

rgw - Bug #40905

rgw: list bucket with delimiter wrongly skip some special keys

07/23/2019 01:49 PM - Tianshan Qu

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category:	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport: nautilus,mimic,luminous	Pull request ID: 29215
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	
Description	
list with delimiter will skip subfile with directory + after_delim_s, but the code wrongly add after_delim_s to next marker regardless it have directory.	
some example: we have objects like	
bucket/0/ bucket/1001 ... bucket/1999 bucket/1999# bucket/1999+ bucket/2000	
first time list get last key = bucket/1999, and the current code will add "/0xff" after 1999 to be next_marker, which cause the 1999# and 1999+ skipped.	
Related issues:	
Related to rgw - Bug #47487: rgw: ordered bucket listing code clean-up	Resolved
Copied to rgw - Backport #41484: mimic: rgw: list bucket with delimiter wrong...	Resolved
Copied to rgw - Backport #41485: nautilus: rgw: list bucket with delimiter wr...	Resolved
Copied to rgw - Backport #41486: luminous: rgw: list bucket with delimiter wr...	Rejected

History

#1 - 07/23/2019 01:52 PM - Tianshan Qu

<https://github.com/ceph/ceph/pull/29215>

#2 - 07/23/2019 02:04 PM - Matt Benjamin

- Status changed from New to Triaged

- Backport set to nautilus,mimic,luminous

need test case(s)

#3 - 07/23/2019 02:18 PM - Casey Bodley

- Pull request ID set to 29215

#4 - 07/25/2019 05:47 PM - Abhishek Lekshmanan

- Status changed from Triaged to In Progress

#5 - 08/08/2019 05:52 PM - Matt Benjamin

- Status changed from In Progress to 17

#6 - 08/15/2019 02:24 PM - Matt Benjamin

@tianshan,

I've been trying to reproduce on an unpatched branch, using RGW NFS. I do have <https://github.com/ceph/ceph/pull/29670>, which fixes a bug with NFS-generated markers in a pseudo-directory immediately below /bucket.

I did the following, which intends to match your test case--/folder/tianshan is a bucket:

```
1211 mount -tnfs lemon:/userx /folder -o sync
1212 cd /folder
1213 ls
1214 mkdir tianshan
1215 cd tianshan/
1216 ls
1217 touch f{1001,1002,1999,1999#,1999+,2000}
```

(unmount and restart of ganesh, then remount--to ensure that a listing is done)

```
[root@lemon /]# ls /folder/tianshan/
f1001 f1002 f1999 f1999# f1999+ f2000
```

I adjusted this slightly by creating files f1000..f1999 as well, so that f1999+ should be next_marker, but did not see it skipped.

I think we need a discrete test case.

Matt

#7 - 08/22/2019 05:53 PM - Matt Benjamin

- Status changed from 17 to Pending Backport

#8 - 08/26/2019 08:06 AM - Tianshan Qu

@Matt Benjamin

need an extra directory before the f1999+ to trigger the skip logical, such as f0/, otherwise the skip_after_delim will not have value.

#9 - 08/26/2019 02:50 PM - Nathan Cutler

- Copied to Backport #41484: mimic: rgw: list bucket with delimiter wrongly skip some special keys added

#10 - 08/26/2019 02:50 PM - Nathan Cutler

- Copied to Backport #41485: nautilus: rgw: list bucket with delimiter wrongly skip some special keys added

#11 - 08/26/2019 02:50 PM - Nathan Cutler

- Copied to Backport #41486: luminous: rgw: list bucket with delimiter wrongly skip some special keys added

#12 - 09/14/2020 09:05 PM - J. Eric Ivancich

- Status changed from Pending Backport to Resolved

#13 - 09/15/2020 08:11 PM - J. Eric Ivancich

- Related to Bug #47487: rgw: ordered bucket listing code clean-up added