

## bluestore - Bug #50965

### In poweroff conditions BlueFS can create corrupted files

05/25/2021 07:12 AM - Adam Kupczyk

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Normal	
<b>Assignee:</b> Adam Kupczyk	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b>	<b>Affected Versions:</b>
<b>Tags:</b>	<b>ceph-qa-suite:</b>
<b>Backport:</b> pacific, octopus, nautilus	<b>Pull request ID:</b> 41501
<b>Regression:</b> No	<b>Crash signature (v1):</b>
<b>Severity:</b> 3 - minor	<b>Crash signature (v2):</b>
<b>Reviewed:</b>	

**Description**

It is possible to create condition in which a BlueFS contains file that is corrupted. It can happen when BlueFS replay log is on device A and we just wrote to device B.

Scenario:

- 1) write to file h1 on SLOW device
- 2) flush h1 (and trigger h1 mark to be added to bluefs replay log, but no fdatasync yet)
- 3) write to file h2 on DB
- 4) fsync h2 (forces replay log to be written, after fdatasync to DB)
- 5) poweroff

In result we have file h1 that is properly declared in replay log, but with uninitialized content.

**Related issues:**

Copied to bluestore - Backport #51128: octopus: In poweroff conditions BlueFS...	<b>Resolved</b>
Copied to bluestore - Backport #51129: nautilus: In poweroff conditions BlueF...	<b>Resolved</b>
Copied to bluestore - Backport #51130: pacific: In poweroff conditions BlueFS...	<b>Resolved</b>

### History

#### #1 - 05/27/2021 02:09 PM - Adam Kupczyk

- Pull request ID set to 41501

#### #2 - 06/05/2021 01:26 PM - Kefu Chai

- Status changed from New to Resolved

#### #3 - 06/08/2021 11:30 AM - Dan van der Ster

Why doesn't this need to be backported? The commit applies cleanly to pacific, at least.

#### #4 - 06/08/2021 11:32 AM - Igor Fedotov

- Status changed from Resolved to Pending Backport

- Backport set to pacific, octopus, nautilus

#### #5 - 06/08/2021 11:35 AM - Backport Bot

- Copied to Backport #51128: octopus: In poweroff conditions BlueFS can create corrupted files added

#### #6 - 06/08/2021 11:35 AM - Backport Bot

- Copied to Backport #51129: nautilus: In poweroff conditions BlueFS can create corrupted files added

**#7 - 06/08/2021 11:35 AM - Backport Bot**

- Copied to Backport #51130: *pacific: In poweroff conditions BlueFS can create corrupted files added*

**#8 - 09/20/2021 04:26 PM - Loïc Dachary**

- Status changed from *Pending Backport* to *Resolved*

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".