

LTTng-UST - Bug #1053

lttng-ust doesn't build on Ubuntu 12.04

07/27/2016 10:35 AM - Simon Marchi

Status:	Resolved	Start date:	10/04/2016
Priority:	Normal	Due date:	
Assignee:	Mathieu Desnoyers	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
UST doesn't build on Ubuntu 12.04, because it uses fields of the perf_event_mmap_page structure that are not available in the version of linux/perf_event.h shipping with that version of Ubuntu.			
Build log: https://launchpadlibrarian.net/275214770/buildlog_ubuntu-precise-amd64.ust_2.8.x+stable+bzr2111+pack66+201607262232~ubuntu12.04.1_BUILDING.txt.gz			
<pre>libtool: compile: gcc -DHAVE_CONFIG_H -I. -I.. -I../include/lttng -I../include -I../include -D_FORTIFY_SOURCE=2 -I/usr/lib/jvm/default-java/include -I/usr/lib/jvm/default-java/include/linux -fno-strict-aliasing -Wall -g -O2 -fstack-protector --param=ssp-buffer-size=4 -Wformat -Wformat-security -Werror=format-security -c lttng-context-perf-counters.c -fPIC -DPIC -o .libs/lttng-context-perf-counters.o lttng-context-perf-counters.c: In function 'read_perf_counter': lttng-context-perf-counters.c:118:22: error: 'struct perf_event_mmap_page' has no member named 'pmc_width' lttng-context-perf-counters.c:119:22: error: 'struct perf_event_mmap_page' has no member named 'pmc_width' make[3]: *** [lttng-context-perf-counters.lo] Error 1</pre>			
Subtasks:			
Bug # 1065: Upgrade failure?			Resolved

History

#1 - 07/27/2016 11:16 AM - Mathieu Desnoyers

One possible solution would be to grab our own copy of perf-event.h into lttng-ust, and use that local header.

As we do that, we should also check the "version" field of the ABI when interacting with the kernel to ensure that the pmc_width is populated, else we should refuse to interact with the kernel.

Thoughts ?

#2 - 07/27/2016 11:51 AM - Simon Marchi

Yeah it makes sense. I suppose that with the stock 12.04 kernel the ABI version will be too old so the perf counter feature won't be available, but at least it would allow using the rest of lttng-ust. It might not be super useful in itself to get lttng-ust working on 12.04 (it's super oooold), but it would be useful to put the infrastructure to check the ABI version in place for future breakages.

#3 - 07/29/2016 05:08 PM - Anonymous

linux/perf_event.h is part of the linux-headers package. I guess many users will have kernel headers installed if they also want to compile lttng-modules, but that may not be the case of all of them. Is UST supposed to have a hard dependency on perf/linux-headers? Could that be a configure option, like --disable-perf-counters ?

#4 - 07/31/2016 07:24 PM - Mathieu Desnoyers

Alex: we're really talking about kernel headers exposing kernel ABI to user-space (make headers_install), not the kernel headers for kernel modules (which are packaged by some distros) here.

The Linux kernel ABI exposed to userspace never breaks, only expands. So it would be safe to build lttng-ust against a local copy of the perf header here.

I don't think we should add yet another configure option. We should instead add a copy of the header into lttng-ust, and dynamically check that the kernel ABI version is appropriate to use the fields we need to use.

#5 - 08/18/2016 02:16 PM - Mathieu Desnoyers

I'm preparing a patch.

#6 - 08/18/2016 08:21 PM - Mathieu Desnoyers

- Status changed from New to Resolved

- Assignee set to Mathieu Desnoyers

The following commit should fix this issue:

```
commit 77d7fa989f79f3c2cb14d26576cbbb54b0637594
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date: Thu Aug 18 14:21:46 2016 -0400
```

Fix: perf counters build against kernel headers < 3.12

Copy Linux kernel perf_event.h installed headers into lttng-ust to know the recent ABI layout, and use the bit description detailed in the following Linux kernel commit:

<https://github.com/torvalds/linux/commit/fa7315871046b9a4c48627905691dbde57e51033>

to check whether the kernel supports rdpmc.

Fall-back on the perf read system call for kernels prior to 3.12, because older kernels have an ABI bug where a union was used for both cap_usr_time and cap_usr_rdpmc.

Ensure setup_perf set the pc pointer value before checking whether we need to the file descriptor open or not.

Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>

#7 - 08/19/2016 10:09 AM - Simon Marchi

The build passed on launchpad, but I haven't tested the runtime.

Thanks!