



```

0593:M 21 Jun 2020 01:31:41.933 * Background
success
0593:M 21 Jun 2020 01:31:41.934 * Synchror
0.0.1:6381 succeeded
C6593:signal-handler (1592730726) Receiv
0593:M 21 Jun 2020 02:12:06.372 * User rec
0593:M 21 Jun 2020 02:12:06.373 * Calling
0593:M 21 Jun 2020 02:12:06.373 * Saving t
0593:M 21 Jun 2020 02:12:06.375 * DB saver
0593:M 21 Jun 2020 02:12:06.375 # Redis i
re...
[root@localhost redis]#

File Edit View Search Terminal Help
6644:X 21 Jun 2020 01:32:55.473 # WARNING: The TCP backlog setting of 511 cannot be enforced because /proc/sys/net/core/somaxconn is set to the lower value of 128.
6644:X 21 Jun 2020 01:32:55.475 # Sentinel ID is 86419e566559c83cd8778ece47e59ec7d73d1e
6644:X 21 Jun 2020 01:32:55.475 # +monitor master mymaster 127.0.0.1 6379 quorum 2
6644:X 21 Jun 2020 01:32:55.476 # +slave slave 127.0.0.1:6380 127.0.0.1 6380 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:32:55.477 # +slave slave 127.0.0.1:6381 127.0.0.1 6381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:34:24.380 # +sentinel sentinel f2d402c35dbf99cc09a9d4d4d7e98335d5e8019 127.0.0.1 26380 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:35:35.211 # +sentinel sentinel 4bfc8f73a6ee913443d0c95d43df7ef99d2d43a 127.0.0.1 26381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 02:12:36.464 # +sdown master mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 02:12:36.549 # +new-epoch 1
6644:X 21 Jun 2020 02:12:36.551 # +vote-for-leader 4bfc8f73a6ee913443d0c95d43df7ef99d2d43a 1
6644:X 21 Jun 2020 02:12:36.555 # +sdown master mymaster 127.0.0.1 6379 quorum 2/2
6644:X 21 Jun 2020 02:12:36.555 # next failover delay: 1 will not start a failover before Sun Jun 21 02:18:37 2020
6644:X 21 Jun 2020 02:12:37.113 # +config-update-from sentinel 4bfc8f73a6ee913443d0c95d43df7ef99d2d43a 127.0.0.1 26381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 02:12:37.113 # +switch-master mymaster 127.0.0.1 6379 127.0.0.1 6381
6644:X 21 Jun 2020 02:12:37.113 # +slave slave 127.0.0.1:6380 127.0.0.1 6380 @ mymaster 127.0.0.1 6381
6644:X 21 Jun 2020 02:12:37.213 # +slave slave 127.0.0.1:6379 127.0.0.1 6379 @ mymaster 127.0.0.1 6381
6644:X 21 Jun 2020 02:13:07.204 # +sdown slave 127.0.0.1:6379 127.0.0.1 6379 @ mymaster 127.0.0.1 6381

File Edit View Search Terminal Help
6677:X 21 Jun 2020 02:12:38.104 # +slave-reconf-inprog slave 127.0.0.1 6380 @ mymaster 127.0.0.1 6379
6677:X 21 Jun 2020 02:12:38.104 # +slave-reconf-done slave 127.0.0.1 6380 @ mymaster 127.0.0.1 6379
6677:X 21 Jun 2020 02:12:38.104 # +failover-end master mymaster 127.0.0.1 6379
6677:X 21 Jun 2020 02:12:38.104 # +switch-master mymaster 127.0.0.1 6381
6677:X 21 Jun 2020 02:12:38.104 # +slave slave 127.0.0.1:6380 127.0.0.1 6380 @ mymaster 127.0.0.1 6381
6677:X 21 Jun 2020 02:12:38.104 # +slave slave 127.0.0.1:6379 127.0.0.1 6379 @ mymaster 127.0.0.1 6381
6677:X 21 Jun 2020 02:13:00.171 # +sdown slave 127.0.0.1:6379 127.0.0.1 6379 @ mymaster 127.0.0.1 6381

# is possible to just rename a command to itself:
#
# SENTINEL rename-command mymaster CONFIG CONFIG
[root@localhost redis]# cat sentinel.conf |grep -v '#' |grep -v '###'
port 26379
daemonize no
pidfile /var/run/redis-sentinel.pid
logfile ""
dir /tmp
sentinel monitor mymaster 127.0.0.1 6379 2
sentinel down-after-milliseconds mymaster 30000
sentinel parallel-syncs mymaster 1
sentinel failover-timeout mymaster 180000
sentinel deny-scripts-reconfig yes
[root@localhost redis]#

```

一、引入:

□

二、哨兵

1.定义

监控

选择

2.作用

监控:

- 通知 (提醒)
- 自动故障转移

PS

3.配置哨兵

conf文件说明：

4.启动哨兵命令

```
[root@localhost redis]# redis-sentinel sentinel26379.conf
6644:X 21 Jun 2020 01:32:55.470 # 0000000000000000 redis is starting o000o000o000o
6644:X 21 Jun 2020 01:32:55.470 # Redis version=5.0.3, bits=64, commit=00000000, modified=0, pid=6644,
just started
6644:X 21 Jun 2020 01:32:55.471 # Configuration loaded
6644:X 21 Jun 2020 01:32:55.471 * Increased maximum number of open files to 10032 (it was originally se
t to 1024).

Redis 5.0.3 (00000000/0) 64 bit

Running in sentinel mode
Port: 26379
PID: 6644

http://redis.io

6644:X 21 Jun 2020 01:32:55.473 # WARNING: The TCP backlog setting of 511 cannot be enforced because /p
roc/sys/net/core/somaxconn is set to the lower value of 128.
6644:X 21 Jun 2020 01:32:55.475 2 Sentinel ID is 86419e56655a9c03cd8778ec647e85aec7d73d1e
6644:X 21 Jun 2020 01:32:55.475 3 +monitor master mymaster 127.0.0.1 6379 quorum 2
6644:X 21 Jun 2020 01:32:55.476 * +slave slave 127.0.0.1:6380 127.0.0.1 6380 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:32:55.477 4 +slave slave 127.0.0.1:6381 127.0.0.1 6381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:34:24.380 * +sentinel sentinel f2d4b2e558bf99cccd9a8d4a0d7e90835d5e8b19 127.0.0.1
26380 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:35:25.211 5 +sentinel sentinel 4befc8f73a6ee913443d0c95d43df7ef99d2d43a 127.0.0.1
26381 @ mymaster 127.0.0.1 6379
```

图解：

修改前:

```
[root@localhost redis]# cat sentinel.conf |grep -v "#" |grep -v "^$"  
port 26379  
daemonize no  
pidfile /var/run/redis-sentinel.pid  
logfile ""  
dir /tmp  
sentinel monitor mymaster 127.0.0.1 6379 2  
sentinel down-after-milliseconds mymaster 30000  
sentinel parallel-syncs mymaster 1  
sentinel failover-timeout mymaster 180000  
sentinel deny-scripts-reconfig yes
```

修改后:

```
[root@localhost redis]# cat sentinel26379.conf  
port 26379  
daemonize no  
pidfile "/var/run/redis-sentinel.pid"  
logfile ""  
dir "/tmp"  
sentinel myid 86419e56655a9c03cd8778ec647e85aec7d73d1e  
sentinel deny-scripts-reconfig yes  
sentinel monitor mymaster 127.0.0.1 6379 2  
sentinel config-epoch mymaster 0  
sentinel leader-epoch mymaster 0  
# Generated by CONFIG REWRITE  
protected-mode no  
sentinel known-replica mymaster 127.0.0.1 6380  
sentinel known-replica mymaster 127.0.0.1 6381  
sentinel known-sentinel mymaster 127.0.0.1 26381 4befc8f73a6ee913443d0c95d43df7ef99d2d43a  
sentinel known-sentinel mymaster 127.0.0.1 26380 f2d4b2e558bf99cced9a8d4a0d7e90835d5e8b19  
sentinel current-epoch 0  
[root@localhost redis]#
```

5.哨兵工作原理

过程：监控、通知、故障转移

```
Redis 5.0.3 (00000000/0) 64 bit
Running in sentinel mode
Port: 26379
PID: 6644

http://redis.io

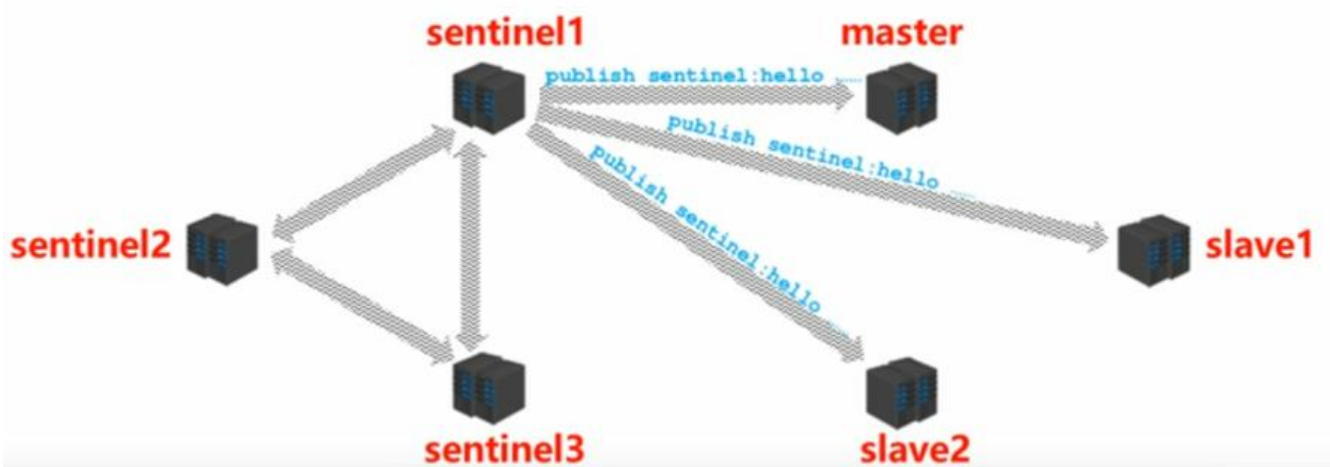
6644:X 21 Jun 2020 01:32:55.473 # WARNING: The TCP backlog setting of 511 cannot be enforced because /proc/sys/net/core/somaxconn is set to the lo
6644:X 21 Jun 2020 01:32:55.475 # Sentinel ID is 86419e56655a9c03cd8778ec647e85aec7d73d1e
6644:X 21 Jun 2020 01:32:55.475 # +monitor master mymaster 127.0.0.1 6379 quorum 2
6644:X 21 Jun 2020 01:32:55.476 * +slave slave 127.0.0.1:6380 127.0.0.1 6380 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:32:55.477 * +slave slave 127.0.0.1:6381 127.0.0.1 6381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:34:24.380 * +sentinel sentinel f2d4b2e558bf99cced9a8d4a0d7e90835d5e8b19 127.0.0.1 26380 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 01:35:25.211 * +sentinel sentinel 4befc8f73a6ee913443d0c95d43df7ef99d2d43a 127.0.0.1 26381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 02:12:36.464 # +sdown master mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 02:12:36.549 # +new-epoch 1
6644:X 21 Jun 2020 02:12:36.551 # +vote-for-leader 4befc8f73a6ee913443d0c95d43df7ef99d2d43a 1
6644:X 21 Jun 2020 02:12:36.555 # +odown master mymaster 127.0.0.1 6379 #quorum 3/2
6644:X 21 Jun 2020 02:12:36.555 # Next failover delay: I will not start a failover before Sun Jun 21 02:18:37 2020
6644:X 21 Jun 2020 02:12:37.113 # +config-update-from sentinel 4befc8f73a6ee913443d0c95d43df7ef99d2d43a 127.0.0.1 26381 @ mymaster 127.0.0.1 6379
6644:X 21 Jun 2020 02:12:37.113 # +switch-master mymaster 127.0.0.1 6379 127.0.0.1 6381
6644:X 21 Jun 2020 02:12:37.113 * +slave slave 127.0.0.1:6380 127.0.0.1 6380 @ mymaster 127.0.0.1 6381
6644:X 21 Jun 2020 02:12:37.113 * +slave slave 127.0.0.1:6379 127.0.0.1 6379 @ mymaster 127.0.0.1 6381
6644:X 21 Jun 2020 02:13:07.204 # +sdown slave 127.0.0.1:6379 127.0.0.1 6379 @ mymaster 127.0.0.1 6381
```

阶段一：监控阶段

作用：

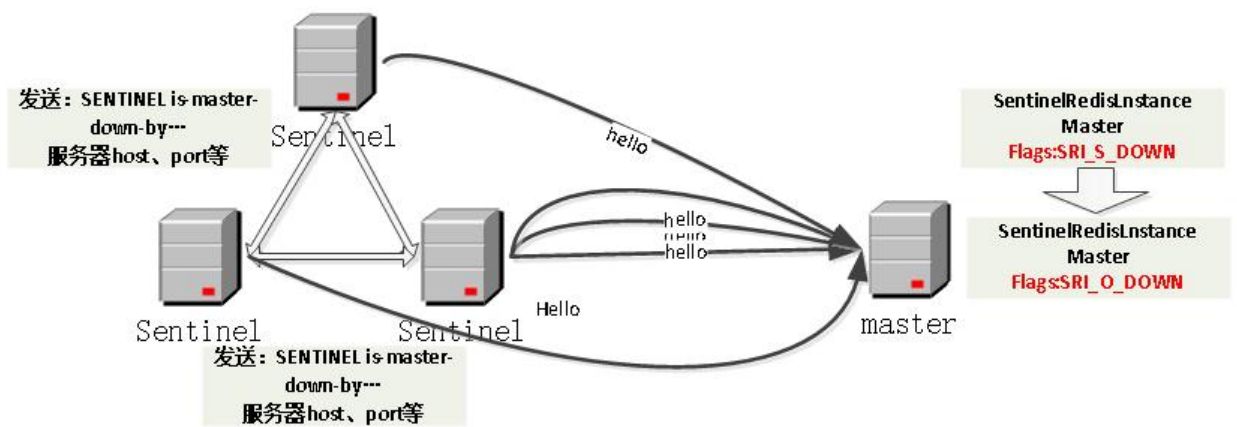
图解:

阶段二：通知阶段



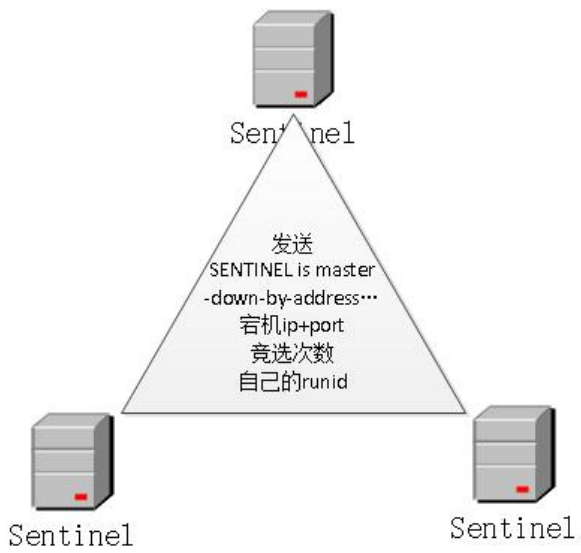
阶段三：故障转移阶段

过程一



图解

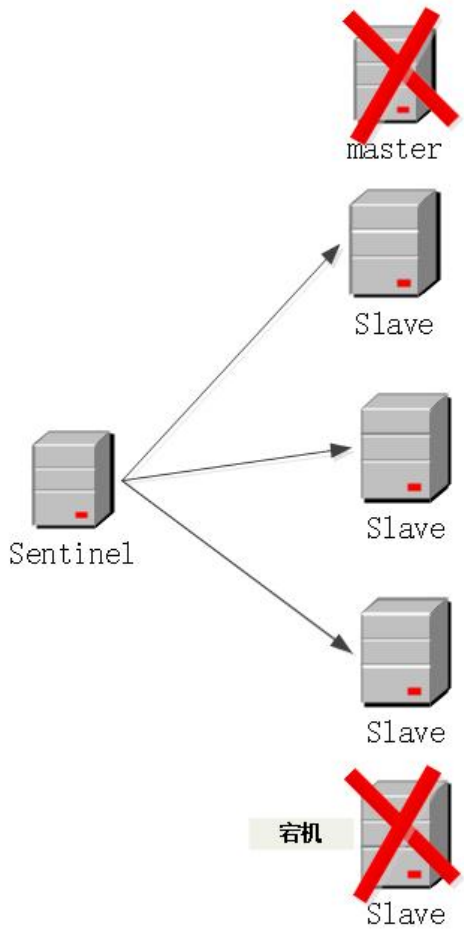
过程二



过程三：

sentinel挑选master的规则:

确定以后发指令:



故事总结：


□

□

□

三、集群

1.定义:



2.作用:

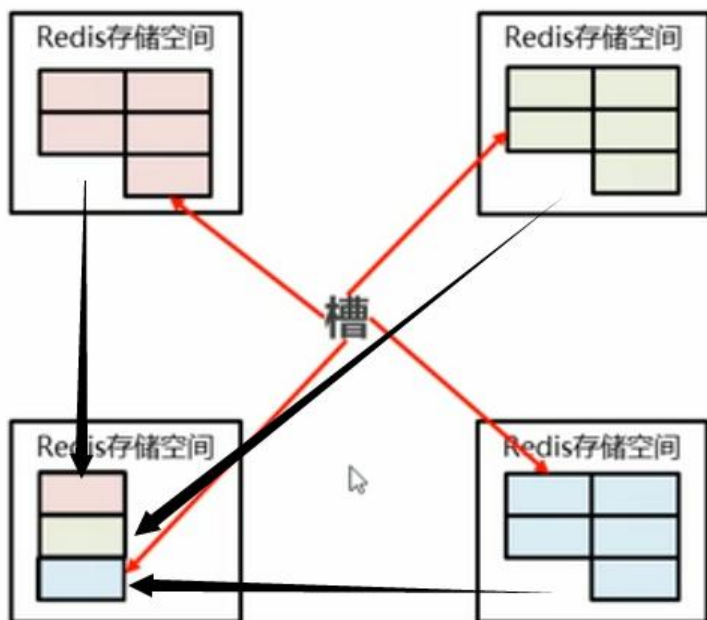
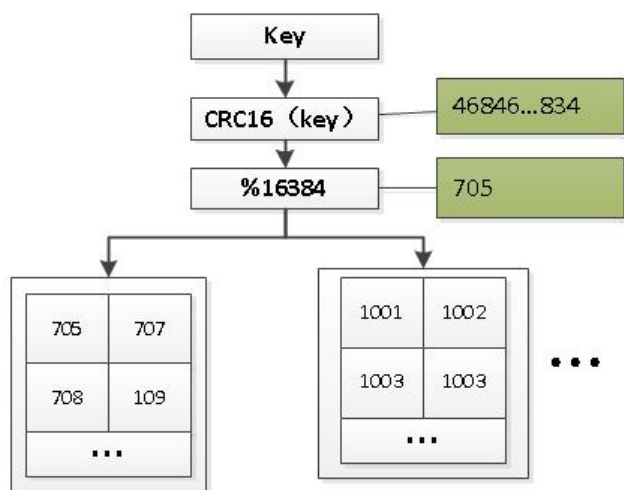
3.redis集群架构设计

数据存储设计:

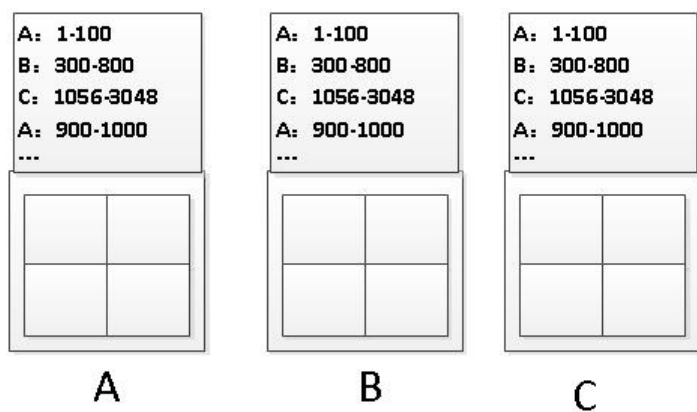
key经过CRC16方法

对16384取模

获取的值就是保存的位置



所谓的增节点、宕机其实就是改变槽所存储的位置不同



图解:

4.redis集群的搭建


配置文件

常用命令

```
[wsl@localhost redis]$ ps -ef |grep redis
root      32047   2849   0 21:47 pts/0    00:00:02 redis-server *:6379 [cluster]
root      32058   31868  0 21:47 pts/1    00:00:02 redis-server *:6380 [cluster]
root      32071   31839  0 21:48 pts/2    00:00:02 redis-server *:6381 [cluster]
root      32088   31810  0 21:48 pts/3    00:00:02 redis-server *:6382 [cluster]
root      32100   31781  0 21:49 pts/4    00:00:02 redis-server *:6383 [cluster]
root      32104   31744  0 21:49 pts/5    00:00:02 redis-server *:6384 [cluster]
wsl       32490   32010  0 22:17 pts/9    00:00:00 grep --color=auto redis
[wsl@localhost redis]$
```

启动过程中提示了不能够使用redis-trib进行启动提示使用redis-cli。因为我用了redis5版本:

```
[wsl@localhost src]$ redis-cli --cluster create 127.0.0.1:6379 127.0.0.1:6380 127.0.0.1:6381 127.0.0.1:6382 127.0.0.1:6383 127.0.0.1:6384 --cluster-replicas 1
>>> Performing hash slots allocation on 6 nodes...
Master[0] -> Slots 0 - 5460
Master[1] -> Slots 5461 - 10922
Master[2] -> Slots 10923 - 16383
Adding replica 127.0.0.1:6382 to 127.0.0.1:6379
Adding replica 127.0.0.1:6383 to 127.0.0.1:6380
Adding replica 127.0.0.1:6384 to 127.0.0.1:6381
>>> Trying to optimize slaves allocation for anti-affinity
[WARNING] Some slaves are in the same host as their master
M: 8a7531ecd8608b14a7262c47097bb30fbaa5cbd3 127.0.0.1:6379
  slots:[0-5460] (5461 slots) master
M: 61639514494f733ef22c71d35b4b1b7cc2f5dcd8 127.0.0.1:6380
  slots:[5461-10922] (5462 slots) master
M: 4ab47a6e61064a05c66787995fbf3e9f7370a60f 127.0.0.1:6381
  slots:[10923-16383] (5461 slots) master
S: b092daa3a44faec9b159e6fe46691982f5ba0dce 127.0.0.1:6382
  replicates 61639514494f733ef22c71d35b4b1b7cc2f5dcd8
S: 4f77b9f3106a0d7b465c0582e79f3824be10a1bb 127.0.0.1:6383
  replicates 4ab47a6e61064a05c66787995fbf3e9f7370a60f
S: f9c69bd18c859cb311f8a005c9fcb0483fadclbd 127.0.0.1:6384
  replicates 8a7531ecd8608b14a7262c47097bb30fbaa5cbd3
Can I set the above configuration? (type 'yes' to accept): yes
>>> Nodes configuration updated
>>> Assign a different config epoch to each node
>>> Sending CLUSTER MEET messages to join the cluster
Waiting for the cluster to join
.....
>>> Performing Cluster Check (using node 127.0.0.1:6379)
M: 8a7531ecd8608b14a7262c47097bb30fbaa5cbd3 127.0.0.1:6379
  slots:[0-5460] (5461 slots) master
```



```
Redis 5.0.3 (00000000/0) 64 bit
Running in cluster mode
Port: 6379
PID: 32047

http://redis.io
```

```
32047:M 21 Jun 2020 21:47:29.848 # WARNING: The TCP backlog setting of 511 cannot be enforced because /proc/sys/net/core/somaxconn is set to the low value of 128. To increase the value, run 'sysctl -w net.core.somaxconn=511' or change it to the desired value in /etc/sysctl.conf and then reboot or run the command 'sysctl vm.overcommit_memory=1' for this to take effect.
32047:M 21 Jun 2020 21:47:29.848 # WARNING overcommit_memory is set to 0! Background save may fail under low memory condition. To fix this issue add 'vm.overcommit_memory=1' in /etc/sysctl.conf and then reboot or run the command 'sysctl vm.overcommit_memory=1' for this to take effect.
32047:M 21 Jun 2020 21:47:29.848 # WARNING you have Transparent Huge Pages (THP) support enabled in your kernel. This will create latency and memory usage issues on your server. To disable THP support, run the command 'echo never > /sys/kernel/mm/transparent_hugepage/enabled' as root, and add it to your /etc/rc.local in order to retain the setting after THP is disabled.
32047:M 21 Jun 2020 21:47:29.848 * Ready to accept connections
32047:M 21 Jun 2020 22:17:05.994 # configEpoch set to 1 via CLUSTER SET-CONFIG-EPOCH
32047:M 21 Jun 2020 22:17:06.032 # IP address for this node updated to 127.0.0.1
32047:M 21 Jun 2020 22:17:11.000 # Cluster state changed: ok
32047:M 21 Jun 2020 22:17:12.912 * Replica 127.0.0.1:6384 asks for synchronization
32047:M 21 Jun 2020 22:17:12.912 * Partial resynchronization not accepted: Replication ID mismatch (Replica asked for 'df649639b1c21b01501ea8376a4580eea45bad806173b7387f7d3fbc3581c050d7' and '0000000000000000000000000000000000000000')
32047:M 21 Jun 2020 22:17:12.912 * Starting BGSAVE for SYNC with target: disk
32047:M 21 Jun 2020 22:17:12.913 * Background saving started by pid 32487
32487:C 21 Jun 2020 22:17:12.914 * DB saved on disk
32487:C 21 Jun 2020 22:17:12.915 * RDB: 0 MB of memory used by copy-on-write
32047:M 21 Jun 2020 22:17:12.935 * Background saving terminated with success
32047:M 21 Jun 2020 22:17:12.936 * Synchronization with replica 127.0.0.1:6384 succeeded
```