

The Status of the Preempt-RT Patch

Embedded Linux Conference Europe 2017

Sebastian A. Siewior

Linutronix GmbH

October 24, 2017

① Overview

② The patchqueue

③ Summary

1 Overview

2 The patchqueue

3 Summary

What is PREEMPT-RT

- ❏ A patchset provided for certain kernels
- ❏ Provides realtime functionality
- ❏ Preempts softirq, spinlocks, IRQ-off regions
- ❏ “Is super deterministic, not super fast”

PREEMPT-RT


- ❏ **Project was without funds for a few years**
- ❏ **LF announced RTL Collaborative Project in October 2015**
- ❏ **Documentation**
<https://rt.wiki.kernel.org>
<https://wiki.linuxfoundation.org/realtime/start>
- ❏ **Testing**
<https://ci-rt.linutronix.de>

① Overview


② **The patchqueue**

③ Summary


v3.10




v3.12

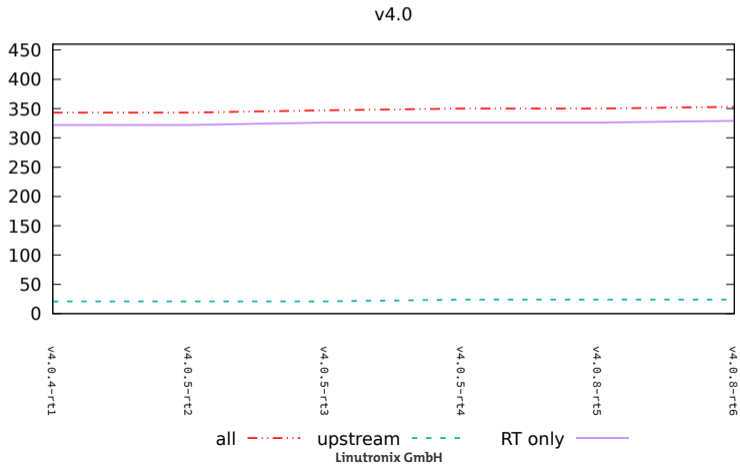


v3.14




v3.18

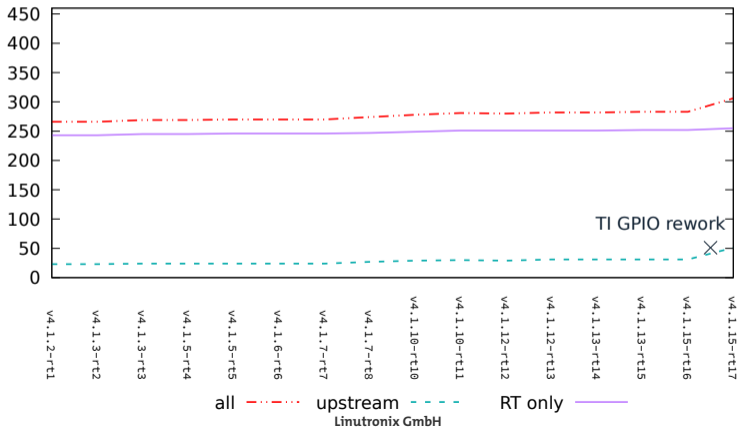





v4.1



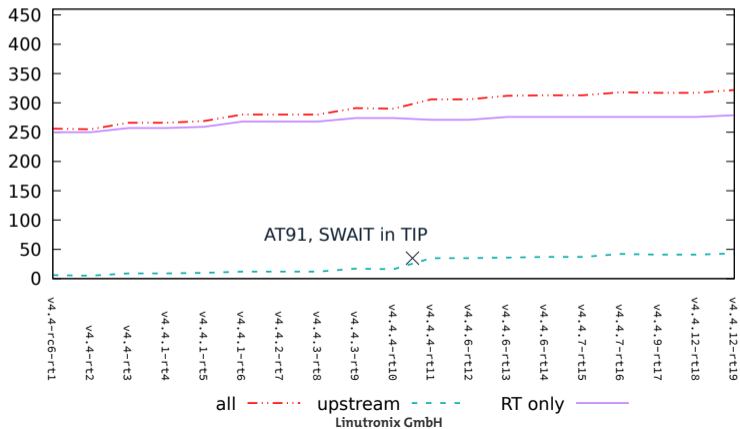
v4.1



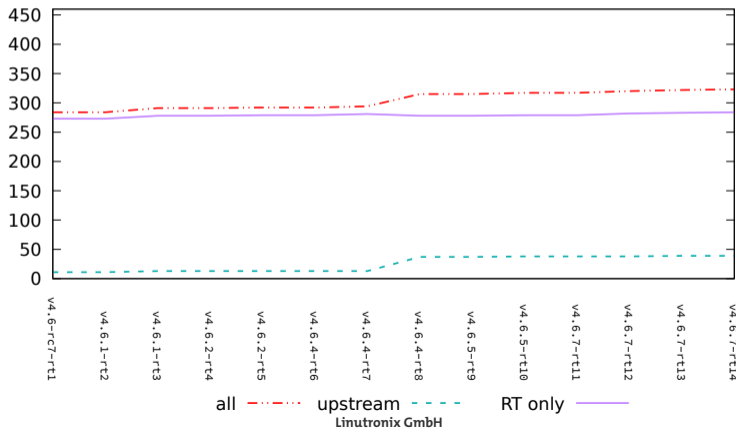
v4.4

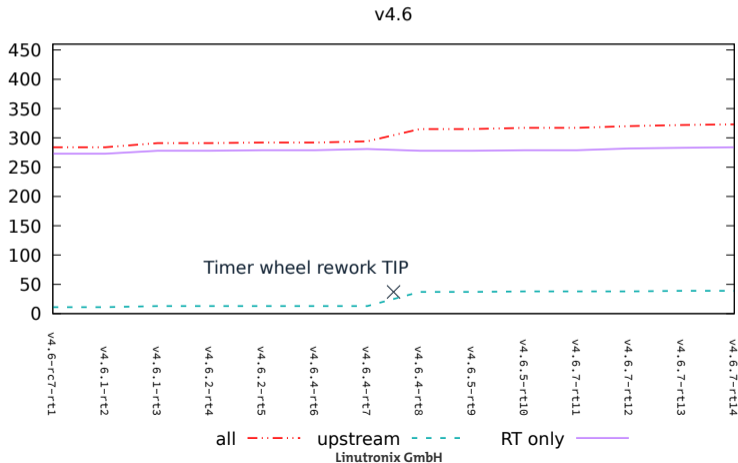


v4.4




v4.6






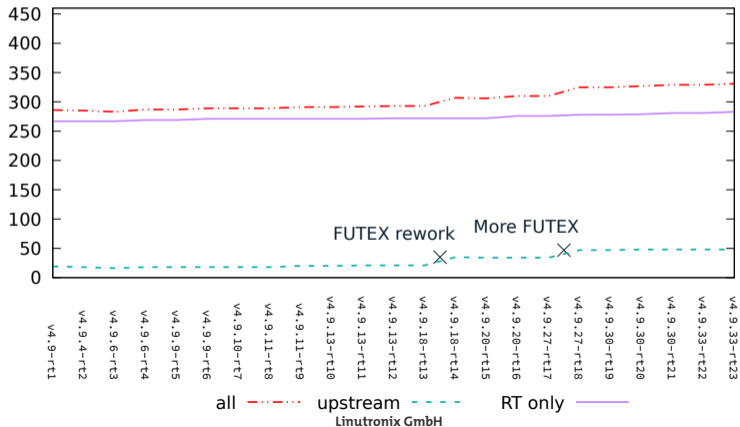
v4.8




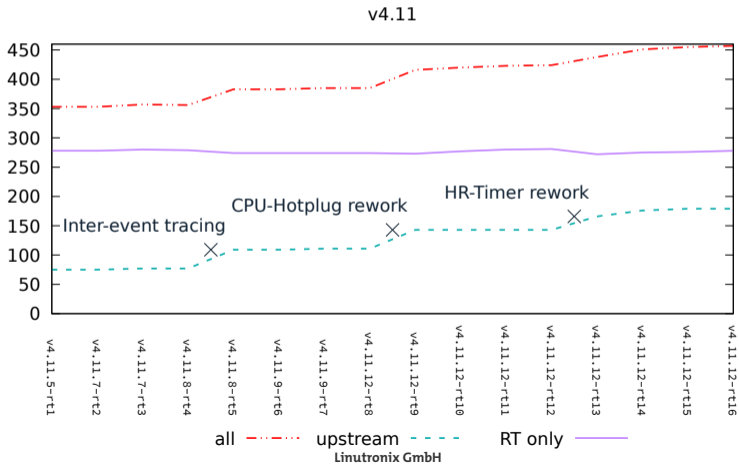
v4.9



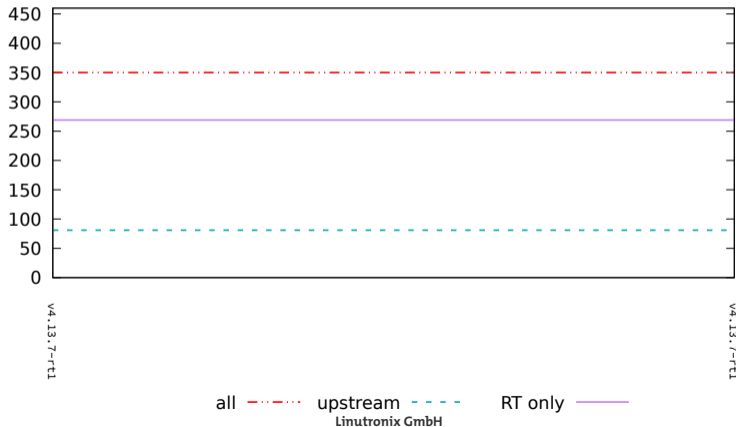
v4.9









v4.13



Summary



Summary



① Overview

② The patchqueue

③ Summary

The focus during development

- ❏ v4.4 to v4.9 RT was mostly quiet
- ❏ Focus on upstream and CPU hotplug
- ❏ A few attempts were made in v4.9, finally in v4.11
- ❏ The strategy changed a few times within RT

What RT version gets released

- ☞ Starting with v4.4 every other version
- ☞ Once Greg-KH decides to make a different LTS, we follow that
- ☞ So we got RT for v4.8 and v4.9
- ☞ we continue every other version until Greg-KH ...
- ☞ ...v4.13 and v4.14

What RT versions are maintained

- 📁 Basically every supported upstream kernel
- 📁 Steven Rostedt maintains most of them
- 📁 Julia Cartwright maintains the v4.1 tree
- 📁 "Features" are hardly backported, mostly fixes

Major pieces

- ❏ Page fault disable by s/390
- ❏ timer wheel rework, made NO_HZ_FULL work
- ❏ CPU hotplug rework, started around v4.1
- ❏ hrtimer rework (pending).
- ❏ Tom Zanussi's "tracing: Inter-event (e.g. latency) support" (pending).

Tiny pieces

- ❏ FUTEX rework. Closes PI-deboost problems, unbreaks SCHED_DEADLINE
- ❏ RW-Semaphore rework (unbreaks the radeon driver)
- ❏ RW-Lock rework. Helps CPU hotplug
- ❏ Fix for a bug exploited by GDB
- ❏ A lot driver and subsystem across tree.

Outstanding pieces

- ❏ MM bits. Shorter atomic sections in buddy / SLUB
- ❏ migrate disable
- ❏ local lock
- ❏ network bits
- ❏ printk, serial drivers
- ❏ signal delivery
- ❏ bit spinlocks / list_bl
- ❏ RCU

Outstanding pieces

- ❏ simple wait queues
- ❏ dcache try loops and `cpu_chill()`
- ❏ workqueue locking
- ❏ crypto, FPU sections
- ❏ sleeping spinlocks
- ❏ softirq bits
- ❏ various driver patches

Required patches vs nice to have

- ☒ **Required features**
 - Sleeping spin locks
- ☒ **Nice to have**
 - lazy preempt support
- ☒ **Should have but...**
 - preempt disabling handling, WARN_ON

Thank you for your attention

Contact

**Linutronix GmbH
Sebastian A. Siewior**

**Auf dem Berg 3
88690 Uhltingen
Germany**

eMail bigeasy@linutronix.de