

GDP - Bug #45

Reading errors (timeout, assertion failure, end of file)

08/21/2016 10:48 PM - Nitesh Mor

Status:	Closed	Start date:	08/21/2016
Priority:	Normal	Due date:	
Assignee:	Eric Allman	% Done:	0%
Category:	libgdp	Estimated time:	0.00 hour
Target version:			
Description			
Reading on certain logs/record combinations fails with one out of three results. (on commit:5da0351b).			
What causes the issue:			
<pre>\$./gdp-reader -D *=20 -t -f 1551471 -n 1 edu.berkeley.eecs.swarmlab.device.c098e5300003</pre>			
What's the result?			
<ul style="list-style-type: none">• Assertion failure:			
<pre>Assertion failed at gdp_gcl_ops.c:447: require: (gcl) != NULL && EP_UT_BITSET(GCLF_INUSE, (gcl)->flags) Aborted</pre>			
<ul style="list-style-type: none">• End of file:			
<pre>exiting after 0 records with status WARNING: end of file [EPLIB:generic:3]</pre>			
<ul style="list-style-type: none">• Connection timed out			
<pre>Couldn't open GCL 4RPmOEEboug46rExalX8F9HdGKB9aH5kAzhTfXNW4DY: ERROR: Connection timed out [EPLIB:errno:110] [ERROR] gdp-reader: Cannot open GCL: ERROR: Connection timed out [EPLIB:errno:110] exiting after 0 records with status ERROR: Connection timed out [EPLIB:errno:110]</pre>			

History

#1 - 08/22/2016 08:55 AM - Eric Allman

The first observation is that that log exists on both gdp-01 and gdp-04. It isn't clear what the semantics of that will be, and it's possible that the multiple routes could be part of the problem. I've already logged this problem as a GDP router issue ([#44](#)).

Whether this is the only problem will require some more investigation. In any case, we need some better way to handle this.

#2 - 08/22/2016 09:39 AM - Nitesh Mor

Eric Allman wrote:

The first observation is that that log exists on both gdp-01 and gdp-04.

Could it be because of the MQTT-gateway startup script attempting to create a log each time it starts? Even though gcl-create checks for existing

logs, but router instability can make those checks unreliable.

#3 - 08/22/2016 09:56 AM - Eric Allman

Yes, I suspect that's exactly the cause. People seem to like the idea of log auto-creation, but it has implications. Perhaps we should pull it out at least until we have replication going, which will reduce the odds of failure.

By the way, this is exactly why I've been making it a point to host the logs on the same physical hardware on which the mqtt-gdp-gateway code runs. It doesn't solve the problem, but it reduces it.

eric

#4 - 08/23/2016 03:16 PM - Eric Allman

- *Category set to libgdp*

- *Assignee set to Eric Allman*

#5 - 08/24/2016 02:32 PM - Eric Allman

- *Status changed from New to Feedback*

This is actually several problems. I think I have them all fixed.

- The gdp library was releasing a GCL handle when a command timed out even if it was going to resend the command. This resulted in an access to a freed data structure. [First symptom]
- There were multiple versions of the log (gdp-01 and gdp-04). If you got the version on gdp-04, it correctly said that the record was not found. [Second symptom]
- The log daemon was raising an assertion failure if there was a corrupt index. The log edu.berkeley.eecs.swarmlab.device.c098e5300003 is corrupt on gdp-01. This caused the server to crash, and even though the log open succeeded, subsequent attempts to read it would return "no route". [Third symptom.]
- If a command responded with a "server NAK", the command was treated as though it should be retried. This doesn't directly correspond to any of the symptoms, but was still related.

Please try this and see if you can reproduce it. If you can, put the issue back to In Progress and I'll give it another try.

#6 - 08/24/2016 05:24 PM - Nitesh Mor

Doesn't work, still. However, this is primarily an issue of corrupt index on the log-server side, which causes it to crash. Reporting it as [#51](#), but I guess we can safely close this one.

#7 - 08/24/2016 05:51 PM - Nitesh Mor

- *Status changed from Feedback to Closed*

My bad, wasn't working on the latest version. Right now, the original gdp-reader command with the specified arguments results in the following

```
gdp_pdu_proc_resp(NAK_S_INTERNAL for CMD_READ): SEVERE: 500 internal server error [Berkeley:Swarm-GDP:500]  
exiting after 0 records with status SEVERE: 500 internal server error [Berkeley:Swarm-GDP:500]
```

Since corruption of index is a server error that we don't currently handle, returning error 500 seems to be the most obvious response. Closing this for now.

Files

timeout.log	3.39 KB	08/22/2016	Nitesh Mor
eof.log	2.19 KB	08/22/2016	Nitesh Mor
assertion.log	4.01 KB	08/22/2016	Nitesh Mor