(readlink \$(readlink \$(which javac))))))
export PATH=\$PATH:\$JAVA_HOME/bin
EOF
source /etc/profile.d/hadoop_java.sh
```

确认：

```
echo $JAVA_HOME
```

### 6、创建 hadoop 用户，并设置密码

```
adduser hadoop
usermod -aG wheel hadoop
passwd hadoop
```

创建 HDFS 本地存放数据的目录：

```
mkdir /home/hadoop/data
chown hadoop: /home/hadoop/data
```

### 7、配置环境变量

```
echo 'export HADOOP_HOME=/home/hadoop/hadoop-3.3.2' >> /etc/profile
echo 'export PATH=\$PATH:\$HADOOP_HOME/bin:\$HADOOP_HOME/sbin' >> /etc/profile
source /etc/profile
```

### 8、配置 SSH

```
yum install openssh
```

切换到 hadoop 用户，执行下面的命令。

```
ssh-keygen
ssh-copy-id node1
ssh-copy-id node2
ssh-copy-id node3
```

每台虚拟机都要执行，执行过程如下：

```
[hadoop@node1 ~]$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/home/hadoop/.ssh/id_rsa):
Created directory '/home/hadoop/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/hadoop/.ssh/id_rsa.
Your public key has been saved in /home/hadoop/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:gFs4NEpc6MIVv7/r5f2rUFdOi7ht11GceM3fd/Uq/nU hadoop@node1.hadoop.com
The key's randomart image is:
+----[RSA 2048]-----+
|..+=          |
|.O+.+        .oo|
|..O +.O      .=*|
|... +..     .* B|
|. .. S o o +*|
|   . . + . =|
|   . o ..o..E|
|   + o.....|
|. +.. o++o |
+----[SHA256]-----+
[hadoop@node1 ~]$ ssh-copy-id node1
/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/hadoop/.ssh/id_rsa.pub"
The authenticity of host 'node1 (192.168.153.11)' can't be established.
ECDSA key fingerprint is SHA256:BxdxJ5ONWI6xkPrFWxy9MIFs/B3IpEgjhFxiwI6KOLU.
ECDSA key fingerprint is MD5:78:ea:2d:36:7e:eb:83:47:8f:61:c6:70:b6:0f:20:d6.
Are you sure you want to continue connecting (yes/no)? yes
/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed
/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys
hadoop@node1's password:
```

Number of key(s) added: 1

Now try logging into the machine, with: "ssh 'node1'"  
and check to make sure that only the key(s) you wanted were added.

```
[hadoop@node1 ~]$ ssh-copy-id node2
/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/hadoop/.ssh/id_rsa.pub"
The authenticity of host 'node2 (192.168.153.12)' can't be established.
ECDSA key fingerprint is SHA256:BxdxJ5ONWI6xkPrFWxy9MIFs/B3IpEgjhFxiwI6KOLU.
ECDSA key fingerprint is MD5:78:ea:2d:36:7e:eb:83:47:8f:61:c6:70:b6:0f:20:d6.
Are you sure you want to continue connecting (yes/no)? yes
/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed
```

```
/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys
hadoop@node2's password:
```

```
Number of key(s) added: 1
```

Now try logging into the machine, with: "ssh 'node2'"  
and check to make sure that only the key(s) you wanted were added.

```
[hadoop@node1 ~]$ ssh-copy-id node3
/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/hadoop/.ssh/id_rsa.pub"
The authenticity of host 'node3 (192.168.153.13)' can't be established.
ECDSA key fingerprint is SHA256:BxdxJ5ONWI6xkPrFWxy9MIFs/B3IpEgjhFxiwI6KOLU.
ECDSA key fingerprint is MD5:78:ea:2d:36:7e:eb:83:47:8f:61:c6:70:b6:0f:20:d6.
Are you sure you want to continue connecting (yes/no)? yes
/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed
/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys
hadoop@node3's password:
```

```
Number of key(s) added: 1
```

Now try logging into the machine, with: "ssh 'node3'"  
and check to make sure that only the key(s) you wanted were added.

```
[hadoop@node1 ~]$
```

## 下载安装

先在 node1 虚拟机进行安装配置，然后把安装好的目录复制到另外两台虚拟机。（使用 hadoop 用户）

### 1、下载并解压

使用 hadoop 用户连接 node1 虚拟机，用下面的命令下载安装包到 `/home/hadoop` 目录。

```
cd /home/hadoop
curl -Ok https://dlcdn.apache.org/hadoop/common/hadoop-3.3.2/hadoop-3.3.2.tar.gz
```

解压：

```
tar xzf hadoop-3.3.2.tar.gz
```

接下来通过配置文件对 Hadoop 进行配置。

Hadoop 的配置文件分为三类：

- 默认配置文件 -- 包括 `core-default.xml`、`hdfs-default.xml`、`yarn-default.xml` 和 `mapred-default.xml`，这些文件是只读的，存放的是参数的默认值。
- 自定义配置文件 -- 包括 `etc/hadoop/core-site.xml`、`etc/hadoop/hdfs-site.xml`、`etc/hadoop/yarn-site.xml` 和 `etc/hadoop/mapred-site.xml`，用来存放自定义配置信息，将会覆盖默认配置。
- 环境配置文件 -- 包括 `etc/hadoop/hadoop-env.sh`、`etc/hadoop/mapred-env.sh` 和 `etc/had`

`op/yarn-env.sh`, 这些文件用来配置各守护进程的 Java 运行环境。

## 2、配置 `hadoop-env.sh` 文件

```
cd hadoop-3.3.2
vi etc/hadoop/hadoop-env.sh
```

添加下面这些内容:

```
export JAVA_HOME=$JAVA_HOME
export HDFS_NAMENODE_USER=hadoop
export HDFS_DATANODE_USER=hadoop
export HDFS_SECONDARYNAMENODE_USER=hadoop
export YARN_RESOURCEMANAGER_USER=hadoop
export YARN_NODEMANAGER_USER=hadoop
```

至少要配置 `JAVA_HOME` 环境变量, 另外可以通过下面这些变量, 为不同的守护进程单独进行配置:

守护进程	环境变量
NameNode	HDFS_NAMENODE_OPTS
DataNode	HDFS_DATANODE_OPTS
Secondary NameNode MENODE_OPTS	HDFS_SECONDARYN
ResourceManager _OPTS	YARN_RESOURCEMANAGE
NodeManager	YARN_NODEMANAGER_OPTS
WebAppProxy	YARN_PROXYSERVER_OPTS
Map Reduce Job History Server ORYSERVER_OPTS	MAPRED_HIS

例如, 给 Namenode 配置使用 `parallelGC` 和 4GB 堆内存:

```
export HDFS_NAMENODE_OPTS="-XX:+UseParallelGC -Xmx4g"
```

## 3、配置 `core-site.xml` 文件

该文件将会覆盖 `core-default.xml` 中的配置。

```
vi etc/hadoop/core-site.xml
```

添加下面的内容:

```
<!-- 设置默认使用的文件系统 Hadoop 支持 file、HDFS、GFS、Ali Cloud、Amazon Cloud 等文  
系统 -->  
<property>  
  <name>fs.defaultFS</name>  
  <value>hdfs://node1:8020</value>  
</property>  
  
<!-- 设置 Hadoop 本地保存数据的路径 -->  
<property>
```

```
<name>hadoop.tmp.dir</name>
<value>/home/hadoop/data</value>
</property>

<!-- 设置 Hadoop web UI 用户身份 -->
<property>
  <name>hadoop.http.staticuser.user</name>
  <value>hadoop</value>
</property>

<!-- 整合 Hive 用户代理设置 -->
<property>
  <name>hadoop.proxyuser.root.hosts</name>
  <value>*</value>
</property>

<!-- 文件垃圾桶保存时间 -->
<property>
  <name>fs.trash.interval</name>
  <value>1440</value>
</property>
```

#### 4、配置 hdfs-site.xml 文件

该文件将会覆盖 [hdfs-default.xml](#) 中的配置。

```
vi etc/hadoop/hdfs-site.xml
```

添加下面的内容：

```
<!-- 设置 SNN 进程运行机器位置信息 -->
<property>
  <name>dfs.namenode.secondary.http-address</name>
  <value>node2:9868</value>
</property>
```

#### 5、配置 mapred-site.xml 文件

该文件将会覆盖 [mapred-default.xml](#) 中的配置。

```
vi etc/hadoop/mapred-site.xml
```

添加下面的内容：

```
<!-- 设置 MR 程序默认运行模式： yarn 集群模式， local 本地模式 -->
<property>
  <name>mapreduce.framework.name</name>
  <value>yarn</value>
</property>

<!-- MR 程序历史服务地址 -->
<property>
  <name>mapreduce.jobhistory.address</name>
  <value>node1:10020</value>
```

```
</property>

<!-- MR 程序历史服务器 web 端地址 -->
<property>
  <name>mapreduce.jobhistory.webapp.address</name>
  <value>node1:19888</value>
</property>

<property>
  <name>yarn.app.mapreduce.am.env</name>
  <value>HADOOP_MAPRED_HOME=${HADOOP_HOME}</value>
</property>

<property>
  <name>mapreduce.map.env</name>
  <value>HADOOP_MAPRED_HOME=${HADOOP_HOME}</value>
</property>

<property>
  <name>mapreduce.reduce.env</name>
  <value>HADOOP_MAPRED_HOME=${HADOOP_HOME}</value>
</property>
```

## 6、配置 yarn-site.xml 文件

该文件将会覆盖 [yarn-default.xml](#) 中的配置。

```
vi etc/hadoop/yarn-site.xml
```

添加下面的内容：

```
<!-- 设置 YARN 集群主角色运行机器位置 -->
<property>
  <name>yarn.resourcemanager.hostname</name>
  <value>node1</value>
</property>

<property>
  <name>yarn.nodemanager.aux-services</name>
  <value>mapreduce_shuffle</value>
</property>

<!-- 是否对容器实施物理内存限制 -->
<property>
  <name>yarn.nodemanager.pmem-check-enabled</name>
  <value>>false</value>
</property>

<!-- 是否对容器实施虚拟内存限制 -->
<property>
  <name>yarn.nodemanager.vmem-check-enabled</name>
  <value>>false</value>
</property>

<!-- 开启日志聚集-->
```



```
<property>
  <name>yarn.log-aggregation-enable</name>
  <value>>true</value>
</property>

<!-- 设置 yarn 历史服务器地址 -->
<property>
  <name>yarn.log.server.url</name>
  <value>http://node1:19888/jobhistory/logs</value>
</property>
```

## 7、配置 workers 文件

```
vi etc/hadoop/workers
```

删除原来内容，并添加下面的内容：

```
node1.hadoop.com
node2.hadoop.com
node3.hadoop.com
```

## 8、将配置好的安装包复制到 node2 和 node3 机器。

```
scp -r /home/hadoop/hadoop-3.3.2 hadoop@node2:/home/hadoop/
scp -r /home/hadoop/hadoop-3.3.2 hadoop@node3:/home/hadoop/
```

# 启动集群

Hadoop 提供了两种启动方式：

- 使用命令逐个启动进程 -- 每台机器都要手动执行命令，可精准控制每个进程的启动。
- 使用脚本一键启动 -- 前提是要配置好机器之间的 SSH 免密登录和 `etc/hadoop/workers` 文件。

### 逐个启动进程的命令：

```
# HDFS 集群
$HADOOP_HOME/bin/hdfs --daemon start namenode | datanode | secondarynamenode

# YARN 集群
$HADOOP_HOME/bin/yarn --daemon start resourcemanager | nodemanager | proxyserver
```

### 启动集群的脚本：

- HDFS 集群 -- `$HADOOP_HOME/sbin/start-dfs.sh`，一键启动 HDFS 集群的所有进程。
- YARN 集群 -- `$HADOOP_HOME/sbin/start-yarn.sh`，一键启动 YARN 集群的所有进程
- Hadoop 集群 -- `$HADOOP_HOME/sbin/start-all.sh`，一键启动 HDFS 集群和 YARN 集群的有进程。

## 1、格式化文件系统

启动集群之前，需要对 HDFS 进行格式化（仅在 node1 机器执行）。

```
[hadoop@node1 ~]$ hdfs namenode -format
```

```
WARNING: /home/hadoop/hadoop-3.3.2/logs does not exist. Creating.
2022-03-17 23:22:55,296 INFO namenode.NameNode: STARTUP_MSG:
/*****
STARTUP_MSG: Starting NameNode
STARTUP_MSG: host = node1/192.168.153.11
STARTUP_MSG: args = [-format]
STARTUP_MSG: version = 3.3.2
STARTUP_MSG: classpath = /home/hadoop/hadoop-3.3.2/etc/hadoop:/home/hadoop/hado
p-3.3.2/share/hadoop/common/lib/accessors-smart-2.4.7.jar:/home/hadoop/hadoop-3.3.2/sh
re/hadoop/common/lib/animal-sniffer-annotations-1.17.jar:/home/hadoop/hadoop-3.3.2/sha
e/hadoop/common/lib/asm-5.0.4.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/li
/audience-annotations-0.5.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/avr
-1.7.7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/checker-qual-2.5.2.jar:/h
me/hadoop/hadoop-3.3.2/share/hadoop/common/lib/commons-beanutils-1.9.4.jar:/home/h
doop/hadoop-3.3.2/share/hadoop/common/lib/commons-cli-1.2.jar:/home/hadoop/hadoop
3.3.2/share/hadoop/common/lib/commons-codec-1.11.jar:/home/hadoop/hadoop-3.3.2/shar
/hadoop/common/lib/commons-collections-3.2.2.jar:/home/hadoop/hadoop-3.3.2/share/ha
oop/common/lib/commons-compress-1.21.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/c
mmon/lib/commons-configuration2-2.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/c
mmon/lib/commons-daemon-1.0.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/commo
/lib/commons-io-2.8.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/common
-lang3-3.12.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/commons-logging
1.1.3.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/commons-math3-3.1.1.jar:
home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/commons-net-3.6.jar:/home/hadoo
/hadoop-3.3.2/share/hadoop/common/lib/commons-text-1.4.jar:/home/hadoop/hadoop-3.3.
/share/hadoop/common/lib/curator-client-4.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hado
p/common/lib/curator-framework-4.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/co
mon/lib/curator-recipes-4.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/d
sjava-2.1.7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/failureaccess-1.0.jar:/
ome/hadoop/hadoop-3.3.2/share/hadoop/common/lib/gson-2.8.9.jar:/home/hadoop/hadoo
-3.3.2/share/hadoop/common/lib/guava-27.0-jre.jar:/home/hadoop/hadoop-3.3.2/share/had
op/common/lib/hadoop-annotations-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/c
mmon/lib/hadoop-auth-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/h
adoop-shaded-guava-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/hado
p-shaded-protobuf_3_7-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/ht
pclient-4.5.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/httpcore-4.4.13.jar
/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/j2objc-annotations-1.1.jar:/home/h
doop/hadoop-3.3.2/share/hadoop/common/lib/jackson-annotations-2.13.0.jar:/home/hadoo
/hadoop-3.3.2/share/hadoop/common/lib/jackson-core-2.13.0.jar:/home/hadoop/hadoop-3.3
2/share/hadoop/common/lib/jackson-core-asl-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/
adoop/common/lib/jackson-databind-2.13.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/
ommon/lib/jackson-jaxrs-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jack
son-mapper-asl-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jacks
n-xc-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jakarta.activation-ap
-1.2.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/javax.servlet-api-3.1.0.jar:
home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jaxb-api-2.2.11.jar:/home/hadoop/h
doop-3.3.2/share/hadoop/common/lib/jaxb-impl-2.2.3-1.jar:/home/hadoop/hadoop-3.3.2/sh
re/hadoop/common/lib/jcip-annotations-1.0-1.jar:/home/hadoop/hadoop-3.3.2/share/hado
p/common/lib/jersey-core-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jer
sey-json-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jersey-server-1.
9.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jersey-servlet-1.19.jar:/home/
adoop/hadoop-3.3.2/share/hadoop/common/lib/jettison-1.1.jar:/home/hadoop/hadoop-3.3.
/share/hadoop/common/lib/jetty-http-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/sha
e/hadoop/common/lib/jetty-io-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/had
```

op/common/lib/jetty-security-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop  
p/common/lib/jetty-server-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop  
common/lib/jetty-servlet-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/c  
mmon/lib/jetty-util-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/comm  
n/lib/jetty-util-ajax-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/commo  
/lib/jetty-webapp-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common  
lib/jetty-xml-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/j  
ch-0.1.55.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/json-smart-2.4.7.jar:/h  
me/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jsp-api-2.1.jar:/home/hadoop/hadoop  
3.3.2/share/hadoop/common/lib/jsr305-3.0.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop  
common/lib/jsr311-api-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/jul-  
o-slf4j-1.7.30.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerb-admin-1.0.1.j  
r:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerb-client-1.0.1.jar:/home/hado  
p/hadoop-3.3.2/share/hadoop/common/lib/kerb-common-1.0.1.jar:/home/hadoop/hadoop-3  
3.2/share/hadoop/common/lib/kerb-core-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hado  
p/common/lib/kerb-crypto-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib  
kerb-identity-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerb-server-1  
0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerb-simplekdc-1.0.1.jar:/h  
me/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerb-util-1.0.1.jar:/home/hadoop/had  
op-3.3.2/share/hadoop/common/lib/kerby-asn1-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share  
hadoop/common/lib/kerby-config-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/co  
mon/lib/kerby-pkix-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerby-  
til-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/kerby-xdr-1.0.1.jar:/ho  
e/hadoop/hadoop-3.3.2/share/hadoop/common/lib/listenablefuture-9999.0-empty-to-avoid-  
onflict-with-guava.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/log4j-1.2.17.j  
r:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/metrics-core-3.2.4.jar:/home/had  
op/hadoop-3.3.2/share/hadoop/common/lib/netty-3.10.6.Final.jar:/home/hadoop/hadoop-3.3  
2/share/hadoop/common/lib/nimbus-jose-jwt-9.8.1.jar:/home/hadoop/hadoop-3.3.2/share/h  
adoop/common/lib/paranamer-2.3.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/l  
b/protobuf-java-2.5.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/re2j-1.1.jar  
/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/slf4j-api-1.7.30.jar:/home/hadoop/  
adoop-3.3.2/share/hadoop/common/lib/slf4j-log4j12-1.7.30.jar:/home/hadoop/hadoop-3.3.2  
share/hadoop/common/lib/snappy-java-1.1.8.2.jar:/home/hadoop/hadoop-3.3.2/share/hado  
p/common/lib/stax2-api-4.2.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/t  
ken-provider-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/woodstox-co  
e-5.3.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/lib/zookeeper-3.5.6.jar:/ho  
e/hadoop/hadoop-3.3.2/share/hadoop/common/lib/zookeeper-jute-3.5.6.jar:/home/hadoop  
hadoop-3.3.2/share/hadoop/common/hadoop-common-3.3.2-tests.jar:/home/hadoop/hado  
p-3.3.2/share/hadoop/common/hadoop-common-3.3.2.jar:/home/hadoop/hadoop-3.3.2/sha  
e/hadoop/common/hadoop-kms-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/com  
on/hadoop-nfs-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/common/hadoop-regist  
y-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs:/home/hadoop/hadoop-3.3.2/sh  
re/hadoop/hdfs/lib/accessors-smart-2.4.7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hd  
s/lib/animal-sniffer-annotations-1.17.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib  
asm-5.0.4.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/audience-annotations-0.5.  
.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/avro-1.7.7.jar:/home/hadoop/hado  
p-3.3.2/share/hadoop/hdfs/lib/checker-qual-2.5.2.jar:/home/hadoop/hadoop-3.3.2/share/ha  
oop/hdfs/lib/commons-beanutils-1.9.4.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/l  
b/commons-cli-1.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-codec  
1.11.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-collections-3.2.2.jar:/  
ome/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-compress-1.21.jar:/home/hado  
p/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-configuration2-2.1.1.jar:/home/hadoop/ha  
oop-3.3.2/share/hadoop/hdfs/lib/commons-daemon-1.0.13.jar:/home/hadoop/hadoop-3.3.2/  
hare/hadoop/hdfs/lib/commons-io-2.8.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdf

/lib/commons-lang3-3.12.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/common-logging-1.1.3.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-math3-3.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-net-3.6.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/commons-text-1.4.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/curator-client-4.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/curator-framework-4.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/curator-recipes-4.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/dnsjava-2.1.7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/failureaccess-1.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/gson-2.8.9.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/guava-27.0-jre.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/hadoop-annotations-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/hadoop-auth-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/hadoop-shaded-guava-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/hadoop-shaded-protobuf\_3\_7-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/httpclient-4.5.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/httpcore-4.4.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/j2objc-annotations-1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-annotations-2.10.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-core-2.13.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-core-asl-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-databind-2.13.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-jaxrs-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-mapper-asl-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jackson-xc-1.9.13.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jakarta.activation-api-1.2.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/javax.servlet-api-3.1.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jaxb-api-2.2.11.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jaxb-impl-2.2.3-1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jci-annotations-1.0-1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jersey-core-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jersey-json-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jersey-server-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jersey-servlet-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jettison-1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-http-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-io-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-security-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-server-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-servlet-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-util-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-util-ajax-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-webapp-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jetty-xml-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jsch-0.1.55.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/json-simple-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/json-smart-2.4.7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/jsr305-3.0.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/js311-api-1.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-admin-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-client-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-common-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-core-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-crypto-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-identity-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-server-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-simplekdc-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerb-util-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerby-asn1-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerby-config-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerby-pkix-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerby-util-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/kerby-xdr-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/leveldbjni-all-1.8.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/listenablefuture-9999.0-empty-t

-avoid-conflict-with-guava.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/log4j-1.2.7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/netty-3.10.6.Final.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/netty-all-4.1.68.Final.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/nimbus-jose-jwt-9.8.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/okhttp-2.7.5.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/okio-1.6.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/paranamer-2.3.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/protobuf-java-2.5.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/re2j-1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/snappy-java-1.1.8.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/stax2-api-4.2.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/token-provider-1.0.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/woodstox-core-5.3.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/zookeeper-3.5.6.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/lib/zookeeper-jute-3.5.6.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-3.3.2-tests.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-client-3.3.2-tests.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-client-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-https-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-native-client-3.3.2-tests.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-native-client-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-nfs-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-rbf-3.3.2-tests.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/hdfs/hadoop-hdfs-rbf-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-app-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-common-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-core-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-hs-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-hs-plugins-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-jobclient-3.3.2-tests.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-jobclient-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-nativetask-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-shuffle-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-client-uploader-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/mapreduce/hadoop-mapreduce-examples-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/HikariCP-java7-2.4.12.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/aopalliance-1.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/asm-analysis-9.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/asm-commons-9.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/asm-tree-9.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/bcprov-jdk15on-1.60.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/bcprov-jdk15on-1.60.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/ehcache-3.3.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/fst-2.50.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/geronimo-jcache\_1.0\_spec-1.0-alpha-1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/guice-4.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/guice-servlet-4.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jackson-jaxrs-base-2.13.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jackson-jaxrs-json-provider-2.13.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jackson-module-jaxb-annotations-2.13.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jakarta.xml.bind-api-2.3.3.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/java-util-1.9.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/javax-websocket-client-impl-9.4.3.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/javax-websocket-server-impl-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/javax.inject-1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/javax.websocket-api-1.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/javax.websocket-client-api-1.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/javax.ws.rs-api-2.1.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jersey-client-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib

```
b/jersey-guice-1.19.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jetty-annotation-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jetty-client-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jetty-jndi-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jetty-plus-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jline-3.9.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/jna-5.2.0.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/json-io-2.5.1.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/metrics-core-3.2.4.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/mssql-jdbc-6.2.1.jre7.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/objenesis-2.6.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/snakeyaml-1.26.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/swagger-annotations-1.5.4.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/websocket-api-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/websocket-client-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/websocket-common-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/websocket-server-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/lib/websocket-servlet-9.4.43.v20210629.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-api-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-applications-distributedshell-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-applications-mawo-core-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-applications-unmanaged-am-launcher-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-client-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-common-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-registry-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-applicationhistoryservice-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-common-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-nodemanager-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-resourcemanager-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-router-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-sharedcachemanager-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-tests-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-timeline-pluginstorage-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-server-web-proxy-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-services-api-3.3.2.jar:/home/hadoop/hadoop-3.3.2/share/hadoop/yarn/hadoop-yarn-services-core-3.3.2.jar
STARTUP_MSG: build = git@github.com:apache/hadoop.git -r 0bcb014209e219273cb6fd412df7df713cbac61; compiled by 'chao' on 2022-02-21T18:39Z
STARTUP_MSG: java = 1.8.0_322
***** /
2022-03-17 23:22:55,312 INFO namenode.NameNode: registered UNIX signal handlers for [TRM, HUP, INT]
2022-03-17 23:22:55,408 INFO namenode.NameNode: createNameNode [-format]
2022-03-17 23:22:55,800 INFO namenode.NameNode: Formatting using clusterid: CID-427170c-605c-44fe-be87-6cbbcbb60338
2022-03-17 23:22:55,834 INFO namenode.FSEditLog: Edit logging is async:true
2022-03-17 23:22:55,870 INFO namenode.FSNamesystem: KeyProvider: null
2022-03-17 23:22:55,872 INFO namenode.FSNamesystem: fsLock is fair: true
2022-03-17 23:22:55,873 INFO namenode.FSNamesystem: Detailed lock hold time metrics enabled: false
2022-03-17 23:22:55,886 INFO namenode.FSNamesystem: fsOwner = hadoop (auth:IMPLE)
2022-03-17 23:22:55,886 INFO namenode.FSNamesystem: supergroup = supergroup
2022-03-17 23:22:55,886 INFO namenode.FSNamesystem: isPermissionEnabled = true
2022-03-17 23:22:55,886 INFO namenode.FSNamesystem: isStoragePolicyEnabled = true
2022-03-17 23:22:55,886 INFO namenode.FSNamesystem: HA Enabled: false
2022-03-17 23:22:55,930 INFO common.Util: dfs.datanode.fileio.profiling.sampling.percentage
```

set to 0. Disabling file IO profiling  
2022-03-17 23:22:55,940 INFO blockmanagement.DatanodeManager: dfs.block.invalidate.limit configured=1000, counted=60, effected=1000  
2022-03-17 23:22:55,941 INFO blockmanagement.DatanodeManager: dfs.namenode.datanoderegistration.ip-hostname-check=true  
2022-03-17 23:22:55,944 INFO blockmanagement.BlockManager: dfs.namenode.startup.delay.block.deletion.sec is set to 000:00:00:00.000  
2022-03-17 23:22:55,944 INFO blockmanagement.BlockManager: The block deletion will start around 2022 Mar 17 23:22:55  
2022-03-17 23:22:55,947 INFO util.GSet: Computing capacity for map BlocksMap  
2022-03-17 23:22:55,947 INFO util.GSet: VM type = 64-bit  
2022-03-17 23:22:55,950 INFO util.GSet: 2.0% max memory 839.5 MB = 16.8 MB  
2022-03-17 23:22:55,950 INFO util.GSet: capacity =  $2^{21}$  = 2097152 entries  
2022-03-17 23:22:55,959 INFO blockmanagement.BlockManager: Storage policy satisfier is disabled  
2022-03-17 23:22:55,959 INFO blockmanagement.BlockManager: dfs.block.access.token.enable = false  
2022-03-17 23:22:55,968 INFO blockmanagement.BlockManagerSafeMode: dfs.namenode.safe.mode.threshold-pct = 0.999  
2022-03-17 23:22:55,968 INFO blockmanagement.BlockManagerSafeMode: dfs.namenode.safe.mode.min.datanodes = 0  
2022-03-17 23:22:55,968 INFO blockmanagement.BlockManagerSafeMode: dfs.namenode.safe.mode.extension = 30000  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: defaultReplication = 3  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: maxReplication = 5  
2  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: minReplication = 1  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: maxReplicationStreams  
2  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: redundancyRecheckInterval = 3000ms  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: encryptDataTransfer =  
alse  
2022-03-17 23:22:55,969 INFO blockmanagement.BlockManager: maxNumBlocksToLog  
= 1000  
2022-03-17 23:22:55,996 INFO namenode.FSDirectory: GLOBAL serial map: bits=29 maxEntries=536870911  
2022-03-17 23:22:55,996 INFO namenode.FSDirectory: USER serial map: bits=24 maxEntries=6777215  
2022-03-17 23:22:55,996 INFO namenode.FSDirectory: GROUP serial map: bits=24 maxEntries=16777215  
2022-03-17 23:22:55,996 INFO namenode.FSDirectory: XATTR serial map: bits=24 maxEntries16777215  
2022-03-17 23:22:56,023 INFO util.GSet: Computing capacity for map INodeMap  
2022-03-17 23:22:56,023 INFO util.GSet: VM type = 64-bit  
2022-03-17 23:22:56,023 INFO util.GSet: 1.0% max memory 839.5 MB = 8.4 MB  
2022-03-17 23:22:56,023 INFO util.GSet: capacity =  $2^{20}$  = 1048576 entries  
2022-03-17 23:22:56,024 INFO namenode.FSDirectory: ACLs enabled? true  
2022-03-17 23:22:56,024 INFO namenode.FSDirectory: POSIX ACL inheritance enabled? true  
2022-03-17 23:22:56,024 INFO namenode.FSDirectory: XAttrs enabled? true  
2022-03-17 23:22:56,025 INFO namenode.NameNode: Caching file names occurring more than 10 times  
2022-03-17 23:22:56,030 INFO snapshot.SnapshotManager: Loaded config captureOpenFiles: also, skipCaptureAccessTimeOnlyChange: false, snapshotDiffAllowSnapRootDescendant: true,

```
maxSnapshotLimit: 65536
2022-03-17 23:22:56,033 INFO snapshot.SnapshotManager: SkipList is disabled
2022-03-17 23:22:56,037 INFO util.GSet: Computing capacity for map cachedBlocks
2022-03-17 23:22:56,037 INFO util.GSet: VM type      = 64-bit
2022-03-17 23:22:56,037 INFO util.GSet: 0.25% max memory 839.5 MB = 2.1 MB
2022-03-17 23:22:56,037 INFO util.GSet: capacity    = 2^18 = 262144 entries
2022-03-17 23:22:56,047 INFO metrics.TopMetrics: NNTop conf: dfs.namenode.top.window.n
m.buckets = 10
2022-03-17 23:22:56,047 INFO metrics.TopMetrics: NNTop conf: dfs.namenode.top.num.users
= 10
2022-03-17 23:22:56,047 INFO metrics.TopMetrics: NNTop conf: dfs.namenode.top.windows.
inutes = 1,5,25
2022-03-17 23:22:56,051 INFO namenode.FSNamesystem: Retry cache on namenode is enabl
ed
2022-03-17 23:22:56,051 INFO namenode.FSNamesystem: Retry cache will use 0.03 of total h
ap and retry cache entry expiry time is 600000 millis
2022-03-17 23:22:56,053 INFO util.GSet: Computing capacity for map NameNodeRetryCache
2022-03-17 23:22:56,053 INFO util.GSet: VM type      = 64-bit
2022-03-17 23:22:56,053 INFO util.GSet: 0.029999999329447746% max memory 839.5 MB =
57.9 KB
2022-03-17 23:22:56,053 INFO util.GSet: capacity    = 2^15 = 32768 entries
2022-03-17 23:22:56,080 INFO namenode.FSImage: Allocated new BlockPoolId: BP-57158312
-192.168.153.11-1647530576071
2022-03-17 23:22:56,101 INFO common.Storage: Storage directory /home/hadoop/data/dfs/
ame has been successfully formatted.
2022-03-17 23:22:56,128 INFO namenode.FSImageFormatProtobuf: Saving image file /home/
adoop/data/dfs/name/current/fsimage.ckpt_00000000000000000000 using no compression
2022-03-17 23:22:56,226 INFO namenode.FSImageFormatProtobuf: Image file /home/hadoop
data/dfs/name/current/fsimage.ckpt_00000000000000000000 of size 401 bytes saved in 0 sec
nds .
2022-03-17 23:22:56,241 INFO namenode.NNStorageRetentionManager: Going to retain 1 im
ges with txid >= 0
2022-03-17 23:22:56,259 INFO namenode.FSNamesystem: Stopping services started for active
state
2022-03-17 23:22:56,260 INFO namenode.FSNamesystem: Stopping services started for stan
by state
2022-03-17 23:22:56,264 INFO namenode.FSImage: FSImageSaver clean checkpoint: txid=0 w
en meet shutdown.
2022-03-17 23:22:56,264 INFO namenode.NameNode: SHUTDOWN_MSG:
/*****
SHUTDOWN_MSG: Shutting down NameNode at node1/192.168.153.11
*****/
[hadoop@node1 ~]$
```

## 2、启动 HDFS 集群

```
start-dfs.sh
```

该脚本将会启动 NameNode 守护进程和 DataNode 守护进程：

```
[hadoop@node1 hadoop-3.3.2]$ start-dfs.sh
Starting namenodes on [node1]
Starting datanodes
node1.hadoop.com: Warning: Permanently added 'node1.hadoop.com' (ECDSA) to the list of
```



nown hosts.

```
node3.hadoop.com: ssh: Could not resolve hostname node3.hadoop.com: Name or service no known
```

```
node2.hadoop.com: ssh: Could not resolve hostname node2.hadoop.com: Name or service no known
```

Starting secondary namenodes [node2]

```
node2: WARNING: /home/hadoop/hadoop-3.3.2/logs does not exist. Creating.
```

```
[hadoop@node1 hadoop-3.3.2]$
```

```
[hadoop@node1 hadoop-3.3.2]$ jps
```

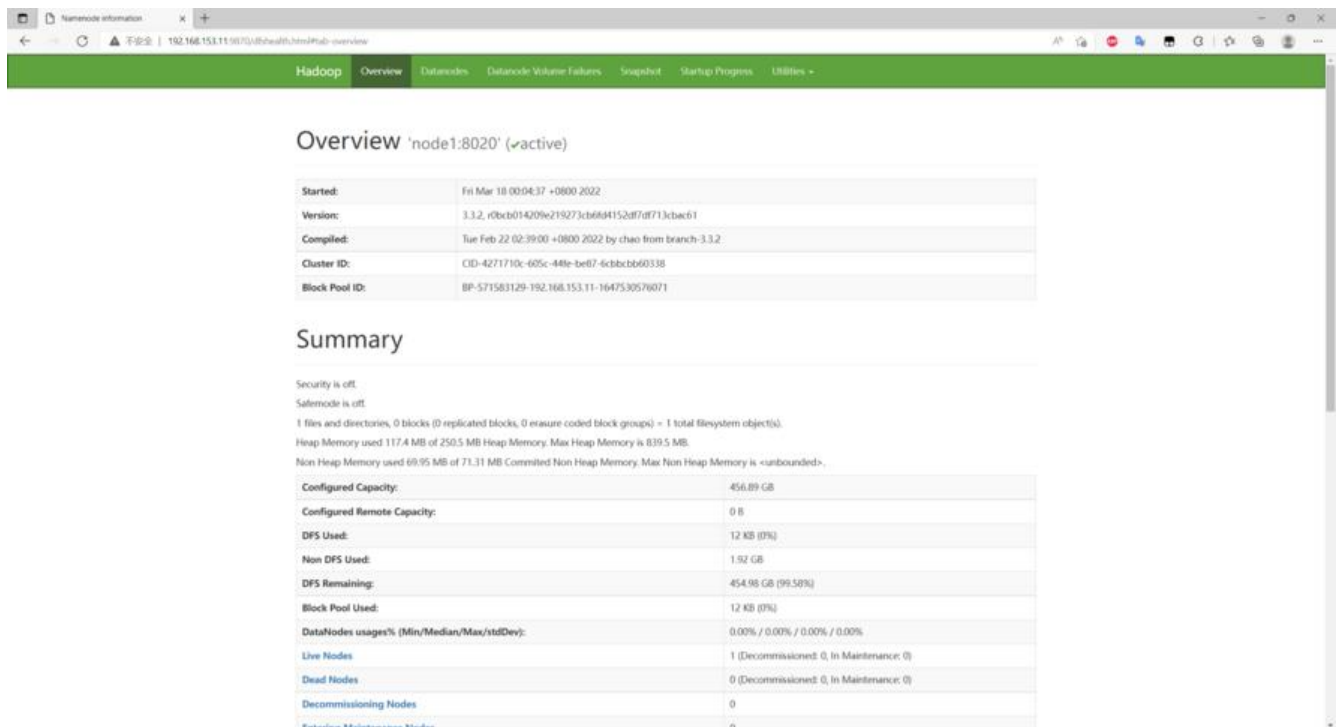
```
5001 DataNode
```

```
5274 Jps
```

```
4863 NameNode
```

```
[hadoop@node1 hadoop-3.3.2]$
```

启动成功后，可以在浏览器访问 NameNode 的 Web 界面（默认端口：9870）：



### 3、启动 YARN 集群

```
start-yarn.sh
```

该脚本将会启动 ResourceManager 守护进程和 NodeManager 守护进程：

```
[hadoop@node1 hadoop-3.3.2]$ start-yarn.sh
```

```
Starting resourcemanager
```

```
Starting nodemanagers
```

```
node3.hadoop.com: ssh: Could not resolve hostname node3.hadoop.com: Name or service no known
```

```
node2.hadoop.com: ssh: Could not resolve hostname node2.hadoop.com: Name or service no known
```

```
[hadoop@node1 hadoop-3.3.2]$
```

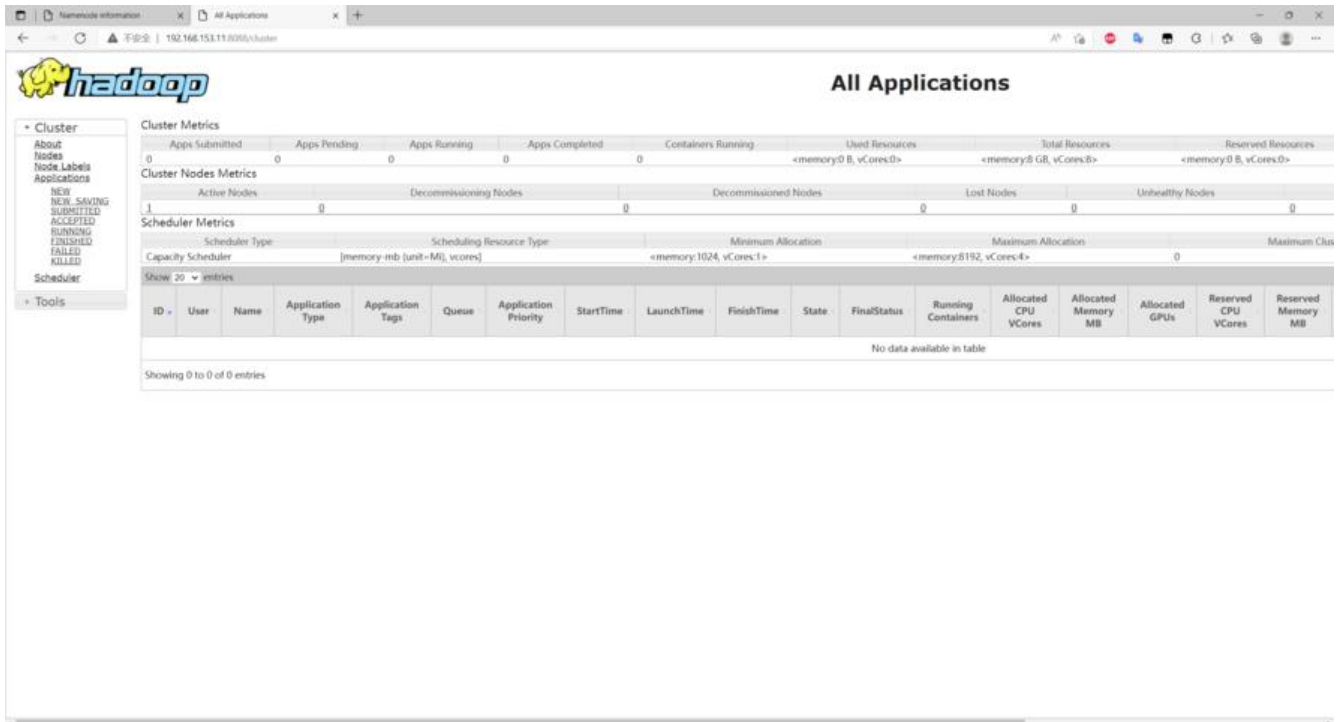
```
[hadoop@node1 hadoop-3.3.2]$ jps
```

```
5536 NodeManager
```

```
5395 ResourceManager
```

```
5001 DataNode
5867 Jps
4863 NameNode
[hadoop@node1 hadoop-3.3.2]$
```

启动成功后，可以在浏览器访问 ResourceManager 的 Web 界面（默认端口：8088）：



除了 `start-dfs.sh` 和 `start-yarn.sh` 脚本，也可使用 `start-all.sh` 脚本，一次性启动 Hadoop 的所有程。

## 停止集群

和启动集群一样，Hadoop 提供了两种方式停止集群。

逐个终止进程的命令：

```
# HDFS 集群
$HADOOP_HOME/bin/hdfs --daemon stop namenode | datanode | secondarynamenode
```

```
# YARN 集群
$HADOOP_HOME/bin/yarn --daemon stop resourcemanager | nodemanager | proxyserver
```

停止集群的脚本：

- HDFS 集群 -- `$HADOOP_HOME/sbin/stop-dfs.sh`，一键终止 HDFS 集群的所有进程。
- YARN 集群 -- `$HADOOP_HOME/sbin/stop-yarn.sh`，一键终止 YARN 集群的所有进程
- Hadoop 集群 -- `$HADOOP_HOME/sbin/stop-all.sh`，一键终止 HDFS 集群和 YARN 集群的有进程。

使用 `stop-all.sh` 脚本，一次性停止 Hadoop 的所有进程。

```
[hadoop@node1 hadoop-3.3.2]$ stop-all.sh
```

```
WARNING: Stopping all Apache Hadoop daemons as hadoop in 10 seconds.
WARNING: Use CTRL-C to abort.
Stopping namenodes on [node1]
Stopping datanodes
node2.hadoop.com: ssh: Could not resolve hostname node2.hadoop.com: Name or service no
known
node3.hadoop.com: ssh: Could not resolve hostname node3.hadoop.com: Name or service no
known
Stopping secondary namenodes [node2]
Stopping nodemanagers
node3.hadoop.com: ssh: Could not resolve hostname node3.hadoop.com: Name or service no
known
node2.hadoop.com: ssh: Could not resolve hostname node2.hadoop.com: Name or service no
known
Stopping resourcemanager
[hadoop@node1 hadoop-3.3.2]$
```

## 相关资料

[Hadoop: Setting up a Single Node Cluster](#)

[Hadoop Cluster Setup](#)

[How To Install Apache Hadoop / HBase on CentOS 7](#)

[2022最新黑马程序员大数据Hadoop入门视频教程\\_哔哩哔哩\\_bilibili](#)