printk: kthreads and atomic consoles for mainline

<john.ogness@linutronix.de>

Linux Plumbers Conference 2021

Real-Time Microconference

printk and real-time

What does printk have to do with PREEMPT_RT?

printk-callers are non-deterministically responsible for printing:

- maybe they print only 1 message
- maybe they print no messages
- maybe they print many messages
- printing to slow consoles
- printk-callers in ”inconvenient” contexts

Just avoid printing from time-sensitive contexts, right?

printk and real-time

Decouple printk-callers from printing!

- implement ”any context” storage of messages
  - new lockless ringbuffer + updated crash tools (5.8)
  - NMI-safe LOG_CONT implementation (5.9)
  - remove logbuf_lock (5.12)
  - remove safe buffers (5.15)

- introduce per-console kthreads to print (not mainline yet)

shutdown, boot, panic, kgdb

Printing when kthreads cannot.

- printk-callers can do the printing for:
  - non-panic termination (shutdown, reboot)
  - earlyprintk-kthread window (boot)
  - pr_flush()

- atomic consoles for panics
  - lockless consoles only activated upon panic
  - use the ”printk cpu lock” to synchronize against kthreads and other atomic consoles
  - atomic vs. polling (kgdb) consoles?

- What about kdb/kgdb?
  - Should kgdb be able to step through cpu locks?
  - Daniel Thompson suggested allowing transfer of ownership when all other CPUs quiesced.

PREEMPT_RT vs. mainline

PREEMPT_RT (now)	mainline (planned)
1. atomic consoles print in pre-kthread boot window	printk-caller non-atomic printing in pre-kthread boot window
2. pr_flush() simply waits	pr_flush() non-atomic printing
3. pr_flush() on non-panic termination	pr_flush() and printk-caller non-atomic printing on non-panic termination
4. only atomic consoles on panic	only atomic consoles on panic until stack trace printed, then also best effort pr_flush() and printk-caller non-atomic printing
5. new atomic console API	re-use polling API ???
6. kgdb CPU requeues if holding cpu lock	??? transfer lock ownership?
7. kgdb only uses atomic consoles	??? keep polling API?