LTTng and Related Projects Update

DORSAL Progress Meeting

May 2024

EfficiOS

Outline

- General Updates
- Project Updates
- Ongoing Work
- SIDE ABI (libside)
- Open Source Community Updates

General Updates

General Updates

- Common Trace Format 2 (CTF 2),
- Monthly HPC Collaboration Meetings,
- Reaching out to users.

Common Trace Format 2 (CTF 2)

Common Trace Format 2 (CTF 2) specification has been finalized and published!

- Announcement - https://lists.lttng.org/pipermail/lttng-dev/2024-March/030743.html
- Specification - https://diamon.org/ctf/CTF2-SPEC-2.0.html
- Test traces (WIP) - https://review.lttng.org/c/babeltrace/+/8645

Monthly HPC Tool Collaboration Meetings

- Started hosting monthly meetings with large HPC labs.
- Goal: Pooling resources when possible
  - Coordinate development efforts for large cluster trace tooling,
  - Find common strategies when possible.

- Attendees:
  - EfficiOS (organizer),
  - Argonne National Lab,
  - Lawrence Livermore National Lab,
  - Oak Ridge National Lab,
  - AMD,
  - Polytechnique.

Reaching Out to Users

- Historically we've mainly had contact with our client's tool developers (e.g. receiving specific feature requests).
- Leads to lack of visibility on impact of changes, extremely long feedback loops (i.e. waterfall).
- Working towards developing active connections with tracing users (e.g. troubleshooters) to:
  - Shorten feedback loops,
  - Have more impact in less time (more effective software development).

Project Updates

Project Updates

- Bug fixes released for project stable branches.
- Added support for Linux kernels 6.{7,8,9} to LTTng-modules.

ExaTracer

- Working in collaboration with AMD to provide a tracing solution for the new El Capitan cluster at Lawrence Livermore National Laboratory (LLNL).

ExaTracer

- LTTng does the tracing heavy lifting.
- Gather tracing data about work distribution (instrumentation of OpenMPI and CrayMPI).
- Integration with GPU tracing (instrumentation of ROCm APIs)
- Starting to be deployed on test systems at AMD and LLNL

ExaTracer Ecosystem Overview 2023

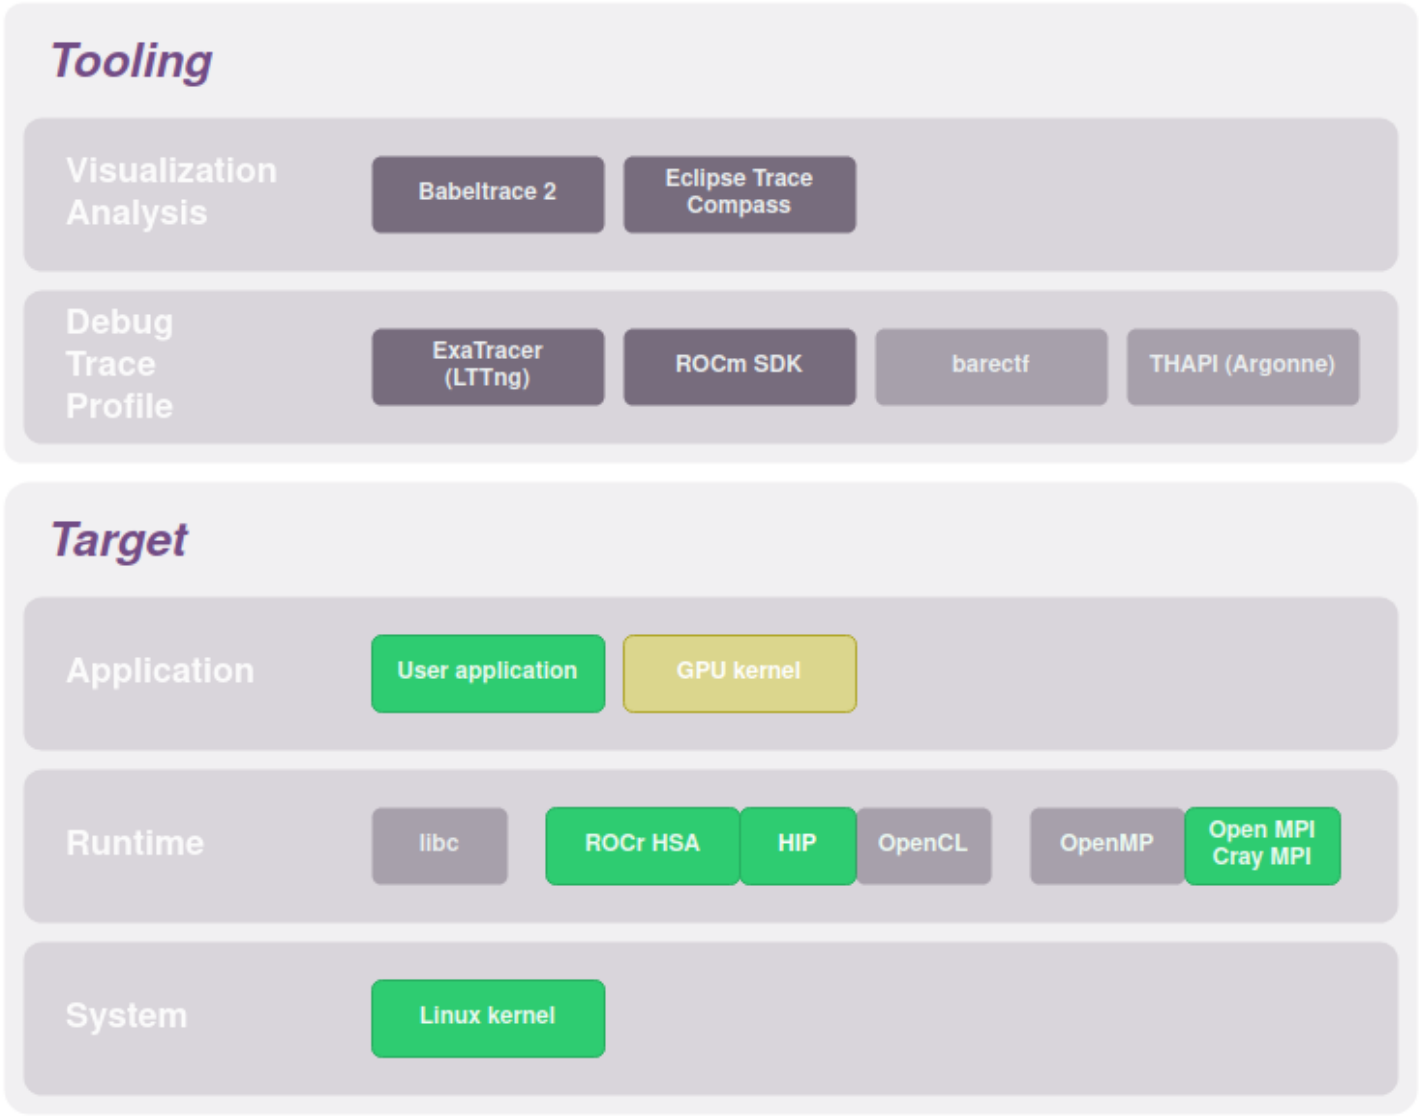

Ongoing Work

Ongoing Work

Babeltrace

- Reading and producing CTF 2 traces (2.1, expected Q4 2024)

LTTng

- Trace Hit Counters (upcoming 2.14, release date TBD)
- Producing CTF 2 traces (upcoming 2.15, release date TBD)

SIDE ABI (libside)

SIDE ABI RFC (libside)

- The SIDE ABI is currently at RFC stage, aiming to create a specification.
  - https://github.com/efficios/libside/blob/master/doc/rfc-side-abi.txt
- Runtime/language agnostic,
- Supports multiple concurrent tracers,
- Instrumentation is not specific to a tracer,
  - No need to rebuild applications if using a different tracer,
- Instrumentation can be either static or dynamic,
- Supports complex/nested types,
- Supports both static and dynamic types,
- Libside is a C/C++ reference implementation for the System V ELF ABI.

Open Source Community Updates

Open Source Community Updates

Linux Kernel

- Accurately identify architectures with virtually aliasing data cache
  - contributed by EfficiOS,
  - re-enables DAX filesystem mount option on 32-bit ARM and MIPS
  - merged upstream, available in Linux v6.9.
- Restore availability of DAX for recovering LTTng-UST traces across crash/reboot on 32-bit ARM and 32-bit/64-bit MIPS.
- Speed up LTTng-modules on 32-bit/64-bit ARM, MIPS, and PowerPC by eliminating superfluous data cache flushing.

Open Source Community Updates

Conferences

- Co-organize a Tracing and Perf Events microconference at Linux Plumbers Conference 2024.
- No Tracing Summit planned for 2024.

Open Source Community Updates

GNU libc

- Patch series introducing support for extensible RSEQ (Linux v6.3) was implemented by EfficiOS.
  - Being reviewed by GNU libc maintainers,
  - Targets inclusion into GNU libc 2.40 (August 1st, 2024).
- Preliminary step needed for future work:
  - Extend RSEQ with new fields allowing indexing ring buffers with a concurrency ID scheme which use memory efficiently in containers.
  - Per-ipc-namespace concurrency ID allocation.

Open Source Community Updates

Librseq

- Introduce CPU-Local Storage allocator (RSEQ per-CPU mempool).
  - Analogous to TLS (Thread-Local Storage) memory,
  - Per-CPU rather than per-thread,
  - For dynamic memory allocation rather than global variables,
  - Prevents false-sharing and memory waste due to cache line size padding.
- Clean up public API in prevision of gradual upstreaming into GNU libc.
- Those are preparatory steps before an official librseq release, which is needed before LTTng-UST can use it to speed up the ring buffer and trace hit counters.

EfficiOS Roadmap

EfficiOS Roadmap

- Babeltrace 2.1: Q4 2024
- Userspace RCU 0.15: Q3 2024
- LTTng 2.14/2.15: release date TBD
- libside: release date TBD
- librseq: release date TBD

Questions ?

- Links:
  - https://www.efficios.com
  - https://lttng.org
  - https://babeltrace.org
  - https://diamon.org
  - https://barectf.org

References

- Common Trace Format 2 Specification

  https://diamon.org/ctf

- libside repository

  https://github.com/efficios/libside