

GDP - Bug #60

New GDP log created on log server inaccessible by gdp-reader.

09/28/2016 04:04 PM - David Mellis

Status:	Closed	Start date:	09/28/2016
Priority:	Urgent	Due date:	
Assignee:	Eric Allman	% Done:	0%
Category:	Routing Layer	Estimated time:	0.00 hour
Target version:			
Description			
I created a new GDP log with:			
<pre>dhcp-45-14:apps mellis\$./gcl-create -C mellis@berkeley.edu -k none edu.berkeley.eecs.bid.mellis.e sp Created new GCL D0nCdKDQxuu3-M9ul5KHGjgAQj0EwJKwTYaEoxfsD4I on log server edu.berkeley.eecs.gdp-03.gdplogd exiting with status OK</pre>			
Then got an error when I tried to read it:			
<pre>dhcp-45-14:apps mellis\$./gdp-reader -s "edu.berkeley.eecs.bid.mellis.esp" Reading GCL D0nCdKDQxuu3-M9ul5KHGjgAQj0EwJKwTYaEoxfsD4I [ERROR] gdp-reader: Cannot open GCL: ERROR: 600 no route available [Berkeley:Swarm-GDP:600] exiting after 0 records with status ERROR: 600 no route available [Berkeley:Swarm-GDP:600]</pre>			
This is with ~/.ep_adm_params/gdp set to swarm.gdp.routers=gdp-01.eecs.berkeley.edu; gdp-02.eecs.berkeley.edu			

History

#1 - 09/28/2016 05:23 PM - Eric Allman

- Status changed from New to In Progress
- Assignee set to Eric Allman

A bit more detail on this problem: The exact sequence must be:

1. gcl-create
2. gdp-reader -s

This causes the gdplogd to crash in post_subscribe with:

```
Assertion failed at logd_proto.c:720: insist:
req->gcl != NULL
```

If there is an intermediate step to just read the log (without -s) then the problem does not occur.

#2 - 09/29/2016 11:36 AM - Eric Allman

- Status changed from In Progress to Feedback

For some reason I am no longer able to reproduce the problem, despite having made no code changes since yesterday. I've also checked out an old version (the same as running on gdp-0[1234], which also works. David, could you check to see if it's still failing for you?

#3 - 09/29/2016 11:43 AM - David Mellis

I can still reproduce, but the problem only seems to happen when the log is created on a log server other than gdp-01 (or maybe gdp-02):

```
dhcp-45-14:apps mellis$ ./gcl-create -C mellis@berkeley.edu -k none edu.berkeley.eecs.bid.mellis.esp3
Created new GCL ipbR0X_iH2ys0EOhFYmJg3nJHuaa_prManMHZsUgvNo
  on log server edu.berkeley.eecs.gdp-01.gdplogd
exiting with status OK
dhcp-45-14:apps mellis$ ./gdp-reader -s edu.berkeley.eecs.bid.mellis.esp3
Reading GCL ipbR0X_iH2ys0EOhFYmJg3nJHuaa_prManMHZsUgvNo
^C
Exiting on signal 2
dhcp-45-14:apps mellis$ ./gcl-create -C mellis@berkeley.edu -k none edu.berkeley.eecs.bid.mellis.esp4
Created new GCL bt-ndEGfEJ6e1TPsUhPBP7Y8IPYZ3j19vnNB9hBSeCE
  on log server edu.berkeley.eecs.gdp-01.gdplogd
exiting with status OK
dhcp-45-14:apps mellis$ ./gdp-reader -s edu.berkeley.eecs.bid.mellis.esp4
Reading GCL bt-ndEGfEJ6e1TPsUhPBP7Y8IPYZ3j19vnNB9hBSeCE
^C
Exiting on signal 2
dhcp-45-14:apps mellis$ ./gcl-create -C mellis@berkeley.edu -k none edu.berkeley.eecs.bid.mellis.esp5
Created new GCL LH0AAfMgh4k59mprk3ELhs9Cb7Ch9eVpWvRCabJ7Kyg
  on log server edu.berkeley.eecs.gdp-04.gdplogd

exiting with status OK
dhcp-45-14:apps mellis$
dhcp-45-14:apps mellis$ ./gdp-reader -s edu.berkeley.eecs.bid.mellis.esp5
Reading GCL LH0AAfMgh4k59mprk3ELhs9Cb7Ch9eVpWvRCabJ7Kyg
[ERROR] gdp-reader: Cannot open GCL:
  ERROR: 600 no route available [Berkeley:Swarm-GDP:600]
exiting after 0 records with status ERROR: 600 no route available [Berkeley:Swarm-GDP:600]
dhcp-45-14:apps mellis$
```

#4 - 09/29/2016 05:21 PM - Nitesh Mor

A bit of detail (not sure how important it is):

- gdp-01 is running on 14.04, where a watcher programs restarts the log-server on a server-crash in a few minutes.
- gdp-0[2,3,4] are running systemd, which is also setup for restarts in case of crashes (not sure of the frequency, something systemd specific).

This may explain the different behavior, but it'll be nice to have a look at the logs if there are server restarts.

#5 - 09/30/2016 07:02 AM - Eric Allman

I did get a chance between meetings yesterday to look through logs, and I didn't see any restarts around the same time as the logs were created. This surprised me because when I did get it to fail, it definitely killed the log daemon, but it appears that David's problem is that the advertisement isn't being processed by the routing layer. I don't know whether the advertisement was never written by the log daemon, if the advertisement is stuck somewhere in the (TCP) network, or if the router read it but didn't process it.

I suspect that a reboot on the servers (which need a systems update anyway) would fix the problem temporarily, but at some point we need to figure out the root problem. I've been trying to figure out how to add debugging hooks before I restart, but I'm running out of ideas.

#6 - 10/04/2016 09:31 AM - Anonymous

- File *gdp600Message.txt* added

FWIW - The Ptolemy II nightly build has been failing for a few days with ERROR: 600 no route available [Berkeley:Swarm-GDP:600].

In particular, it looks like

`$PTII/bin/vergil $PTII/ptolemy/actor/lib/jjs/modules/gdp/test/auto/GDPLogSubscribe.xml` has been failing since at least Sep 30, 2016 9:38:00 PM

I poked around and it looks like using `edu.berkeley.eecs.gdp-01.gdplogd` for `logdname` is failing, see *gdp600Message.txt*

If, after running, I change to `gdp-03`, then I get the assert error and the entire process fails:

```
Assertion failed at gdp_gcl_ops.c:453: require:  
(gcl) != NULL && EP_UT_BITSET(GCLF_INUSE, (gcl)->flags)
```

There was another bug about this. Asserts are really risky.

However, if I change to `gdp-03` and never run using `gdp-01`, then the model works.

I'm hoping the log for `gdp-01` helps.

#7 - 10/04/2016 03:31 PM - Eric Allman

- Project changed from *GDP* to *Click GDP Routers*

- Status changed from *Feedback* to *In Progress*

- Priority changed from Normal to Urgent

Just to bring everyone up to date, I'm looking at this. The problem is clearly at the routing layer, so I'm moving this to that project. It's persistent (reboots do not fix the problem, although they sometimes make it move), and the problem seems to come and go without any obvious reason.

Briefly, regardless of which router (gdp-01, gdp-02, gdp-03, and gdp-04) you connect to, the same set of logs should be available. What I'm finding is that there are three outcomes for each router: (a) it works as expected; (b) the router claims the log is unknown ("No Route"); or (c) it times out, giving an error such as the one that Christopher reported. The logs having problems come and go, and most logs seem to work (i.e., give consistent results) most of the time.

This problem seems to have started recently, which leads me to suspect that it may have something to do with the absolute number of logs in the GDP, which is constantly growing. The other clue is that I'm seeing some "false routes", i.e., routes for obviously incorrect log names (such as AAAAAAAAAAB4AAAAfAAAAMDdSf39fwAAAAAAAAAAAA).

#8 - 10/06/2016 07:07 AM - Eric Allman

I found a definite bug in the router code that caused log names to get mangled in some circumstances if they had zero octets. That fix has been running since about 16:30PT yesterday afternoon, and so far so good. However, I'm not clear how this relates to routes that seem to appear and disappear, so I'm going to leave this as "In Progress" for at least a few days.

To be a bit more explicit, this turned out to be several intertwined problems. David's original problem was caused by his extraordinary bad luck in choosing names that had zero octets in the internal name. His second instance of the problem (which only worked on logs created on gdp-01) was again just a coincidence. The assertion failure in `post_subscribe` appears to have been an unrelated problem where certain error conditions on `subscribe` commands weren't handled properly. The other problems seem to be because, due to routing issues, duplicate logs were sometimes being created, which badly confused the routing layer.

It's a complete mystery to me why the problem only appeared recently. My guess that it was just that, as we add more logs, the chances of a log with a zero octet in the name goes up.

#9 - 02/08/2017 12:02 PM - Eric Allman

- Status changed from In Progress to Closed

This problem seems to have gone away, so I'm closing it. It was fixed in `commit:0e072ada`.

#10 - 04/09/2019 12:28 PM - Eric Allman

- Project changed from Click GDP Routers to GDP

- Category set to Routing Layer

Files

gdp600Message.txt	7.93 KB	10/04/2016	Anonymous
-------------------	---------	------------	-----------