

问题： 一块磁盘不能读写

数据目录状态 (CDH-PRODUCT-V1 , HDFS , DataNode , cdh81-33)

测试该 DataNode 是否出现卷故障。

不良: 该 DataNode 当前有 1 个卷故障。 临界阈值: 任意。

操作

系统错误日志:

```
Oct 31 23:20:02 cdh81-33 systemd: Starting Session 540562 of user root.
Oct 31 23:20:02 cdh81-33 systemd-logind: Removed session 540562.
Oct 31 23:20:07 cdh81-33 kernel: XFS (sdc): xfs_log_force: error -5 returned.
Oct 31 23:20:10 cdh81-33 systemd-logind: New session 540563 of user root.
Oct 31 23:20:10 cdh81-33 systemd: Started Session 540563 of user root.
Oct 31 23:20:10 cdh81-33 systemd: Starting Session 540563 of user root.
Oct 31 23:20:10 cdh81-33 systemd-logind: Removed session 540563.
Oct 31 23:20:27 cdh81-33 kernel: XFS (sdc): xfs_log_force: error -5 returned.
Oct 31 23:20:28 cdh81-33 systemd-logind: New session 540564 of user root.
Oct 31 23:20:28 cdh81-33 systemd: Started Session 540564 of user root.
Oct 31 23:20:28 cdh81-33 systemd: Starting Session 540564 of user root.
Oct 31 23:20:28 cdh81-33 systemd-logind: Removed session 540564.
Oct 31 23:20:31 cdh81-33 systemd-logind: New session 540565 of user root.
Oct 31 23:20:31 cdh81-33 systemd: Started Session 540565 of user root.
Oct 31 23:20:31 cdh81-33 systemd: Starting Session 540565 of user root.
Oct 31 23:20:31 cdh81-33 systemd-logind: Removed session 540565.
Oct 31 23:20:31 cdh81-33 kernel: XFS (sdc): Unmounting Filesystem
Oct 31 23:20:31 cdh81-33 kernel: XFS (sdc): xfs_log_force: error -5 returned.
Oct 31 23:20:32 cdh81-33 kernel: XFS (sdc): xfs_log_force: error -5 returned.
Oct 31 23:20:41 cdh81-33 kernel: XFS (sdc): Mounting V4 Filesystem
Oct 31 23:20:49 cdh81-33 kernel: sd 0:0:4:0: [sdc] tag#0 FAILED Result: hostbyte=DID_OK driverbyte=DRIVER_SENSE
```

解决方法:

重启机器

```
umount /dfs/data3
```

```
mke2fs -L data /dev/sdc
```

```
mks.xfs -f /dev/sdc
```

```
mount -t xfs /dev/sdc /dfs/data3
```

下面是处理步骤:

df -h 查看:

```
[root@cdh81-33 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        252G   0  252G   0% /dev
tmpfs           252G  592K  252G   1% /dev/shm
tmpfs           252G  4.0G  248G   2% /run
tmpfs           252G   0  252G   0% /sys/fs/cgroup
/dev/mapper/ol-root 554G  21G  533G   4% /
/dev/sdg1       1014M 176M  839M  18% /boot
tmpfs           51G   0   51G   0% /run/user/0
/dev/sda        1.1T 902G  217G  81% /dfs/data1
/dev/sdb        1.1T 903G  215G  81% /dfs/data2
/dev/sdc        1.1T 908G  210G  82% /dfs/data3
/dev/sdd        1.1T 904G  215G  81% /dfs/data4
/dev/sde        1.1T 928G  190G  84% /dfs/data5
/dev/sdf        1.1T 913G  205G  82% /dfs/data6
cm_processes    252G  115M  252G   1% /run/cloudera-scm-agent/process
[root@cdh81-33 ~]#
```

但是data3不能写入数据:

```
drwxr-xr-x. 4 root root 50 Dec 6 2017 yarn
[root@cdh81-33 data1]# cp aa.txt /dfs/data2
[root@cdh81-33 data1]# cp aa.txt /dfs/data3
cp: failed to access '/dfs/data3': Input/output error
[root@cdh81-33 data1]# cp aa.txt /dfs/data4
[root@cdh81-33 data1]# cp aa.txt /dfs/data5
[root@cdh81-33 data1]# cp aa.txt /dfs/data6
[root@cdh81-33 data1]#
```

不能umount, 而且挂载目录data3 权限用户都是“?” 号

```
[root@cdh81-33 data3]# umount /dfs/data3
-bash: umount: command not found
[root@cdh81-33 data3]# mount /dfs/data3
mount: /dfs/data3: target is busy.
(In some cases useful info about processes that use
the device is found by lsof(8) or fuser(1))
[root@cdh81-33 data3]# cd ..
```

```
[root@cdh81-33 dfs]# umount -f /dfs/data3
mount: /dfs/data3: target is busy.
(In some cases useful info about processes that use
the device is found by lsof(8) or fuser(1))
[root@cdh81-33 dfs]# mount /dev/sdc /dfs/data3
mount: /dev/sdc is already mounted or /dfs/data3 busy
/dev/sdc is already mounted on /dfs/data3
[root@cdh81-33 dfs]# ll
ls: cannot access data3: Input/output error
total 0
drwxr-xr-x. 5 root root 57 Nov 1 10:54 data1
drwxr-xr-x. 5 root root 57 Nov 1 10:54 data2
d????????? ? ? ? ? ? data3
drwxr-xr-x. 5 root root 57 Nov 1 10:54 data4
drwxr-xr-x. 5 root root 57 Nov 1 10:54 data5
drwxr-xr-x. 5 root root 57 Nov 1 10:54 data6
```

查看占用的进程：

```
[root@cdh81-33 dfs]# fuser -cu /dfs/data3
Cannot stat /dfs/data3: Input/output error
Cannot stat file /proc/1447878/fd/3: Input/output error
Cannot stat file /proc/1447907/fd/341: Input/output error
Cannot stat file /proc/1448008/fd/1: Input/output error
Cannot stat file /proc/1448008/fd/2: Input/output error
Cannot stat file /proc/1448008/fd/349: Input/output error
Cannot stat file /proc/1448010/fd/1: Input/output error
Cannot stat file /proc/1448010/fd/2: Input/output error
Cannot stat file /proc/1448010/fd/334: Input/output error
Cannot stat file /proc/1593075/fd/3: Input/output error
Cannot stat file /proc/1593078/fd/4: Input/output error
Cannot stat file /proc/1593081/fd/3: Input/output error
Cannot stat file /proc/1593101/fd/327: Input/output error
Cannot stat file /proc/1593177/fd/327: Input/output error
Cannot stat file /proc/1593183/fd/327: Input/output error
Cannot stat file /proc/1593186/fd/327: Input/output error
Cannot stat file /proc/1593192/fd/327: Input/output error
[root@cdh81-33 dfs]#
[root@cdh81-33 dfs]#
[root@cdh81-33 dfs]#
[root@cdh81-33 dfs]# fuser -ck /dfs/data3
Cannot stat /dfs/data3: Input/output error
Cannot stat file /proc/1447878/fd/3: Input/output error
Cannot stat file /proc/1447907/fd/341: Input/output error
Cannot stat file /proc/1448008/fd/1: Input/output error
Cannot stat file /proc/1448008/fd/2: Input/output error
Cannot stat file /proc/1448008/fd/349: Input/output error
Cannot stat file /proc/1448010/fd/1: Input/output error
Cannot stat file /proc/1448010/fd/2: Input/output error
Cannot stat file /proc/1448010/fd/334: Input/output error
Cannot stat file /proc/1593075/fd/3: Input/output error
Cannot stat file /proc/1593078/fd/4: Input/output error
Cannot stat file /proc/1593081/fd/3: Input/output error
Cannot stat file /proc/1593101/fd/327: Input/output error
Cannot stat file /proc/1593177/fd/327: Input/output error
Cannot stat file /proc/1593183/fd/327: Input/output error
Cannot stat file /proc/1593186/fd/327: Input/output error
Cannot stat file /proc/1593192/fd/327: Input/output error
[root@cdh81-33 dfs]#
[root@cdh81-33 dfs]#
```

Kill 1447878 . . .

手动 kill 进程之后 可以umount了

```
[root@cdh81-33 dfs]# umount /dfs/data3
[root@cdh81-33 dfs]# mount /dev/sdc /dfs/data3
mount: mount /dev/sdc on /dfs/data3 failed: Remote I/O error
[root@cdh81-33 dfs]# ll
total 0
```

但是 mount的时候报错

```
[root@cdh81-33 ~]# mount -t xfs /dev/sdc /dfs/data3
mount: mount /dev/sdc on /dfs/data3 failed: Remote I/O error
```

```
[root@cdh81-33 ~]#
```

重新格式化也不行：

```
[root@cdh81-33 ~]# mkfs.xfs /dev/sdc
```

```
mkfs.xfs: /dev/sdc appears to contain an existing filesystem (xfs).
```

```
mkfs.xfs: Use the -f option to force overwrite.
```

```
[root@cdh81-33 ~]# mkfs.xfs -f /dev/sdc
```

```
meta-data=/dev/sdc          isize=256      agcount=4,
agsize=73257062 blks
                =                      sectsz=512
                attr=2, projid32bit=1
                =                      crc=0
                finobt=0, sparse=0
data                =                  bsize=4096
                blocks=293028246, imaxpct=5
                =                      sunit=0
                swidth=0 blks
naming              =version 2          bsize=4096   ascii-ci=0
ftype=1
log                 =internal log      bsize=4096   blocks=143080,
version=2
                =                      sectsz=512
                sunit=0 blks, lazy-count=1
realtime =none      extsz=4096   blocks=0,
rtextents=0
mkfs.xfs: pwrite64 failed: Input/output error
[root@cdh81-33 ~]#
```

检查磁盘，没有坏道

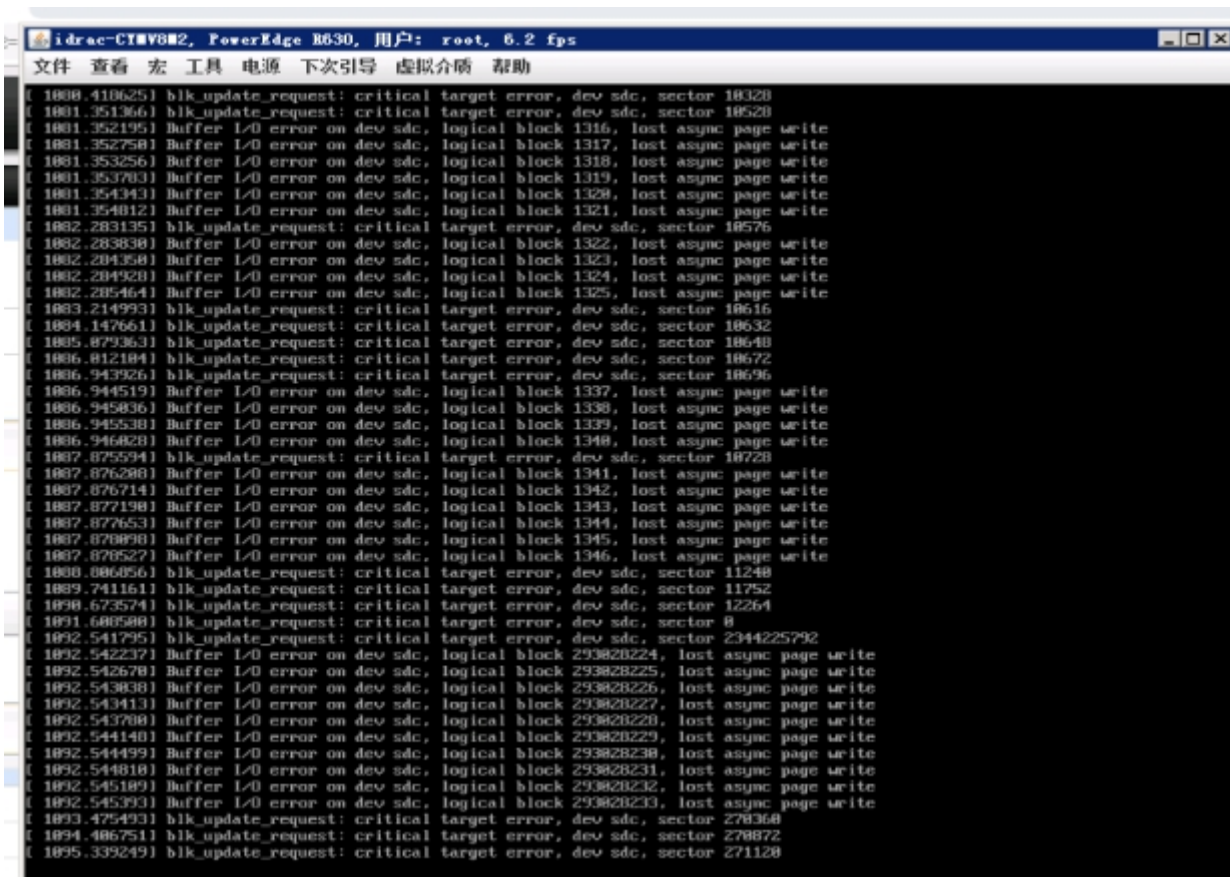
```
badblocks -s -v -o /opt/badblocks.log /dev/sdc
```

```
[root@cdh81-33 ~]# badblocks -s -v -o /opt/badblocks.log /dev/sdc
Checking blocks 0 to 1172112983
Checking for bad blocks (read-only test): done
Pass completed, 0 bad blocks found. (0/0/0 errors)
[root@cdh81-33 ~]#
```

重启机器

格式化磁盘

```
mke2fs -L data /dev/sdc
```



The screenshot shows a terminal window titled "idrac-CY8V8M2, PowerEdge R630, 用户: root, 6.2 fps". The terminal output displays a list of disk errors on /dev/sdc, including "blk_update_request: critical target error" and "Buffer I/O error on dev sdc, logical block". The errors are listed with their corresponding logical block numbers and sector numbers.

```
1080.410625] blk_update_request: critical target error, dev sdc, sector 10320
1081.351366] blk_update_request: critical target error, dev sdc, sector 10520
1081.352195] Buffer I/O error on dev sdc, logical block 1316, lost async page write
1081.352750] Buffer I/O error on dev sdc, logical block 1317, lost async page write
1081.353256] Buffer I/O error on dev sdc, logical block 1318, lost async page write
1081.353703] Buffer I/O error on dev sdc, logical block 1319, lost async page write
1081.354343] Buffer I/O error on dev sdc, logical block 1320, lost async page write
1081.354812] Buffer I/O error on dev sdc, logical block 1321, lost async page write
1082.283135] blk_update_request: critical target error, dev sdc, sector 10576
1082.283830] Buffer I/O error on dev sdc, logical block 1322, lost async page write
1082.284350] Buffer I/O error on dev sdc, logical block 1323, lost async page write
1082.284920] Buffer I/O error on dev sdc, logical block 1324, lost async page write
1082.285464] Buffer I/O error on dev sdc, logical block 1325, lost async page write
1083.214993] blk_update_request: critical target error, dev sdc, sector 10616
1084.147661] blk_update_request: critical target error, dev sdc, sector 10632
1085.079363] blk_update_request: critical target error, dev sdc, sector 10640
1086.012104] blk_update_request: critical target error, dev sdc, sector 10672
1086.943926] blk_update_request: critical target error, dev sdc, sector 10696
1086.944519] Buffer I/O error on dev sdc, logical block 1337, lost async page write
1086.945036] Buffer I/O error on dev sdc, logical block 1338, lost async page write
1086.945530] Buffer I/O error on dev sdc, logical block 1339, lost async page write
1086.946020] Buffer I/O error on dev sdc, logical block 1340, lost async page write
1087.075594] blk_update_request: critical target error, dev sdc, sector 10720
1087.076200] Buffer I/O error on dev sdc, logical block 1341, lost async page write
1087.076714] Buffer I/O error on dev sdc, logical block 1342, lost async page write
1087.077190] Buffer I/O error on dev sdc, logical block 1343, lost async page write
1087.077653] Buffer I/O error on dev sdc, logical block 1344, lost async page write
1087.078090] Buffer I/O error on dev sdc, logical block 1345, lost async page write
1087.078527] Buffer I/O error on dev sdc, logical block 1346, lost async page write
1088.086056] blk_update_request: critical target error, dev sdc, sector 11248
1089.741161] blk_update_request: critical target error, dev sdc, sector 11752
1090.673574] blk_update_request: critical target error, dev sdc, sector 12264
1091.600500] blk_update_request: critical target error, dev sdc, sector 0
1092.541795] blk_update_request: critical target error, dev sdc, sector 2944225792
1092.542237] Buffer I/O error on dev sdc, logical block 293820224, lost async page write
1092.542670] Buffer I/O error on dev sdc, logical block 293820225, lost async page write
1092.543030] Buffer I/O error on dev sdc, logical block 293820226, lost async page write
1092.543413] Buffer I/O error on dev sdc, logical block 293820227, lost async page write
1092.543700] Buffer I/O error on dev sdc, logical block 293820228, lost async page write
1092.544140] Buffer I/O error on dev sdc, logical block 293820229, lost async page write
1092.544499] Buffer I/O error on dev sdc, logical block 293820230, lost async page write
1092.544810] Buffer I/O error on dev sdc, logical block 293820231, lost async page write
1092.545109] Buffer I/O error on dev sdc, logical block 293820232, lost async page write
1092.545393] Buffer I/O error on dev sdc, logical block 293820233, lost async page write
1093.475493] blk_update_request: critical target error, dev sdc, sector 278360
1094.406751] blk_update_request: critical target error, dev sdc, sector 278872
1095.339249] blk_update_request: critical target error, dev sdc, sector 271120
```

(have trouble writing out superblocks)

```
mks. xfs -f /dev/sdc
```

```
mount -t xfs /dev/sdc /dfs/data3
```

最终问题解决了。

```
umount /dfs/data3 如果一次不行 重试多次。
```