

rgw - Bug #19870

rgw: bytes_send and bytes_rcv in the msg of usage show returning is 0 in master branch

05/05/2017 10:30 AM - Enming Zhang

Status:	Resolved	Start date:	05/05/2017
Priority:	High	Due date:	
Assignee:	Marcus Watts	% 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	Pull request ID:	
Severity:	3 - minor		

Description

```
⊠ build git:(master) ⊠ ./bin/radosgw-admin usage show {
```

```
"entries": [
  {
    "user": "tester",
    "buckets": [
      {
        "bucket": "",
        "time": "2017-05-05 06:00:00.000000Z",
        "epoch": 1493964000,
        "owner": "tester",
        "categories": [
          {
            "category": "list_buckets",
            "bytes_sent": 0,
            "bytes_received": 0,
            "ops": 2,
            "successful_ops": 2
          }
        ]
      },
      {
        "bucket": "222",
        "time": "2017-05-05 06:00:00.000000Z",
        "epoch": 1493964000,
        "owner": "tester",
        "categories": [
          {
            "category": "create_bucket",
            "bytes_sent": 0,
            "bytes_received": 0,
            "ops": 1,
            "successful_ops": 1
          }
        ]
      }
    ]
  },
  {
    "summary": [
      {
        "user": "tester",
        "categories": [
          {
            "category": "create_bucket",
            "bytes_sent": 0,
            "bytes_received": 0,
            "ops": 1,
            "successful_ops": 1
          }
        ],
        "category": "list_buckets",
        "bytes_sent": 0,
```

```
"bytes_received": 0,
"ops": 2,
"successful_ops": 2
},
],
"total": {
"bytes_sent": 0,
"bytes_received": 0,
"ops": 3,
"successful_ops": 3
}
}
]
```

Related issues:

Related to rgw - Bug #19194: after s3cmd put file to bucket, 'radosgw-admin u...	Duplicate	03/06/2017
Copied to rgw - Backport #20968: luminous: rgw: bytes_send and bytes_recv in ...	Resolved	

History

#1 - 05/05/2017 05:10 PM - Greg Farnum

- Project changed from Ceph to rgw

#2 - 05/25/2017 05:38 PM - Orit Wasserman

- Assignee set to Marcus Watts

#3 - 05/25/2017 05:45 PM - Casey Bodley

- Related to Bug #19194: after s3cmd put file to bucket, 'radosgw-admin usage show --categories=put_obj' bytes_received=0 added

#4 - 05/25/2017 05:46 PM - Orit Wasserman

duplicate of <http://tracker.ceph.com/issues/19194>

#5 - 06/06/2017 05:53 PM - Yehuda Sadeh

- Subject changed from bytes_send and bytes_recv in the msg of usage show returning is 0 in master branch to rgw: bytes_send and bytes_recv in the msg of usage show returning is 0 in master branch

#6 - 06/06/2017 05:57 PM - Yehuda Sadeh

while we're at it, we need to make sure that there's sufficient teuthology coverage for this issue. I think we do test usage, but maybe not look at the bytes_received / bytes_sent fields.

#7 - 06/07/2017 07:15 AM - Marcus Watts

I may have found the teuthology tests. in ceph, master branch, "qa/tasks/radosgw_admin.py" and "qa/tasks/radosgw_admin_rest.py". Is that the current thing? Are there any other places I should look? It looks like that just exercises radosgw-admin interfaces, so may not fully cover all the s3/swift possible accounting paths. There's no logic there to look at bytes_sent|bytes_received. Looks like it's mainly checking existence and "successful_ops".

jewel has relatively straightforward code for this. Does it calculate bytes_sent and bytes_received correctly?

Looking at master, it appears it has significantly more complex code using c++ classes, templates, and such. One of the classes involved is "BufferingFilter" which appears to have the property of shifting byte counts from send_body() to a later complete_request() call. So far my gdb experiments suggest that this later count is not wrapped in AccountingFilter and so winds up not being accumulated.

#8 - 06/08/2017 08:28 PM - Marcus Watts

cbodley has a fix for a similar problem with swift:

<https://github.com/ceph/ceph/pull/14042>

I believe it may be correct that "create_bucket" lists 0 bytes sent/received. The count appears to be intended to cover the body of the request/response only, not the headers. create_bucket may not have any data for either. I'll verify that this is the case with jewel/boto.

"list_buckets" should definitely have a nz sent_bytes count. Jewel certainly shows a nz count. I've been able to confirm with debug code that the count is being lost in process_request() when it calls complete_request.

In my experiments with s3cmd, I believe get/put object do update byte counts right.

I'm looking at what it takes to run teuthology's radmin_admin command standalone. I appear to have managed this once upon a time, so it should be just a matter of updating my process. Once I can do that, adding code to test out usage should be much more feasible.

#9 - 06/13/2017 07:58 AM - Marcus Watts

I've been able to run a modified version of radosgw_admin.py against jewel & master. Modifications included instrumenting boto calls to track sent & received bytes, moving the "usage show" logic to near the end (after all the boto) calls, and added a few more boto calls to exercise more interfaces.

Both jewel & master seem to have various oddities in counts.

jewel,

list_buckets: actually received 299 bytes, but ceph claims to have sent 320.

create_bucket: in 9 ops, 8 successes returned 0 bytes, one failed returning 218 bytes. Ceph claims to have sent 152 bytes in 8 ops.

get_acls: got 1026 bytes back, but ceph claims to have sent 1068.

list_bucket: got 484 bytes, but ceph claims to have sent 505.

put_acls: got 0 bytes back, but ceph claims to have sent 38.

master,

create_bucket: received 220 bytes in 1 failure of 9 ops total, but ceph claims to have sent 0 bytes in 8 ops.

get_acls: received 1026 bytes, but ceph claim to have sent 0 bytes.

list_bucket: received 484 bytes, but ceph claims to have sent 0 bytes.

list_buckets: received 299 bytes, but ceph claims to have sent 0 bytes.

My build of ceph master doesn't have casey's patch. I'll have to look at that to see if it fixes s3 too.

Ideally these tests should be doing swift too, and be capable of testing casey's patch more directly.

#10 - 06/17/2017 07:31 AM - Marcus Watts

AccountingFilter is missing "complete_request". So that's why master drops "bytes_sent" in most cases. Implementing that function in reveals another problem: BufferingFilter does some trickery of its own. The big problem is that in complete_request, it calls

send_content_length/complete_header - which return counts that should not be counted. The mechanism that's supposed to control when bytes are counted, "AccountingFilter::set_account()", is only supposed to be called after complete_header() is called, but BufferingFilter cleverly defers calling that. I have a not very elegant fix for this, which seems to be giving me mostly correct counts.

#11 - 06/18/2017 10:29 PM - Marcus Watts

github.com:mdw-at-linuxbox/ceph-1 policy is a copy of my current changes to master to do this. Lots of temp debug messages and deadend code get_pending_data_count() to remove--needs significant cleanup before making a PR. Also there's a dmclock build fix which at the very least has a different sha1 in master.

#12 - 06/22/2017 05:44 PM - Matt Benjamin

per Marcus 6/21, he has a candidate fix with some "ugly" aspects, is working on revised version

#13 - 07/20/2017 06:01 PM - Matt Benjamin

- Status changed from New to In Progress

#14 - 07/27/2017 05:40 PM - Casey Bodley

can Marcus propose a simple fix so other developers can review and suggest alternatives?

#15 - 08/04/2017 11:29 PM - Marcus Watts

I've posted a fix against master for the bytes usage fix.

<https://github.com/ceph/ceph/pull/16834>

I got rid of some unnecessary detours for fixing this bug (fixes based on theories of the bug that turned out to be wrong.)

This passes modified teuthology test code that I'll shortly post as a PR against teuthology.

#16 - 08/05/2017 01:23 AM - Marcus Watts

Turns out the actual teuthologic logic to test usage now lives in the ceph repo! So

<https://github.com/ceph/ceph/pull/16835>

This points to a changed radosgw_admin.py file that checks usage a lot more thoroughly.

I tested this logic out against a jewel without any special patches - it passes this test as well. So we could use this to catch future regressions in jewel if we want.

#17 - 08/09/2017 06:24 PM - Matt Benjamin

does this need backport (L?)

#18 - 08/09/2017 06:25 PM - Matt Benjamin

- Status changed from In Progress to Pending Backport

#19 - 08/09/2017 06:36 PM - Nathan Cutler

- Backport set to luminous

#20 - 08/09/2017 07:31 PM - Nathan Cutler

- Copied to Backport #20968: luminous: rgw: bytes_send and bytes_rcv in the msg of usage show returning is 0 in master branch added

#21 - 09/11/2017 07:50 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved