

rgw - Bug #21200

rgw: setxattrs call leads to different mtimes for bucket index and object

08/31/2017 11:34 AM - Abhishek Lekshmanan

Status:	Resolved	Start date:	08/31/2017
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	luminous	ceph-qa-suite:	
Regression:	No	Release:	
Severity:	3 - minor	Needs Doc:	No

Description

calling `RGWRados::set_attrs()` will lead to different mtimes for the object and that stored in the bucket index, currently this is consumed by swift: `setxattrs`, `s3`: object tagging (probably `acls` as well).

Needs <https://github.com/ceph/ceph/pull/17395> for verification (or a hand written application using `librados::stat2` calls)

Initial state

```
$ ./bin/rados -p default.rgw.buckets.index getomapval .dir.7d1d633f-6404-489d-9653-648e393014f5.4133.1 obj obj_initial.omapval
```

```
$ ./bin/ceph-dencoder type rgw_bucket_dir_entry import obj_initial.omapval decode dump_json
```

```
{
  "name": "obj",
  "instance": "",
  "ver": {
    "pool": 6,
    "epoch": 1
  },
  "locator": "",
  "exists": "true",
  "meta": {
    "category": 1,
    "size": 333,
    "mtime": "2017-08-31 11:25:16.024856Z",
    "etag": "4c693106c9ae67cc4d857fb7c11a89bf",
    "owner": "test",
    "owner_display_name": "Tester-Subuser",
    "content_type": "",
    "accounted_size": 333,
    "user_data": ""
  },
  "tag": "7d1d633f-6404-489d-9653-648e393014f5.4133.5",
  "flags": 0,
  "pending_map": [],
  "versioned_epoch": 0
}
```

```
$ ./bin/rados -p default.rgw.buckets.data stat2 7d1d633f-6404-489d-9653-648e393014f5.4133.1_obj
2017-08-31 13:28:37.235041 7f463313fe00 -1 WARNING: all dangerous and experimental features are enabled.
2017-08-31 13:28:37.235527 7f463313fe00 -1 WARNING: all dangerous and experimental features are enabled.
2017-08-31 13:28:37.237181 7f463313fe00 -1 WARNING: all dangerous and experimental features are enabled.
default.rgw.buckets.data/7d1d633f-6404-489d-9653-648e393014f5.4133.1_obj mtime 2017-08-31 13:25:16
```

```
.024856, size 333

$ swift -A http://localhost:8000/auth -U "test:tester" -K testing post cont1 obj -m foo:bar

$ ./bin/rados -p default.rgw.buckets.index getomapval .dir.7d1d633f-6404-489d-9653-648e393014f5.4133.1_obj_obj_final.omapval
Writing to obj_final.omapval
$ ./bin/ceph-dencoder type rgw_bucket_dir_entry import obj_final.omapval decode dump_json
{
  "name": "obj",
  "instance": "",
  "ver": {
    "pool": 6,
    "epoch": 2
  },
  "locator": "",
  "exists": "true",
  "meta": {
    "category": 1,
    "size": 333,
    "mtime": "2017-08-31 11:29:26.008730Z",
    "etag": "4c693106c9ae67cc4d857fb7c11a89bf\u0000",
    "owner": "test",
    "owner_display_name": "Tester-Subuser",
    "content_type": "\u0000",
    "accounted_size": 333,
    "user_data": ""
  },
  "tag": "_R9cfI4zUfPFpZsEEaWv6arRljs4Zjs6",
  "flags": 0,
  "pending_map": [],
  "versioned_epoch": 0
}
$ ./bin/rados -p default.rgw.buckets.data stat2 7d1d633f-6404-489d-9653-648e393014f5.4133.1_obj
default.rgw.buckets.data/7d1d633f-6404-489d-9653-648e393014f5.4133.1_obj mtime 2017-08-31 13:29:25.994271, size 333
```

After a swift setxattrs call, bi reports the mtime at 2017-08-31 13:29:25.994271 compared to obj's mtime at 2017-08-31 11:29:26.008730Z (ignore the hour difference UTC vs local time)

Related issues:

Copied to rgw - Backport #21444: luminous: rgw: setxattrs call leads to diffe...

Resolved

History

#1 - 08/31/2017 04:09 PM - Abhishek Lekshmanan

- Status changed from New to Need Review

master PR: <https://github.com/ceph/ceph/pull/17400>

#2 - 09/14/2017 06:16 PM - Abhishek Lekshmanan

- Status changed from Need Review to Pending Backport

- Backport set to luminous

#3 - 09/19/2017 11:37 AM - Nathan Cutler

- Copied to Backport #21444: luminous: rgw: setxattrs call leads to different mtimes for bucket index and object added

#4 - 10/27/2017 03:02 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved