

Linux kernel client - Feature #2374

ceph-client: start laying the groundwork for Linux tracepoints

05/02/2012 05:46 AM - Alex Elder

Status:	New	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Alex Elder		
Category:			
Target version:			
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:			
Description <p>Linux supports a mechanism for very efficiently inserting trace points in code, which allow for an arbitrary function to be executed at each point when enabled. This can be used for fine-grained tracing of activity in the code at runtime, including recording and exporting to user space data that is available at the time an event occurs. Control of whether a particular trace point (or group of them) is enabled is provided by pseudo files under /sys/kernel/debug/tracing.</p> <p>A trace point "call" occupies only a few bytes of instruction at the place it's used. Conceptually it amounts to a test of a memory location, and if it's non-zero, call a function whose address is recorded in a separate memory location. This allows both memory locations to exist but hold zero value for the default inactive case. This construct also allows trace points can be installed at run-time (via a separate kernel module, for example), so the actual jumped-to functions do not even need to exist in memory unless and until that is desired.</p> <p>Event trace points are structured to gather some data when they are called, and that data can be extracted via other utilities and then graphed or aggregated.</p> <p>There is a lot of flexibility, but there's some sort of baseline infrastructure that you need to put in place for a module to begin using these. Once the baseline is done you can add new trace points easily.</p> <p>I would really like to do this soon, given the trouble I've had understanding what's going on for bug 2260.</p>			

History

#1 - 05/02/2012 09:03 AM - Sage Weil

- Status changed from New to 12

Sounds like a great idea to me!

#2 - 06/28/2012 09:00 AM - Alex Elder

- File 0001-ceph-add-tracepoints-for-message-submission-on-read-.patch added
- File 0002-ceph-add-tracepoints-for-message-send-queueing-and-c.patch added

No progress on this yet.

However, I got this e-mail from Jim Schutt shortly after creating this bug, and just want to document it here. I said I would gladly

use his work as a starting point once I get a chance to work on this.
(The two patch messages--ported to the current state of the testing
branch--are attached.)

Hi Alex,

I ran across tracker [#2374](#) today - I've been carrying these two
tracepoint patches for a while. Perhaps you'll find them useful.

Jim Schutt (2):
ceph: add tracepoints for message submission on read/write requests
ceph: add tracepoints for message send queueing and completion, reply
handling

```
fs/ceph/addr.c          |      8  +±
fs/ceph/file.c          |      6  +
include/trace/events/ceph.h |  144  ++++++
net/ceph/messenger.c    |      9  ++=
net/ceph/osd_client.c    |      8  ++
5 files changed, 174 insertions(+), 1 deletions(-)
create mode 100644 include/trace/events/ceph.h
```

#3 - 08/20/2012 03:35 PM - Sage Weil
- Priority changed from High to Normal

#4 - 12/05/2019 09:34 PM - Patrick Donnelly
- Status changed from 12 to New

Files

0001-ceph-add-tracepoints-for-message-submission-on-read-.patch	5.41 KB	06/28/2012	Alex Elder
0002-ceph-add-tracepoints-for-message-send-queueing-and-c.patch	4.75 KB	06/28/2012	Alex Elder