

fs - Bug #45553

mds: rstats on snapshot are updated by changes to HEAD

05/14/2020 11:40 AM - Kotresh Hiremath Ravishankar

Status:	Duplicate	% Done:	0%
Priority:	Urgent		
Assignee:			
Category:	Correctness/Safety		
Target version:	v16.0.0		
Source:	Development	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Component(FS):	MDS
Regression:	No	Labels (FS):	
Severity:	2 - major	Pull request ID:	
Reviewed:		Crash signature:	

Description

The "ceph.dir.rbytes" xattr on .snap/<snap_name> directory is getting updated on already taken snapshots. Check the details below.

On a vstart environment created with "env MDS=3 ../src/vstart.sh -d -b -n --without-dashboard"

1. Create subvolume

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# bin/ceph fs subvolume create a sub_0
*** DEVELOPER MODE: setting PATH, PYTHONPATH and LD_LIBRARY_PATH ***
2020-05-14T15:39:13.132+0530 7f88d9b27700 -1 WARNING: all dangerous and experimental features are
enabled.
2020-05-14T15:39:13.156+0530 7f88d9b27700 -1 WARNING: all dangerous and experimental features are
enabled.
```

2. List created subvolume

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# bin/ceph fs subvolume ls a
* DEVELOPER MODE: setting PATH, PYTHONPATH and LD_LIBRARY_PATH *
2020-05-14T15:39:28.432+0530 7f03423b1700 -1 WARNING: all dangerous and experimental features are
enabled.
2020-05-14T15:39:28.452+0530 7f03423b1700 -1 WARNING: all dangerous and experimental features are
enabled.
[
  {
    "name": "sub_0"
  }
]
```

3. Write 1GB file to the subvolume

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# dd if=/dev/zero of=/mnt/volumes/_nogroup/sub
_0/af4abd14-41e6-4843-9945-96a0d0b48128/file1 bs=1M count=1
1+0 records in
1+0 records out
1048576 bytes (1.0 MB, 1.0 MiB) copied, 0.0127988 s, 81.9 MB/s
```

4. Create a snapshot

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# bin/ceph fs subvolume snapshot create a sub_
0 snap_sub_0
*** DEVELOPER MODE: setting PATH, PYTHONPATH and LD_LIBRARY_PATH ***
2020-05-14T15:41:52.948+0530 7f6f40135700 -1 WARNING: all dangerous and experimental features are
```

enabled.

```
2020-05-14T15:41:52.972+0530 7f6f40135700 -1 WARNING: all dangerous and experimental features are enabled.
```

5. List snapshot

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# bin/ceph fs subvolume snapshot ls a sub_0
*** DEVELOPER MODE: setting PATH, PYTHONPATH and LD_LIBRARY_PATH ***
2020-05-14T15:42:04.940+0530 7fe14cb06700 -1 WARNING: all dangerous and experimental features are enabled.
2020-05-14T15:42:04.964+0530 7fe14cb06700 -1 WARNING: all dangerous and experimental features are enabled.
[
  {
    "name": "snap_sub_0"
  }
]
```

6. Get rbytes on subvolume directory

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# getfattr -d -m . -e text /mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/
getfattr: Removing leading '/' from absolute path names
file: mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/
ceph.dir.entries="1"
ceph.dir.files="1"
ceph.dir.rbytes="1048576"
ceph.dir.rctime="1589455578.148494821"
ceph.dir.rentries="2"
ceph.dir.rfiles="1"
ceph.dir.rsubdirs="1"
ceph.dir.subdirs="0"
```

7. Get rbytes on snap directory

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# getfattr -d -m . -e text /mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/.snap/snap_sub_0/
getfattr: Removing leading '/' from absolute path names
file: mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/.snap/snap_sub_0/
ceph.dir.entries="1"
ceph.dir.files="1"
ceph.dir.rbytes="1048576"
ceph.dir.rctime="1589455578.148494821"
ceph.dir.rentries="2"
ceph.dir.rfiles="1"
ceph.dir.rsubdirs="1"
ceph.dir.subdirs="0"
ceph.snap.btime="1589455578.148494821"
```

8. Create another 1G file

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# dd if=/dev/zero of=/mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/file2 bs=1M count=1
1+0 records in
1+0 records out
1048576 bytes (1.0 MB, 1.0 MiB) copied, 0.00110732 s, 947 MB/s
```

9. Get rbytes on subvolume. It gets updated to reflect 2GB

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# getfattr -d -m . -e text /mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/
getfattr: Removing leading '/' from absolute path names
file: mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/
ceph.dir.entries="2"
ceph.dir.files="2"
ceph.dir.rbytes="2097152"
ceph.dir.rctime="1589455807.534189066"
```

```
ceph.dir.rentries="3"  
ceph.dir.rfiles="2"  
ceph.dir.rsubdirs="1"  
ceph.dir.subdirs="0"
```

10. Get rbytes on snap directory. It also is updated with 2GB.

```
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build#  
root@kotresh-T490s:~/sandbox/kotresh-ceph/ceph/build# getfattr -d -m . -e text /mnt/volumes/_nogro  
up/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/.snap/snap_sub_0/  
getfattr: Removing leading '/' from absolute path names  
# file: mnt/volumes/_nogroup/sub_0/af4abd14-41e6-4843-9945-96a0d0b48128/.snap/snap_sub_0/  
ceph.dir.entries="1"  
ceph.dir.files="1"  
ceph.dir.rbytes="2097152"  
ceph.dir.rctime="1589455807.534189066"  
ceph.dir.rentries="3"  
ceph.dir.rfiles="2"  
ceph.dir.rsubdirs="1"  
ceph.dir.subdirs="0"  
ceph.snap.btime="1589455578.148494821"
```

Related issues:

Duplicates fs - Feature #24725: mds: propagate rstats from the leaf dirs up t...

Fix Under Review

History

#1 - 05/14/2020 09:10 PM - Patrick Donnelly

- Subject changed from CephFS Snapshots: "ceph.dir.rbytes" on .snap/<snap_name> directory is getting updated on older snapshots to mds: rstats on snapshot are updated by changes to HEAD
- Priority changed from Normal to Urgent
- Target version set to v16.0.0
- Severity changed from 3 - minor to 2 - major
- Component(FS) MDS added

#2 - 05/18/2020 01:46 PM - Patrick Donnelly

- Status changed from New to Duplicate

#3 - 05/18/2020 01:46 PM - Patrick Donnelly

- Duplicates Feature #24725: mds: propagate rstats from the leaf dirs up to the specified diretory added