

LTTng-tools - Bug #1017

consumerd memory leak

06/03/2016 07:28 PM - Florea Irinel

Status:	Feedback	Start date:	06/03/2016
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
Consumerd memory usage increases in the following scenario:			
<pre>lttng create mysession --snapshot -o /var/snapshots/ lttng enable-channel mychannel -u -s mysession lttng enable-event -u -a -c mychannel -s mysession lttng start mysession sleep 1 lttng destroy mysession</pre>			
Starting tracing the session adds 20 kb to the memory usage of consumerd. Running this in a loop, consumerd starts to eat serious memory. Make sure there are traced application running. The bug doesn't reproduce without. Also from what I've tested it only happens in snapshot mode. Bug reproduces on multiple machines with different architectures(tried armv7 with 3.10.62 kernel and ubuntu 14.04 3.19.0 x86_64), lttng 2.7.0 on armv7 and lttng 2.7.3 on x86_64.			

History

#1 - 06/03/2016 07:41 PM - Jonathan Rajotte Julien

- Status changed from New to Feedback

Hi,

Were you able to reproduce/observe on stable v2.8?

Cheers

#2 - 06/04/2016 03:38 AM - Florea Irinel

Hi,

Yes, I did just now on the ubuntu machine. Same behaviour.

Regards.

#3 - 06/13/2016 08:29 AM - Florea Irinel

Hi,

I've been debugging a bit. Within the sessiond_poll thread ,on receiving command 11 - LTTNG_CONSUMER_CLOSE_METADATA, the wait fd/poll pipe of the metadata stream is not closed when using snapshot sessions. Because the channel key passed to close_metadata() (calling consumer_find_channel() and passing the channel ref to lttng_ustconsumer_close_metadata()) corresponds to a channel for which the condition !channel->metadata_stream (within lttng_ustconsumer_close_metadata()) evaluates as true. Causing the metadata_poll thread to not get any LPOLLHUP event over the pipe and it doesn't call consumer_del_metadata_stream() to free the shm_table.
Does any of this make any sense? Do you think this is possibly the reason for the behaviour I'm getting?

Regards.

#4 - 06/13/2016 10:19 AM - Jonathan Rajotte Julien

Florea Irinel wrote:

Hi,

I've been debugging a bit.

Good!

Within the sessiond_poll thread ,on receiving command 11 - LTTNG_CONSUMER_CLOSE_METADATA, the wait fd/poll pipe of the metadata stream is not closed when using snapshot sessions. Because the channel key passed to close_metadata() (calling consumer_find_channel() and passing the channel ref to lttng_ustconsumer_close_metadata()) corresponds to a channel for which the condition !channel->metadata_stream (within lttng_ustconsumer_close_metadata()) evaluates as true. Causing the metadata_poll thread to not get any LPOLLHUP event over the pipe and it doesn't call consumer_del_metadata_stream() to free the shm_table.

Does any of this make any sense? Do you think this is possibly the reason for the behaviour I'm getting?

Well, were you able to fix your issue ? If so do you have a patch we can look at?

Since you have a reproducer I would suggest that you poke around and try a fix while we find some time to look at it.

Cheers

Regards.

#5 - 06/16/2016 07:20 AM - Florea Irinel

Hi,

Since on snapshot sessions the metadata stream is not monitored by the metadata thread the stream has to be freed somewhere else (the closing of the wait fd's does not suffice). Do you agree? If yes, what is the best approach :

- to make sessiond to issue a LTTNG_CONSUMER_DESTROY_CHANNEL command after the LTTNG_CONSUMER_CLOSE_METADATA?
- in consumerd at receiving LTTNG_CONSUMER_CLOSE_METADATA if the metadata stream is not monitored , notify channel thread using notify_thread_del_channel() ?
- something else.

Regards.

#6 - 06/16/2016 07:21 AM - Florea Irinel

Florea Irinel wrote:

Hi,

Since on snapshot sessions the metadata stream is not monitored by the metadata thread the stream has to be freed somewhere else (the closing of the wait fd's does not suffice). Do you agree? If yes, what is the best approach :

- to make sessiond to issue a LTTNG_CONSUMER_DESTROY_CHANNEL command after the LTTNG_CONSUMER_CLOSE_METADATA if the session is in snapshot mode?
- in consumerd at receiving LTTNG_CONSUMER_CLOSE_METADATA if the metadata stream is not monitored , notify channel thread using notify_thread_del_channel() ?
- something else.

Regards.

#7 - 06/21/2016 09:31 AM - Florea Irinel

- File 0001-Fix-consumerd-memory-leak-when-using-snapshot-sessio.patch added

Hi,

Went with the second approach and created the patch attached. The memory allocated when starting tracing on snapshot sessions is now freed when destroying the session. From what I tested it does not interfere with other modes of sessions. Please let me know what you make of this .

#8 - 06/21/2016 10:08 AM - Jonathan Rajotte Julien

Hi Florea,

I would recommend that you send a proper email to the mailing with the patch inline with a more verbose message explaining everything in the commit message (valgrind report, causes etc.) Please read <http://lttng.org/community/#contributors-guide> for more detail.

I'm very glad that you went the patch way! We are always looking for new contributors.

Cheers!

Files
