

SolrCloudTomcat

Example setup for 3 node solr cloud cluster with 3 node Zookeeper ensemble

Solr Cloud machines : s1.solr.com,s2.solr.com,s3.solr.com

Zookeeper machines: z1.com, z2.com, z3.com

The default directory for all the setup /home/myuser/

First bring up zookeeper ensemble

Do following steps for z1.com,z2.com,z3.com to setup zookeeper

1) Download a [ZooKeeper](http://hadoop.apache.org/zookeeper/releases.html) from <http://hadoop.apache.org/zookeeper/releases.html>

2) Untar the zookeeper inside zookeeper directory. tar -xf -C /home/myuser/zookeeper/

3) Edit zookeeper config file

vi /home/myuser/zookeeper/conf/zoo.cfg and put something similar to

```
# The number of milliseconds of each tick tickTime=2000
```

```
# The number of ticks that the initial synchronization phase can take initLimit=10
```

```
# The number of ticks that can pass between
```

```
# sending a request and getting an acknowledgement syncLimit=5
```

```
# the directory where the snapshot is stored.
```

```
# Choose appropriately for your environment dataDir=/home/myuser/zookeeper/data/
```

```
# the port at which the clients will connect clientPort=2181
```

```
# the directory where transaction log is stored.
```

```
# this parameter provides dedicated log device for ZooKeeper dataLogDir=/home/myuser/zookeeper/log/
```

```
# ZooKeeper server and its port no.
```

```
# ZooKeeper ensemble should know about every other machine in the ensemble
```

```
# specify server id by creating 'myid' file in the dataDir
```

```
# use hostname instead of IP address for convenient maintenance server.1=z1.com:2888:3888 server.2=z2.com:2888:3888 server.3=z3.com:2888:3888
```

4) For each server create an id file in data directory and put appropriate id in the file. For example

vi /home/myuser/zookeeper/data/myid and write just 1 for z1.com in the file. For z2.com server write 2 in its myid and 3 in z3.com myid file. myid file is used to identify the zookeeper server.

4)Now start zookeeper for each server.

```
# cd /home/myuser/zookeeper/bin/
```

```
# ./zkServer.sh start
```

Setup tomcat

Now once zookeeper cluster is ready lets setup Solr Cloud machines : s1.solr.com,s2.solr.com,s3.solr.com

Lets assume all the 3 solr machines have tomcat installed in /home/myuser/tomcat/

First lets load solr config from one of the solr machine

1) Get solr.war to s1.solr.com machine

2) Now lets copy solr.war i to solr-war folder – cp solr.war /home/myuser/solr-war/

3) cd /home/myuser/solr-war/

4) jar xvf solr.war

5) cp WEB-INF/lib/* /home/myuser//solr-war-lib/

6) get all the solrconfig files like solrconfig.xml, schema.xml etc in /home/myuser/solr-config/

7) Upload config to zookeeper using solr zookeeper cli

```
java -classpath ./home/myuser/solr-war-lib/* org.apache.solr.cloud.ZkCLI -cmd upconfig -zkhost z1.com:2181,z2.com:2181,z3.com:2181 -confdir /home/myuser/solr-config -confname myconf
```

8) Link uploaded config with target collection (assume new collection name will be mycollection)

```
java -classpath ./home/myuser/solr-war-lib/* org.apache.solr.cloud.ZkCLI -cmd linkconfig -collection mycollection -confname myconf -zkhost z1.com:2181,
z2.com:2181,z3.com:2181
```

Now in all the solr machines do following

9) create /home/myuser/solr-cores/ for storing the solr index

10) create solr.xml in solr-cores

```
<?xml version="1.0" encoding="UTF-8" ?>
<solr persistent="true" sharedLib="lib">
<cores adminPath="/admin/cores" zkClientTimeout="20000" hostPort="8080" hostContext="solr">
</cores>
</solr>
```

11) Start the tomcat with following jvm options

```
-DzkHost=z1.com:2181,z2.com:2181,z3.com:2181
```

12) Put solr.war in /home/myuser/tomcat/webapps/ and restart tomcat

13) edit /home/myuser/tomcat/webapps/solr/META-INF/context.xml and restart tomcat

```
<?xml version="1.0" encoding="utf-8"?>
<Context docBase="solr.war" debug="0" crossContext="false">
<Environment name="solr/home" type="java.lang.String" value="/home/myuser/solr-cores" override="true"/>
</Context>
```

14) curl 'http://s1.solr.com:8080/solr/admin/collections?action=CREATE&name=mycollection&numShards=3&replicationFactor=1'