

Redis Cluster 搭建与测试 (一)

作者: [wangsch](#)

原文链接: <https://ld246.com/article/1477019413416>

来源网站: [链滴](#)

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

Redis Cluster搭建与测试 (一)

OS: CentOS 6.5 64位

Redis版本: redis 3.2.4, 官网: <http://redis.io>

安装

安装

下载源码包后, 解压, 看里面的README.md文件, 有更具体的编译和安装步骤。我的步骤:

```
cd /opt/src
```

```
tar -xvzf redis-3.2.4.tar.gz
```

```
cd redis-3.2.4
```

```
make
```

```
make test
```

```
make PREFIX=/opt/redis-3.2.4 install
```

```
ln -s /opt/redis-3.2.4/ /opt/redis
```

最后两步, 执行安装的时候, 指定特定的安装目录。最后, 做一个软连接。这两步, 主要是之后升级redis版本方便。

注: "make test"可能会报错"You need tcl 8.5...", 解决方法, 文章末尾; 附: 安装tcl

测试Redis Cluster

创建并启动

```
cd /opt/src/redis-3.2.4/src
```

```
./redis-trib.rb create --replicas 1 127.0.0.1:7000 127.0.0.1:7001 127.0.0.1:7002 127.0.0.1:7003 127.0.0.1:7004 127.0.0.1:7005
```

成功输出:

```
[OK] All 16384 slots covered.
```

7000端口redis日志输出:

```
25979:M 21 Oct 14:10:28.444 # configEpoch set to 1 via CLUSTER STATE-CONFIG-EPOCH
```

```
25979:M 21 Oct 14:10:28.544 # IP address for this node updated to 127.0.0.1
```

```
25979:M 21 Oct 14:10:32.854 * Slave 127.0.0.1:7003 asks for synchronization
```

```
25979:M 21 Oct 14:10:32.854 * Full resync requested by slave 127.0.0.1:7003
```

```
25979:M 21 Oct 14:10:32.854 * Starting BGSAVE for SYNC with target: disk
```

```
25979:M 21 Oct 14:10:32.855 * Background saving started by pid 26408
```

```
26408:C 21 Oct 14:10:32.874 * DB saved on disk
```

```
26408:C 21 Oct 14:10:32.875 * RDB: 6 MB of memory used by copy-on-write
```

```
25979:M 21 Oct 14:10:32.952 * Background saving terminated with success
```

```
25979:M 21 Oct 14:10:32.953 * Synchronization with slave 127.0.0.1:7003 succeeded
```

```
25979:M 21 Oct 14:10:33.352 # Cluster state changed: ok
```

7003端口redis输出:

```
26003:M 21 Oct 14:10:28.445 # configEpoch set to 4 via CLUSTER STATE-CONFIG-EPOCH
```


OK

```
127.0.0.1:7001> dbsize
```

```
(integer) 1
```

```
127.0.0.1:7001> get key1
```

```
"value1"
```

<p>可以看出，操作被自动重定向</p>

<p> </p>

<h3>测试failover</h3>

<p>逐步关闭redis cluster中的redis实例，看看是什么效果。 </p>

<p> </p>

<p>1. 关闭7000的redis实例 </p>

<p>7001</p>

```
<pre class="brush: bash">25986:M 21 Oct 14:28:16.408 * Marking node 86ca55c465645c184f1dbc48336e0aae5620948 as failing (quorum reached).
```

```
25986:M 21 Oct 14:28:16.408 # Cluster state changed: fail
```

```
25986:M 21 Oct 14:28:17.217 # Failover auth granted to eb5000f25af13145a04ee8ec8b9ecfe826ae833 for epoch 7
```

```
25986:M 21 Oct 14:28:17.256 # Cluster state changed: ok</pre>
```

<p> </p>

<p>7002</p>

```
<pre class="brush: bash">25996:M 21 Oct 14:28:16.408 * Marking node 86ca55c465645c184f1dbc48336e0aae5620948 as failing (quorum reached).
```

```
25996:M 21 Oct 14:28:16.408 # Cluster state changed: fail
```

```
25996:M 21 Oct 14:28:17.217 # Failover auth granted to eb5000f25af13145a04ee8ec8b9ecfe826ae833 for epoch 7
```

```
25996:M 21 Oct 14:28:17.256 # Cluster state changed: ok</pre>
```

<p> </p>

<p>7003是7000的slave，控制台输出： </p>

```
<pre class="brush: bash">26003:S 21 Oct 14:28:10.288 # Connection with master lost.
```

```
26003:S 21 Oct 14:28:10.289 * Caching the disconnected master state.
```

```
26003:S 21 Oct 14:28:10.891 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:10.891 * MASTER &lt;-&gt; SLAVE sync started
```

```
26003:S 21 Oct 14:28:10.891 # Error condition on socket for SYNC: Connection refused
```

```
26003:S 21 Oct 14:28:11.893 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:11.893 * MASTER &lt;-&gt; SLAVE sync started
```

```
26003:S 21 Oct 14:28:11.893 # Error condition on socket for SYNC: Connection refused
```

```
26003:S 21 Oct 14:28:12.896 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:12.896 * MASTER &lt;-&gt; SLAVE sync started
```

```
26003:S 21 Oct 14:28:12.896 # Error condition on socket for SYNC: Connection refused
```

```
26003:S 21 Oct 14:28:13.898 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:13.898 * MASTER &lt;-&gt; SLAVE sync started
```

```
26003:S 21 Oct 14:28:13.898 # Error condition on socket for SYNC: Connection refused
```

```
26003:S 21 Oct 14:28:14.903 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:14.903 * MASTER &lt;-&gt; SLAVE sync started
```

```
26003:S 21 Oct 14:28:14.903 # Error condition on socket for SYNC: Connection refused
```

```
26003:S 21 Oct 14:28:15.905 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:15.905 * MASTER &lt;-&gt; SLAVE sync started
```

```
26003:S 21 Oct 14:28:15.905 # Error condition on socket for SYNC: Connection refused
```

```
26003:S 21 Oct 14:28:16.307 * Marking node 86ca55c465645c184fd1dbc48336e0aae5620948 as failing (quorum reached).
```

```
26003:S 21 Oct 14:28:16.308 # Cluster state changed: fail
```

```
26003:S 21 Oct 14:28:16.407 # Start of election delayed for 713 milliseconds (rank #0, offset 146).
```

```
26003:S 21 Oct 14:28:16.909 * Connecting to MASTER 127.0.0.1:7000
```

```
26003:S 21 Oct 14:28:16.909 * MASTER &lt;-&gt; SLAVE sync started
26003:S 21 Oct 14:28:16.909 # Error condition on socket for SYNC: Connection refused
26003:S 21 Oct 14:28:17.210 # Starting a failover election for epoch 7.
26003:S 21 Oct 14:28:17.217 # Failover election won: I'm the new master.
26003:S 21 Oct 14:28:17.217 # configEpoch set to 7 after successful failover
26003:M 21 Oct 14:28:17.217 * Discarding previously cached master state.
26003:M 21 Oct 14:28:17.217 # Cluster state changed: ok</pre>
<p>&nbsp;</p>
<p>7004</p>
<pre class="brush: bash">26010:S 21 Oct 14:28:16.307 * Marking node 86ca55c465645c184f
1dbc48336e0aae5620948 as failing (quorum reached).
26010:S 21 Oct 14:28:16.307 # Cluster state changed: fail
26010:S 21 Oct 14:28:17.257 # Cluster state changed: ok</pre>
<p>&nbsp;</p>
<p>7005</p>
<pre class="brush: bash">26020:S 21 Oct 14:28:16.407 * Marking node 86ca55c465645c184f
1dbc48336e0aae5620948 as failing (quorum reached).
26020:S 21 Oct 14:28:16.407 # Cluster state changed: fail
26020:S 21 Oct 14:28:17.258 # Cluster state changed: ok</pre>
<p>&nbsp;</p>
<p>client &ldquo;CLUSTER NODES&rdquo;命令输出: </p>
<pre class="brush: bash">127.0.0.1:7002&gt; CLUSTER NODES
be619fa382c8471eae6222d735ed05e6cab5e7a 127.0.0.1:7004 slave 688a1ac1bae1a24327cd
6ac4b9efc69888885c9 0 1477031314864 5 connected
eb5000f25af13145a04ee8ec8b9ecfe8426ae833 127.0.0.1:7003 master - 0 1477031314363 7 c
nected 0-5460
86ca55c465645c184fd1dbc48336e0aae5620948 127.0.0.1:7000 master,fail - 1477031290388 1
77031289787 1 disconnected
088390b994557c8a94cab38cfded4ec60fce724 127.0.0.1:7005 slave 8dc1c8679a46ba85750a
d388937932689e4963f 0 1477031313361 6 connected
688a1ac1bae1a24327cdd6ac4b9efc69888885c9 127.0.0.1:7001 master - 0 1477031315366 2 c
nected 5461-10922
8dc1c8679a46ba85750a5d388937932689e4963f 127.0.0.1:7002 myself,master - 0 0 3 connect
d 10923-16383 </pre>
<p>&nbsp;</p>
<p><strong>2. 再退出一个slave, redis实例7004</strong></p>
<p>剩下的7001、7002、7003、7005均会收到类似消息: </p>
<pre class="brush: bash">26020:S 21 Oct 14:34:09.473 * FAIL message received from 8dc1c8
79a46ba85750a5d388937932689e4963f about be619fa382c8471eae6222d735ed05e6cab5e
a</pre>
<p>&nbsp;</p>
<p><strong>3. 再退出一个master, <strong>redis实例</strong>7001</strong></p>
<p>剩下的7002、7003、7005均会收到类似消息: <strong><br /></strong></p>
<pre class="brush: bash">25996:M 21 Oct 14:37:00.001 * Marking node 688a1ac1bae1a2432
cdd6ac4b9efc698888885c9 as failing (quorum reached).
25996:M 21 Oct 14:37:00.001 # Cluster state changed: fail</pre>
<p><br />过去一段时间, 保持在&ldquo;Cluster state changed: fail&rdquo;这个状态不再变动
此刻, redis cluster已经完全不可用。client执行命令: </p>
<pre class="brush: bash">127.0.0.1:7002&gt; set a b
(error) CLUSTERDOWN The cluster is down</pre>
<p>&nbsp;</p>
<p>QUES: Redis Cluster如何平滑的关闭和重启, 是否有这个需求? </p>
<p>&nbsp;</p>
<p>&nbsp;</p>
```

```
<p>-----</p>
<p>&nbsp;</p>
<h2>附</h2>
<h3>安装tcl</h3>
<p>编译redis后，推荐&ldquo;make test&rdquo;测试编译结果，然而： </p>
<p><em>[root@wangsch-serv redis-3.2.4]# make test</em><br /><em>cd src &amp;&am
; make test</em><br /><em>make[1]: Entering directory `/opt/src/redis-3.2.4/src'</em><br
/><em>You need tcl 8.5 or newer in order to run the Redis test</em><br /><em>make[1]: *
* [test] 错误 1</em><br /><em>make[1]: Leaving directory `/opt/src/redis-3.2.4/src'</em><
r /><em>make: *** [test] 错误 2</em></p>
<p>"You need tcl 8.5 or newer in order to run the Redis test"，说明要执行这个测试，还要依赖t
l，要下载最新版并编译安装。官网：<a href="http://www.tcl.tk">http://www.tcl.tk</a>，在<a h
ef="https://sourceforge.net/projects/tcl/">https://sourceforge.net/projects/tcl/</a>下载</p>
<p>最新版源码包下载到： /opt/src/tcl8.6.6-src.tar.gz，解压后进到解压目录的unix子目录，然后
接编译安装三板斧。</p>
<pre class="brush: bash">cd tcl8.6.6/unix
```

```
./configure &amp;&amp; make &amp;
&amp; make install </pre>
```

```
<p>&nbsp;</p>
<p>再次执行&ldquo;make test&rdquo;，成功： </p>
<p><em>All tests passed without errors!</em><br /><em>Cleanup: may take some time...
OK</em></p>
<p>&nbsp;</p>
<h3>安装redis为daemon</h3>
<p>可选步骤，可以通过&ldquo;service&rdquo;命令管理redis。进入到redis解压目录的utils目录
执行：</p>
<pre class="brush: java">./install_server.sh
Welcome to the redis service installer
This script will help you easily set up a running redis server
```

```
Please select the redis port for this instance: [6379]
Selecting default: 6379
Please select the redis config file name [/etc/redis/6379.conf] /opt/redis-conf/6379
Please select the redis log file name [/var/log/redis_6379.log] /data/log/redis/6379
Please select the data directory for this instance [/var/lib/redis/6379] /data/redis/6379
Please select the redis executable path [/usr/local/bin/redis-server] /opt/redis/bin/redis-server
Selected config:
Port      : 6379
Config file  : /opt/redis-conf/6379
Log file    : /data/log/redis/6379
Data dir    : /data/redis/6379
Executable  : /opt/redis/bin/redis-server
Cli Executable : /usr/local/bin/redis-cli
Is this ok? Then press ENTER to go on or Ctrl-C to abort.
```

Copied /tmp/6379.conf => /etc/init.d/redis_6379

Installing service...

Successfully added to chkconfig!

Successfully added to runlevels 345!

Starting Redis server...

Installation successful! </pre>

<p>上面的步骤，采用交互方式，指定端口号、数据目录等配置项。最后点击确认后，生成脚本文件 “/etc/init.d/redis_<portnumber>”， portnumber为redis实例端口号。 </p>

```
<pre class="brush: bash">service redis_6379 status
```

```
Redis is running (25080)</pre>
```

<p> 可以看到redis进程已经启动，支持“start”、 “stop”、 &lquo;stop”和“status”四种指令。 </p>

<p> </p>

<p> </p>