

rgw - Bug #11851

Civetweb RGW appears to report full size of object as downloaded when only partially downloaded

06/02/2015 07:26 PM - Matthew Wodrich

Status:	Resolved	Start date:	06/02/2015
Priority:	Urgent	Due date:	
Assignee:	Orit Wasserman	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	Community (user)	Reviewed:	
Tags:		Affected Versions:	
Backport:	hammer, firefly	ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Description			
Looking at the radosgw-admin usage show get_obj bytes_sent output for a user with 500MB to ~1.1GB objects being streamed or partially downloaded, the Civetweb RGW appears to be recording that the entire size of an object has been sent for every single request, even where it's clear from logs that the connection was aborted mid-download. This results in usage numbers orders of magnitude over true usage. This has been observed on Firefly.			
Related issues:			
Copied to rgw - Backport #12243: Civetweb RGW appears to report full size of ...		Resolved	06/02/2015
Copied to rgw - Backport #12244: Civetweb RGW appears to report full size of ...		Resolved	06/02/2015

Associated revisions

Revision daa679c3 - 06/25/2015 09:35 PM - Yehuda Sadeh

rgw: error out if frontend did not send all data

Fixes: #11851

The civetweb mg_write() doesn't return error when it can't flush all data to the user, it just sends the total number of bytes written. Modified the client io to return total number of bytes and return an error if didn't send anything.

Signed-off-by: Yehuda Sadeh <yehuda@redhat.com>

Revision ec705336 - 07/14/2015 09:27 AM - Yehuda Sadeh

rgw: error out if frontend did not send all data

Fixes: #11851

The civetweb mg_write() doesn't return error when it can't flush all data to the user, it just sends the total number of bytes written. Modified the client io to return total number of bytes and return an error if didn't send anything.

Signed-off-by: Yehuda Sadeh <yehuda@redhat.com>

(cherry picked from commit daa679c3dd3770a6d5421e2cc9a36924f4550439)

Revision 4e7de5b5 - 07/14/2015 10:09 AM - Yehuda Sadeh

rgw: error out if frontend did not send all data

Fixes: #11851

The civetweb mg_write() doesn't return error when it can't flush all data to the user, it just sends the total number of bytes written. Modified the client io to return total number of bytes and return an error if didn't send anything.

Signed-off-by: Yehuda Sadeh <yehuda@redhat.com>
(cherry picked from commit daa679c3dd3770a6d5421e2cc9a36924f4550439)

History

#1 - 06/04/2015 06:34 PM - Sage Weil

- Priority changed from Normal to Urgent
- Source changed from other to Community (user)
- Backport set to hammer, firefly

#2 - 06/16/2015 05:39 PM - Yehuda Sadeh

Can this be reproduced with a more concrete info? Looking at the code I don't think we report the full object size for partial upload. What might be happening is that civetweb is buffering a lot so we end up overreporting if nothing went on the wire, but this is just a guess.

#3 - 06/25/2015 09:37 PM - Yehuda Sadeh

- Status changed from New to Need Review
- Assignee set to Orit Wasserman

#4 - 06/30/2015 05:58 PM - Yehuda Sadeh

- Status changed from Need Review to Pending Backport
- Assignee changed from Orit Wasserman to Loic Dachary

#5 - 07/08/2015 12:42 PM - Loic Dachary

- Assignee changed from Loic Dachary to Orit Wasserman

#6 - 07/14/2015 09:29 AM - Abhishek Lekshmanan

Master PR: <https://github.com/ceph/ceph/pull/5081>

#7 - 07/14/2015 09:32 AM - Abhishek Lekshmanan

hammer backport: <https://github.com/ceph/ceph/pull/5231>

#8 - 08/11/2015 06:06 PM - Yehuda Sadeh

- Status changed from Pending Backport to Resolved