



银河麒麟桌面操作系统
文件系统只读问题分析报告



麒麟软件有限公司
2021 年 02 月

目 录

1 问题概述.....	3
2 问题定位.....	3
2.1 内核 ext4 文件系统 bug 排查.....	3
2.2 母盘问题排查.....	3
2.3 排查总结.....	8
3 解决方案.....	8



文件系统只读问题分析报告

1 问题概述

通过母盘 dd 拷贝过来刚安装好的机器,经过大约 10 分钟左右的运行测试后出现文件系统报错,变只读文件系统。

2 问题定位

客户现场系统问题分析是文件系统出现了损坏,出现该问题分析有可能是如下两个方面导致的:

- 1、内核有 ext4 文件系统 BUG 导致
- 2、结合安装环境,可能是安装母盘出了问题

2.1 内核 ext4 文件系统 bug 排查

经对比测试内核相关补丁,排除内核层面的 ext4 文件系统问题。

2.2 母盘问题排查

根据现场反馈,用一台刚安装好的机器(系统是直接从母盘 dd 拷贝过来),开始测试,大约 10 分钟后出现文件系统报错,变只读文件系统。

具体排查过程:

- 1、从母盘(指新机器中通过拷贝机拷贝的系统盘)先 dd 一个新的系统盘;
- 2、使用 fsck.ext4 去检查新系统盘的 sdb2(根分区)与 sdb6(数据分区),可以看到这两个分区的 ext4 文件系统的元数据已经出现了不一致的问题,也就是说 ext4 文件系统已经有损坏了(测试母盘也是同样的现象),如下表 1 检测过程(注意以下检测过程提示**加粗字体**与**标红字体**)

母盘检测:

表 1 通过 dd 安装的系统盘及数据盘监测结果

<pre># fsck.ext4 -f /dev/sdb2 e2fsck 1.45.3 (14-Jul-2019) Pass 1: Checking inodes, blocks, and sizes Pass 2: Checking directory structure Pass 3: Checking directory connectivity Pass 4: Checking reference counts Pass 5: Checking group summary information</pre>

```
Block bitmap differences:      +24641536  +(24641546--24641547)
+(24641557--24642069) +(24646678--24647189)
Fix<y>? no
Free blocks count wrong for group #752 (27114, counted=28142).
Fix<y>? no
Free blocks count wrong (22406025, counted=22407053).
Fix<y>? no
Padding at end of block bitmap is not set. Fix<y>? no

/dev/sdb2: ***** WARNING: Filesystem still has errors *****

/dev/sdb2: 245602/6250496 files (0.1% non-contiguous), 2593911/24999936
blocks

# fsck.ext4 -f /dev/sdb6
e2fsck 1.45.3 (14-Jul-2019)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
Block bitmap differences:      +14155776 +14155792 +(14155808--14156320)
+14680064  +14680080  +(14680096--14680608)  +15204352  +15204368
+(15204384--15204896) +15728640 +15728656 +(15728672--15729184) +16252928
+16252944 +(16252960--16253472) +16777216 +16777232 +(16777248--16777760)
+17301504  +17301520  +(17301536--17302048)  +17825792  +17825808
+(17825824--17826336) +18350080 +18350096 +(18350112--18350624) +18874368
+18874384 +(18874400--18874912) +19398656 +19398672 +(19398688--19399200)
+19922944  +19922960  +(19922976--19923488)  +20447232  +20447248
+(20447264--20447776) +20971520 +20971536 +(20971552--20972064) +21495808
+21495824 +(21495840--21496352) +22020096 +22020112 +(22020128--22020640)
+22544384  +22544400  +(22544416--22544928)  +23068672  +23068688
+(23068704--23069216) +23592960 +23592976 +(23592992--23593504) +24117248
```

+24117264 +(24117280--24117792) +24641536 +24641552 +(24641568--24642080)
+25165824 +25165840 +(25165856--25166368) +25690112 +25690128
+(25690144--25690656) +26214400 +26214416 +(26214432--26214944) +26738688
+(26738703--26738704) +(26738719--26739232) +(26746415--26746927)

Fix<y>? no

Free blocks count wrong for group #432 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #448 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #464 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #480 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #496 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #512 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #528 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #544 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #560 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #576 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #592 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #608 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #624 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #640 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #656 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #672 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #688 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #704 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #720 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #736 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #752 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #768 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #784 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #800 (24528, counted=25043).

Fix<y>? no

Free blocks count wrong for group #816 (24528, counted=25558).

Fix<y>? no

Free blocks count wrong (26786347, counted=26799737).

Fix<y>? no

数据盘: *** WARNING: Filesystem still has errors *******

数据盘: 11/6829056 files (0.0% non-contiguous), 476629/27262976 blocks

3、使用 `fsck.ext4 -f -y /dev/sdb2` 与 `fsck.ext4 -f -y /dev/sdb6` 对新系统盘进行修复，同样的测试方法，测试多次再也不会出现文件系统损坏了。

下表 2 中是系统运行中出现文件系统损坏后用 `fsck` 检查结果，对比上图(母盘 `fsck` 检查结果)可以看出运行过程中的文件系统损坏进一步损坏了(对比 `sdb2`，元数据损坏的地方更多了)。

表 2 正常文件系统 `fsck` 检查结果

```
系统运行中出现文件系统损坏后用 fsck 检查结果
# fsck.ext4 /dev/sdb2
e2fsck 1.45.3 (14-Jul-2019)
/dev/sdb2: recovering journal
Superblock last mount time is in the future.
    (by less than a day, probably due to the hardware clock being incorrectly
set)
/dev/sdb2 contains a file system with errors, check forced.
Pass 1: Checking inodes, blocks, and sizes
Inodes that were part of a corrupted orphan linked list found.  Fix<y>? no
Inode 671237 was part of the orphaned inode list.  IGNORED.
Inode 1452588 was part of the orphaned inode list.  IGNORED.
Inode 1453095 was part of the orphaned inode list.  IGNORED.
Inode 1453100 was part of the orphaned inode list.  IGNORED.
Inode 1453114 was part of the orphaned inode list.  IGNORED.
Inode 1453115 was part of the orphaned inode list.  IGNORED.
Inode 1453116 was part of the orphaned inode list.  IGNORED.
Inode 1453117 was part of the orphaned inode list.  IGNORED.
Inode 1453118 was part of the orphaned inode list.  IGNORED.
Inode 1984708 was part of the orphaned inode list.  IGNORED.
Inode 1984709 was part of the orphaned inode list.  IGNORED.
Inode 3932162 was part of the orphaned inode list.  IGNORED.
Deleted inode 3932163 has zero dtime.  Fix<y>? no
Inode 3932169 was part of the orphaned inode list.  IGNORED.
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
Block bitmap differences:  -(1111137--1111138)  -(1111563--1111571)
-(1121284--1121285)        -(2933544--2933613)        -(7404037--7404038)
-(7404043--7404044)        -(7404054--7404055)        -15736865        -15736877
-(21028902--21028903) -(21028914--21028915) -(21028953--21028954) +24641536
+(24641546--24641547) +(24641557--24642069) +(24646678--24647189)
```

```
Fix<y>? no
Free blocks count wrong for group #752 (27114, counted=28142).
Fix<y>? no
Free blocks count wrong (22406025, counted=22428080).
Fix<y>? no
Inode bitmap differences:      -671237  -1452588  -1453095  -1453100
-(1453114--1453118) -(1984708--1984709) -(3932162--3932163) -3932169
Fix<y>? no
Directories count wrong for group #480 (10, counted=8).
Fix<y>? no
Free inodes count wrong (6004894, counted=6004854).
Fix<y>? no
Padding at end of block bitmap is not set. Fix<y>? no

/dev/sdb2: ***** FILE SYSTEM WAS MODIFIED *****

/dev/sdb2: ***** WARNING: Filesystem still has errors *****

/dev/sdb2: 245602/6250496 files (0.1% non-contiguous), 2593911/24999936
blocks
```

2.3 排查总结

系统出现文件系统损坏问题的根本原因是母盘有问题（母盘的 ext4 文件系统的元数据损坏了），系统在运行过程中访问到了损坏部分或相关连部分导致损坏进一步扩大，并最终被系统检测到，变成只读文件系统，重启系统时因文件系统损坏出现挂不上盘，从而进入 initramfs 黑屏。（因为目前 0 号母盘找不到了，不能排除是拷贝机在拷贝过程中出了错）

3 解决方案

根据排查的结论，解决此类问题，可使用以下两种方式处理：

1. 用 ISO 重新安装系统
2. 用 ISO 重新制作一个母盘，并确保正确（如需要用 fsck 检查等）