

Subject: Re: URGENT: New deadlock issue with the .53 build (RHGS 3.0.4 GA candidate) for random writes with IOZone.

From: Vivek Agarwal <vagarwal@redhat.com>

Date: 03/24/2015 11:57 AM

To: Ben Turner <bturner@redhat.com>, rejoy@redhat.com

CC: Nagaprasad Sathyanarayana <nsathyan@redhat.com>, Anoop Kumar Nair <annair@redhat.com>, Sayan Saha <ssaha@redhat.com>, Pranith Kumar Karampuri <pkarampu@redhat.com>, Shyam <srangana@redhat.com>, Krutika Dhananjay <kdhananj@redhat.com>, Vijay Bellur <vbellur@redhat.com>, Alok Srivastava <asrivast@redhat.com>, Satish Mohan <smohan@redhat.com>, Rajesh Joseph <rjoseph@redhat.com>

On 03/24/2015 08:58 AM, Ben Turner wrote:

----- Original Message -----

From: "Rejoy M Cyriac" <rcyriac@redhat.com>

To: "Ben Turner" <bturner@redhat.com>, "Nagaprasad Sathyanarayana" <nsathyan@redhat.com>, "Anoop Kumar Nair" <annair@redhat.com>

Cc: "Sayan Saha" <ssaha@redhat.com>, "Pranith Kumar Karampuri" <pkarampu@redhat.com>, "Shyam" <srangana@redhat.com>, "Krutika Dhananjay" <kdhananj@redhat.com>, "Vijay Bellur" <vbellur@redhat.com>, "Alok Srivastava" <asrivast@redhat.com>, rejoy@redhat.com, "Satish Mohan" <smohan@redhat.com>, "Vivek Agarwal" <vagarwal@redhat.com>

Sent: Monday, March 23, 2015 1:06:26 AM

Subject: Re: URGENT: New deadlock issue with the .53 build (RHGS 3.0.4 GA candidate) for random writes with IOZone.

@devel and qe

I think it would be helpful if we capture at least a summary of the root cause investigation done so far, at the BZ - https://bugzilla.redhat.com/show_bug.cgi?id=1203850 . There is nothing at the BZ beyond the initial description and comments put in by BenT.

As I have said before, I think the hangs on glusterfs, NFS and cifs are caused by the same issue manifesting a bit differently on each the 3 mount types. I put all of the data in the glusterfs BZ and I opened the NFS BZ to cover my bases. I put anything NFS specific in the NFS BZ, with all I have todo before release I couldn't duplicate all the troubleshooting. Now we are seeing the issue on cifs as well, should I open up a 3rd BZ for CIFS or do we want to work all 3 in the glusterfs BZ? I don't have any evidence the hang is the same, its just my gut.

+Rajesh

Clearly as the issue is not easy to reproduce, this is not a blocker for 3.04. It is fair to open anothe bz for CIFS and put the relevant details there.

Regards,
Vivek

Also I do not see any dependent upstream BZ for this. If we have not opened one, we ideally need to have that opened as well.

Sorry that is my fault. I was holding until we had more of an RCA and someone else was able to repro. I was starting to believe that I had something wrong in my env and I

should look into potentially closing things, but now I believe it is just difficult to trigger.

-b

- rejoy (rmc)

On 03/22/2015 09:49 PM, Ben Turner wrote:

----- Original Message -----

From: "Sayan Saha" <ssaha@redhat.com>

To: "Nagaprasad Sathyanarayana" <nsathyan@redhat.com>, "Ben Turner" <bturner@redhat.com>

Cc: "Pranith Kumar Karampuri" <pkarampu@redhat.com>, "Shyam"

<srangana@redhat.com>, "Krutika Dhananjay"

<kdhananj@redhat.com>, "Vijay Bellur" <vbellur@redhat.com>, "Alok

Srivastava" <asrivast@redhat.com>,

rejoy@redhat.com, "Satish Mohan" <smohan@redhat.com>, "Anoop Kumar Nair"

<annair@redhat.com>, "Vivek Agarwal"

<vagarwal@redhat.com>

Sent: Sunday, March 22, 2015 11:58:20 AM

Subject: Re: URGENT: New deadlock issue with the .53 build (RHGS 3.0.4 GA candidate) for random writes with IOZone.

@ Ben

Were you able to capture some performance data with 16 kb block & file sizes in the runs where we are not hitting the deadlock issue?

Yes and I have been able to consistently run without the issue since last wednesday.

@ Naga

Since this issue is not consistently reproducible I am willing to waive this for 3.0.4 GA provided Ben has been able to collect some data for the 16 kb runs.

+1 here. I have the data and this issue is no where near as reproducible as I first thought when running. I really do think we need to put some resources on investigating why these smaller block / record sizes choke gluster like this. I do think there is a race condition that gets triggered when the back end is really bogged down.

-b

Thanks
Sayan

On 3/22/15, 2:49 AM, Nagaprasad Sathyanarayana wrote:

Hi Shyam, Ben, Pranith, Krutika,

It is so great to see so much of perseverance in root causing this issue. As we all understand, we are yet to have a consistently reproducible setup to root cause the issue.

Along with this we should start asking question: should we hold the 3.0.4 release for this BZ or Fix this upstream first ASAP, without blocking 3.0.4 release.

Given 3.0.4 GA, as per the current schedule, is just 2 days away, it is very important we answer this question.

Some of the data points that can help us are

1. Is the scenario in which the issue occurred a standard setup our customers use?
2. Could this be an existing issue or is it a regression introduced in 3.0.4.

Thanks
Naga

On 22-Mar-2015, at 9:01 am, Pranith Kumar Karampuri<pkarampu@redhat.com> wrote:

On 03/22/2015 12:38 AM, Shyam wrote:

On 03/21/2015 01:17 PM, Pranith Kumar Karampuri wrote:

On 03/21/2015 09:29 PM, Pranith Kumar Karampuri wrote:

On 03/21/2015 08:42 PM, Ben Turner wrote:

@Pranith - Hmm I have been trying to repro in my 2 server 2 client replica 2 + 1 brick per server config(running on 4 total systems). So

are you saying to create 2 bricks on a single node and mount that volume to localhost instead of a client and I should it hit more regularly?

At least I hit it regularly and made a little progress 😊.

Not the same issue as the one Ben saw. The issue is happening on my machine because of slow fsyncs. Not the case with Ben's bug.

Pranith, why would the slow fsync be not what Ben observed as well?

Basically we have a situation where the server is not responding to any RPC for a few seconds, which is big enough for the client to ping timeout. As the server is showing behavior where it is rearming the socket (which would happen as the first RPC is responded to and the throttle count is reduced) there is no clear hang/deadlock on the brick process that is causing the pig timeout.

Shyam,

On all the runs, there were submit_reply failures for 61 fsyncs everytime the ping timeout happened. Not in the sosreports of the bugs.

So it is not the same issue.

I also went through the code. Not able to find any mistakes at the moment.

Will let you guys know if I find something. We definitely need a reasonably consistent reproducer to figure things out for this bug. Because after the issue happens, the state of the system is lost, because of the lock state cleanup etc.

Pranith

My 2 c's follow,

I think the issue here is that the storage system at times is not responding fast enough (too much load?) and that causes the client to ping timeout.

Why the storage system is not responding quick enough I am not sure, but that seems to be the trigger for the ping timeout.

So bottom line, we really do not have a fix for this that is straightforward and could occur even in 3.0.3 (unless there is something else that has changed in between, that I am missing).

I would say that we may need to rework the ping timeout as a mechanism not to check a sockets health but health of the server/deamon (of sorts) so maybe we should move that to a separate socket (or create a socket pair per socket, one for regular IO and the other just for ping packets). This way of the IO threads get stuck, that does not mean it is a server not responding (it does mean server is too busy doing something).

Anyway the above would not be for 3.0.4 (in my eyes). I would vote for 3.0.4 to push forward with this known issue in place, thoughts?

Shyam

Pranith

Pranith

@Naga - None of my tests worked with replication until the .50 build so I don't have the exact build this was introduced. My guess is that it was introduced with .50 but things always segfaulted before I could gather results or even get to the point to see this issue. It could have been there since day 1 of 3.0.4 but my tests never made it far enough to notice until .50.

-b

----- Original Message -----

From: "Pranith Kumar Karampuri" <pkarampu@redhat.com>
To: "Shyam" <srangana@redhat.com>, "Ben Turner" <bturner@redhat.com>
Cc: "Krutika Dhananjay" <kdhananj@redhat.com>, "Sayan Saha" <ssaha@redhat.com>, "Vijay Bellur" <vbellur@redhat.com>, "Alok Srivastava" <asrivast@redhat.com>, rejoy@redhat.com, "Satish Mohan" <smohan@redhat.com>, "Anoop Kumar Nair" <annair@redhat.com>, "Nagaprasad Sathyanarayana" <nsathyan@redhat.com>, "Vivek Agarwal" <vagarwal@redhat.com>
Sent: Saturday, March 21, 2015 8:44:59 AM
Subject: Re: URGENT: New deadlock issue with the .53 build (RHGS 3.0.4 GA candidate) for random writes with IOZone.

On 03/21/2015 06:48 AM, Pranith Kumar Karampuri wrote:

On 03/21/2015 05:45 AM, Shyam wrote:

Tried to reproduce this issue (pretty much the whole day) and

we
(i.e
BenT and myself) were not successful.

The intention was to capture rolling tcpdumps and also hold the gluster mount in gdb so that when the issue occurs, we have some level of debugability to see what the issue could be, but as we did not see the issue occurring, I have not progressed any further on this.

BenT was still attempting to reproduce, will let him comment on that.

Shyam
P.S: Pranith, as we were talking in my morning, ping me for any updates.

I will try to recreate it today. Will let you know my progress.

Hi,

Here is the progress Krutika and I made today so far:
We were able to recreate the issue 4/4 times with the following steps:

- 1) Create a plain replicate volume on a single machine.
- 2) Run "iozone -t 4 -r 16k -s 2g -i 0 -i 2" on the mount point.

It will either fail at random read or random write stage. Best case will fail in 20 minutes. Worst case it fails after an hour. But it always failed. We tried to check if the epoll_ctl optimization could be causing this issue by commenting out the optimization. It took little over an hour but failed. With the optimization it failed in ~40 minutes.

Here is something we gathered by adding some logs:

```
[2015-03-21 05:21:05.766594] I
[event-epoll.c:623:event_dispatch_epoll_handler] 0-epoll: epollin
changed 1 0 0xcfaa60<<---- It shuts EPOLLIN here
[2015-03-21 05:22:05.221265] I
[server-handshake.c:578:server_setvolume]
0-r2-server: accepted client from
pranithk-laptop-15528-2015/03/21-04:30:29:75
0660-r2-client-0-0-1 (version: 3.6.0.53)
[2015-03-21 05:22:42.841443] I
[event-epoll.c:513:event_select_on_epoll]
0-epoll: epollin changed 0 1 0xcfaa60<<--- After a minute it tries
to
open it again.
[2015-03-21 05:22:42.841455] W [socket.c:546:__socket_rvw]
0-tcp.r2-server: writev on 192.168.1.5:994 failed (Broken pipe)
```

```
[2015-03-21 05:22:42.841489] I
[event-epoll.c:618:event_dispatch_epoll_handler] 0-epoll:
epoll-ctl
called after 37 0xcfaaa8
[2015-03-21 05:22:42.841671] I
[socket.c:2290:socket_event_handler]
0-transport: disconnecting now
[2015-03-21 05:22:42.841692] I [server.c:519:server_rpc_notify]
0-r2-server: disconnecting connection from
pranithk-laptop-15528-2015/03/21-04:30:29:750660-r2-client-0-0-0
```

Client-log at the same time when the disconnect happened, so it is because of epoll-in off/on which is leading to the issue:

```
[2015-03-21 05:22:04.078439] C
[rpc-clnt-ping.c:109:rpc_clnt_ping_timer_expired] 0-r2-client-0:
server
192.168.1.5:49153 has not responded in the last 42 seconds,
disconnecting.
```

We are reading the code to come up with possible scenarios which can lead to this issue at this point. Feel free to let me know if you have any theories you want me to verify. At this moment we are sure that optimization in 'event_select_on_epoll' is not a contributing factor.

Pranith

Pranith

On 03/20/2015 10:08 AM, Ben Turner wrote:

----- Original Message -----

From: "Pranith Kumar Karampuri" <pkarampu@redhat.com>
 To: "Shyamsundar Ranganathan" <srangana@redhat.com>
 Cc: "Ben Turner" <bturner@redhat.com>, "Krutika Dhananjay" <kdhananj@redhat.com>, "Sayan Saha" <ssaha@redhat.com>, "Vijay Bellur" <vbellur@redhat.com>, "Alok Srivastava" <asrivast@redhat.com>, reivy@redhat.com, "Satish Mohan" <smohan@redhat.com>, "Anoop Kumar Nair" <annair@redhat.com>, "Nagaprasad Sathyanarayana" <nsathyan@redhat.com>, "Vivek Agarwal" <vagarwal@redhat.com>
 Sent: Friday, March 20, 2015 8:58:04 AM
 Subject: Re: URGENT: New deadlock issue with the .53 build (RHGS 3.0.4 GA candidate) for random writes with IOZone.

Shyam,

So far this is what I gathered from inspecting the variables in gdb:

- 1) Server thinks it is still connected to client xlator:
- 2) client xlator thinks it is disconnected.
- 3) client-xlator's conn object thinks it is connected.

Ping me once you come online.

Same with me 😊 Here is my take on what is happening, I think

it
is
a combination of things that are on their own OK but when
chained
together things break down:

-I am only hitting the issue with smaller block sizes, when
I go
above 16k everything is OK. I ran some tests on the
backend:

<https://mojo.redhat.com/docs/D0C-1020792>

And sure enough 16k reads are really choked, and IIRC reads
run
first with IOzone and that is the one that is hitting the
issue(IOZone does both random tests in the same run so I
don't
know
for sure yet if that is what is hitting the problem). I am
investigating now if this is thinp related and wasn't
identified
as
part of our previous fixes or if this is new. I would think
that
this array can do better than 38 MB / sec rand reads...

-Couple the slow backend with the huge pipe that these fast
NICs
can
dump we started seeing the throttling issues. I didn't see
the
throttling issue with larger blocks because the back end was
able
to
service things timely enough to not hit it. I think this is
one
bug
and we need to identify the issue with throttling, we may be
able
to
tune our way out with tunibles but I havent had any success
as
of
yet.

-At some point the client hits the timeout and thinks is
disconnected. There is likely a bug here like Pranith was
describing, a race condition during the disconnect(that was
caused
by throttling?).

-From there we need to figure out if there is something else
happening on why the system doesnt recover / recovery is
taking
so
long. There may be a timeout that needs tweaked or a race
condition
or something here as well. Or it could be that the client

just
thinks its still connected and everything is working
properly,
not
sure.

LMK if you need anything from me.

-b

Pranith

On 03/20/2015 05:34 PM, Ben Turner wrote:

----- Original Message -----

From: "Pranith Kumar Karampuri" <pkarampu@redhat.com>
To: "Krutika Dhananjay" <kdhananj@redhat.com>, "Sayan
Saha"

<ssaha@redhat.com>

Cc: "Shyamsundar Ranganathan" <srangana@redhat.com>,
"Vijay

Bellur"

<vbellur@redhat.com>, "Alok Srivastava"

<asrivast@redhat.com>, reiy@redhat.com, "Satish

Mohan"

<smohan@redhat.com>, "Anoop Kumar

Nair" <annair@redhat.com>,
"Nagaprasad Sathyanarayana" <nsathyan@redhat.com>,
"Vivek

Agarwal"

<vagarwal@redhat.com>, "Ben Turner"

<bturner@redhat.com>

Sent: Friday, March 20, 2015 7:56:06 AM

Subject: Re: URGENT: New deadlock issue with the .53
build

(RHGS

3.0.4 GA

candidate) for random writes with IOZone.

With krutika, Rafi, manoj's help we were able to
re-create
the

issue.

The setup is in that state. Please do no disturb the

setup.

Great news! And eek I just ran one last test on the

back end

as

soon as I

woke up, I don't _think_it should affect you but if you

see a

self

heal on

files it was me running on the back end. I ran at

about 7:45

my

time if

you see stuff in the logs. I won't touch them again,

sry.

-b

Pranith

On 03/20/2015 10:23 AM, Krutika Dhananjay wrote:
Pranith had fixed a similar issue in upstream @
<http://review.gluster.org/7531> about a year ago
which has

also
made it
to downstream.

This one seems similar where an UNLOCK request
doesn't get
through due
to RPC throttling on the server side, because of
which
subsequent lock
requests get blocked.

Will need to reproduce the issue and capture the
statedump
of
the
bricks to start with.

Not sure if the last change to RPC throttling -

commit c48cbccfafbcf71aaad4ed7d868dbac609bc34fe

Author: Shyam<srangana@redhat.com>

Date: Mon Feb 23 10:00:39 2015 -0500

epoll: Fix broken RPC throttling due to MT epoll

- has anything to do with it. Will investigate
further.

-Krutika

*From: *"Sayan Saha"<ssaha@redhat.com>
*To: *"Shyamsundar Ranganathan"
<srangana@redhat.com>,
"Pranith
Kumar Karampuri"<pkarampu@redhat.com>,
"Vijay Bellur"
<vbellur@redhat.com>, "Krutika Dhananjay"
<kdhananji@redhat.com>
*Cc: *"Alok Srivastava"
<asrivast@redhat.com>,
<reijy@redhat.com>,
"Satish Mohan"<smohan@redhat.com>, "Anoop
Kumar Nair"
<annair@redhat.com>, "Nagaprasad
Sathyanarayana"
<nsathyan@redhat.com>, "Vivek Agarwal"
<vagarwal@redhat.com>, "Ben
Turner"<bturner@redhat.com>
*Sent: *Friday, March 20, 2015 2:26:43 AM

*Subject: *URGENT: New deadlock issue with
the .53
build
(RHGS
3.0.4 GA candidate) for random writes with
IOZone.

+ Adding Ben Turner

All,

Not sure who can help but I am emailing a
set of
folks
who
I think
can help.

BZ here -
https://bugzilla.redhat.com/show_bug.cgi?id=1203850

Ben T has the environment up for people who
want to
look.

Thanks
Sayan