Red Hat
Enterprise Linux

Real-time Linux: What is, what is not and what is next.

Daniel Bristot de Oliveira
Principal Software Engineer, Red Hat
Researcher, Retis Lab – Scuola Superiore Sant'Anna

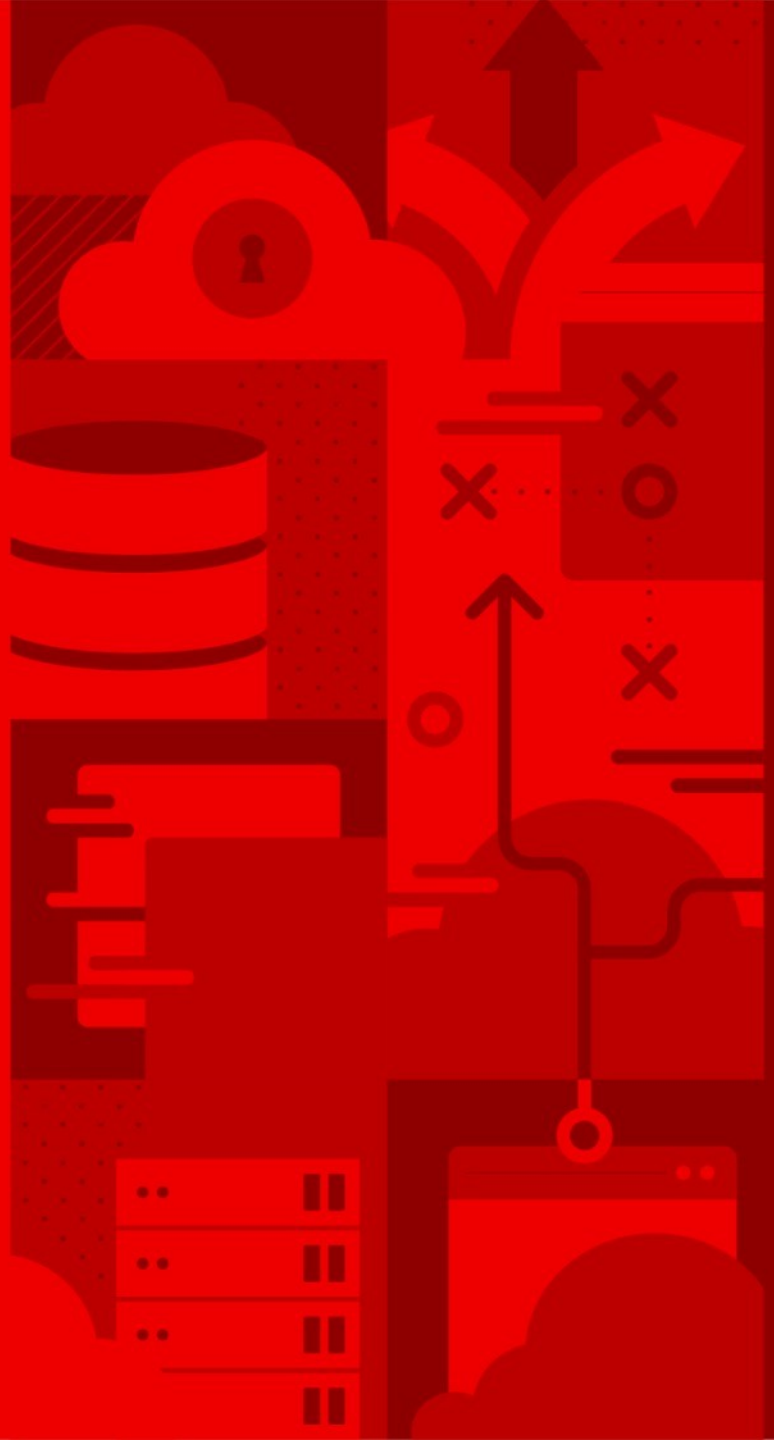

PREEMPT_RT is mainline!

Red Hat

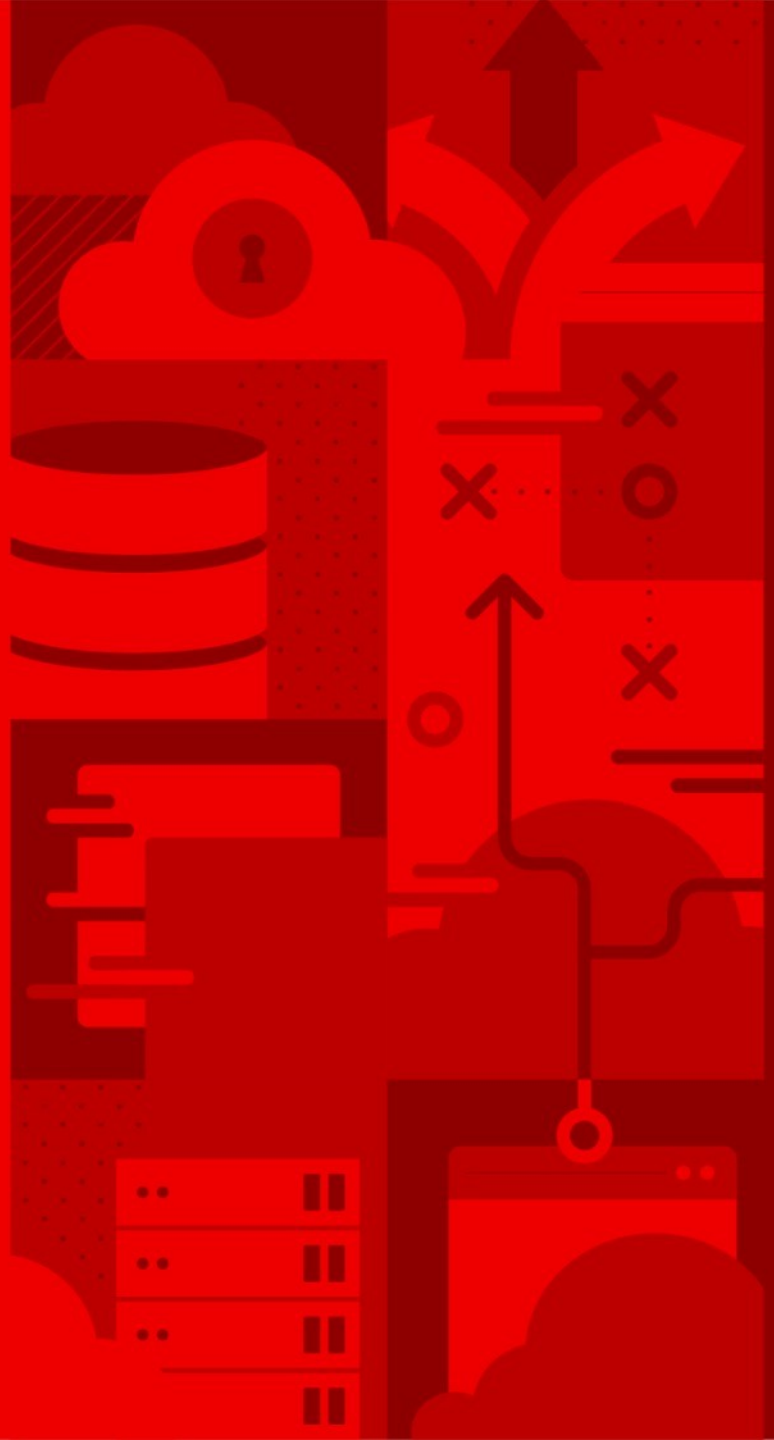

So... Real-time Linux is *done*?

Red Hat

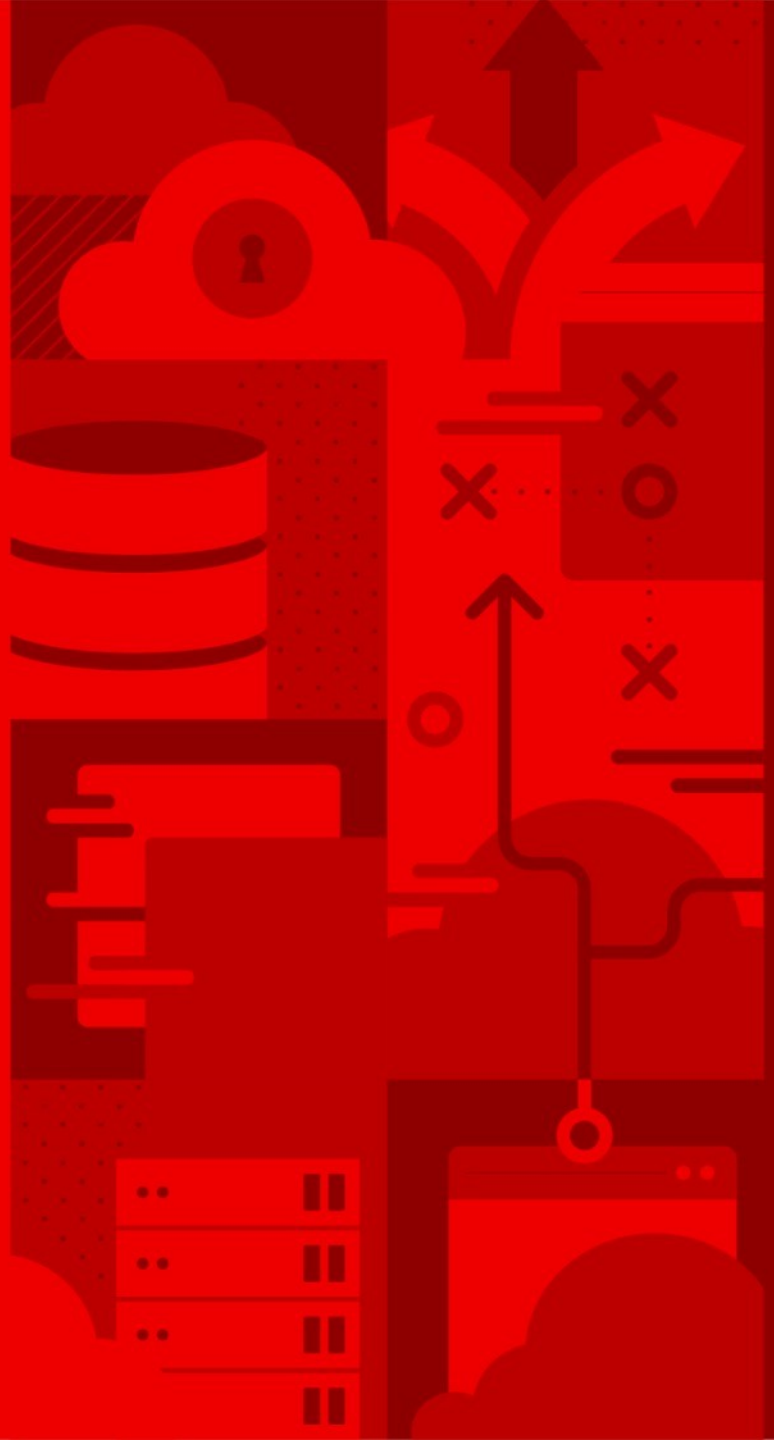

Let’s discuss…

Red Hat

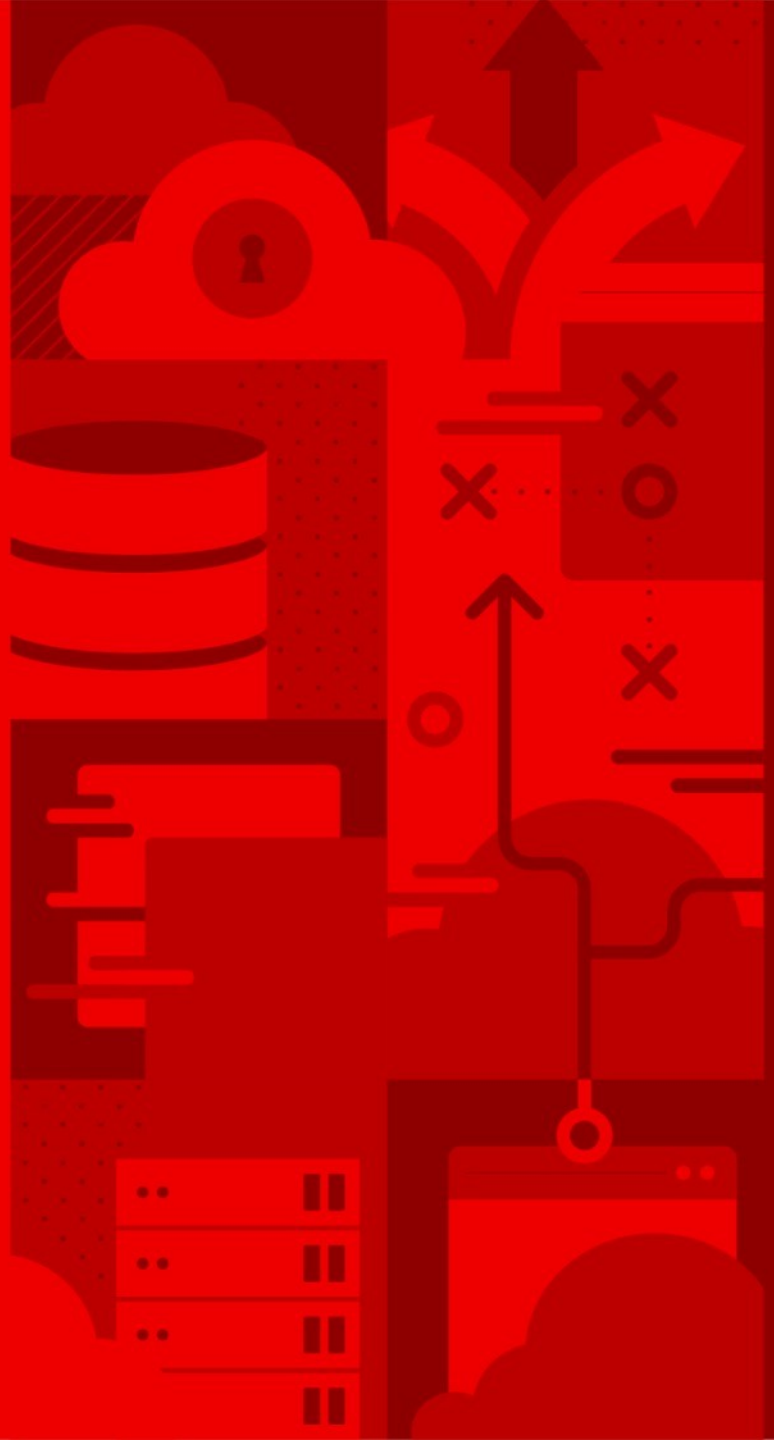

But before start...

Red Hat

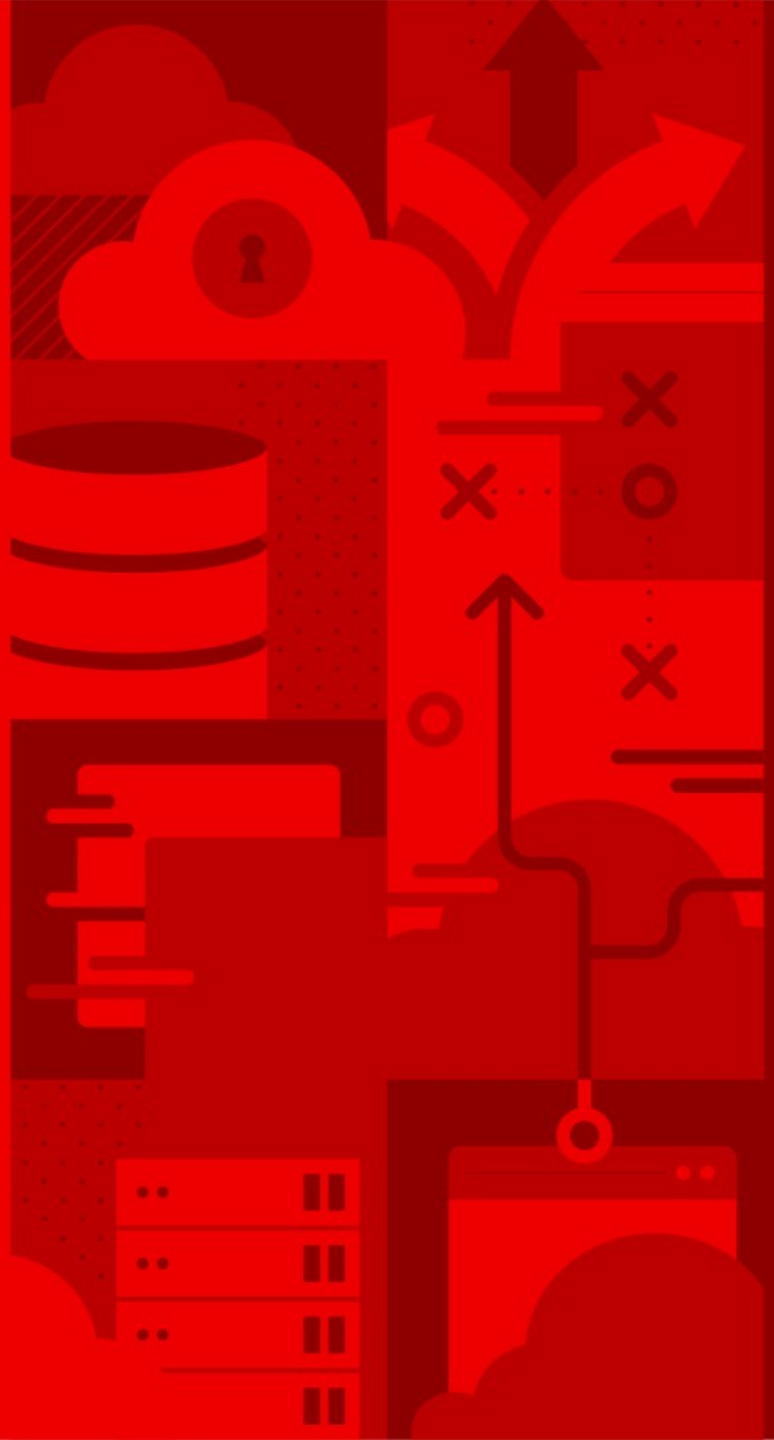

What is?

Red Hat

Non-rt Timeline

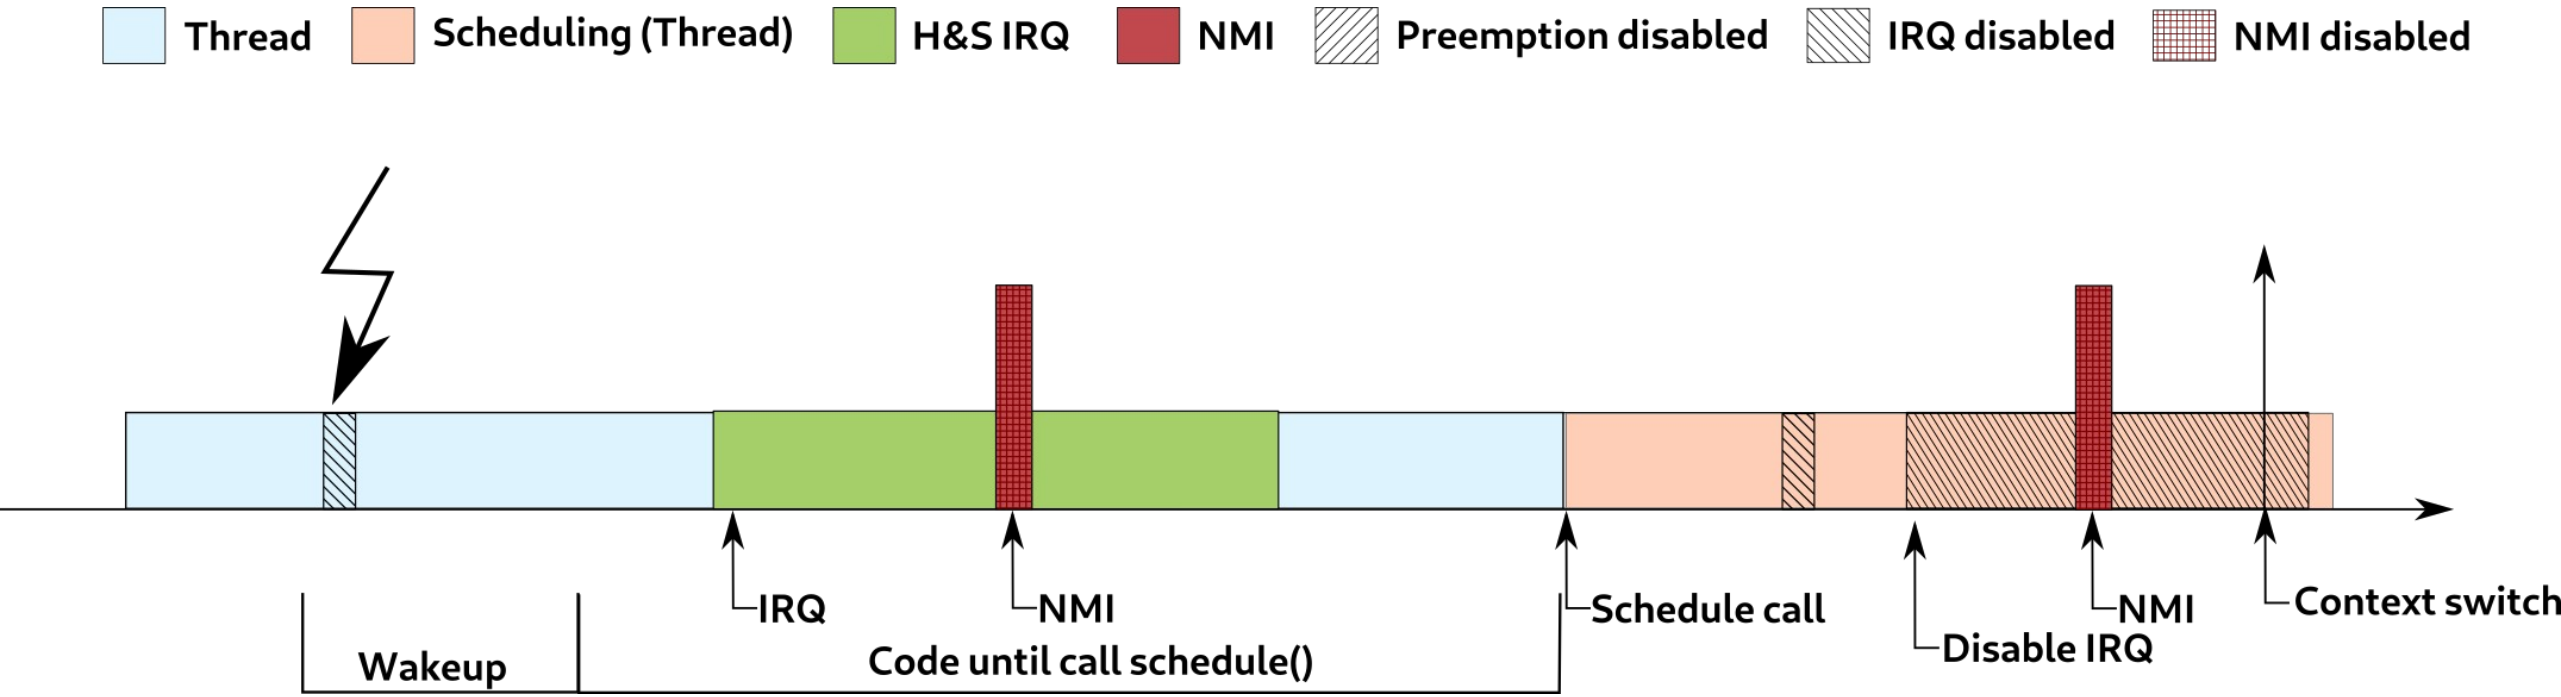

Red Hat

Non-rt Timeline + thread IRQ

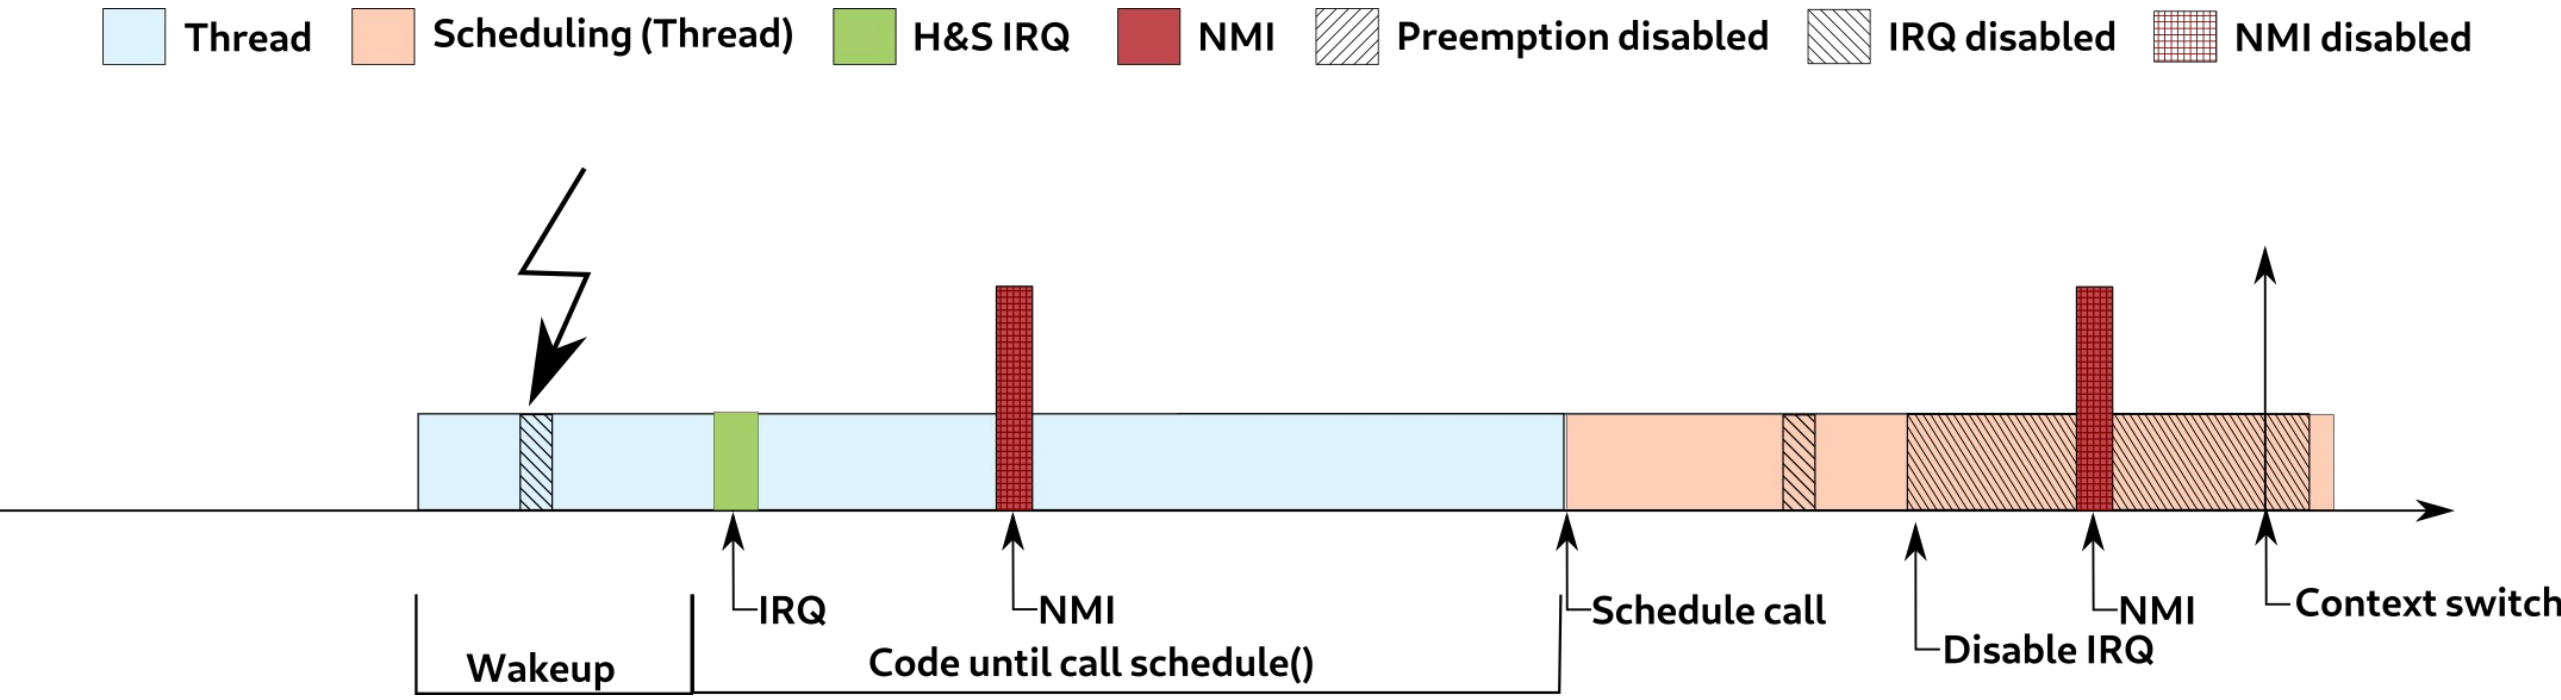

Red Hat

Non-preempt -> preempt

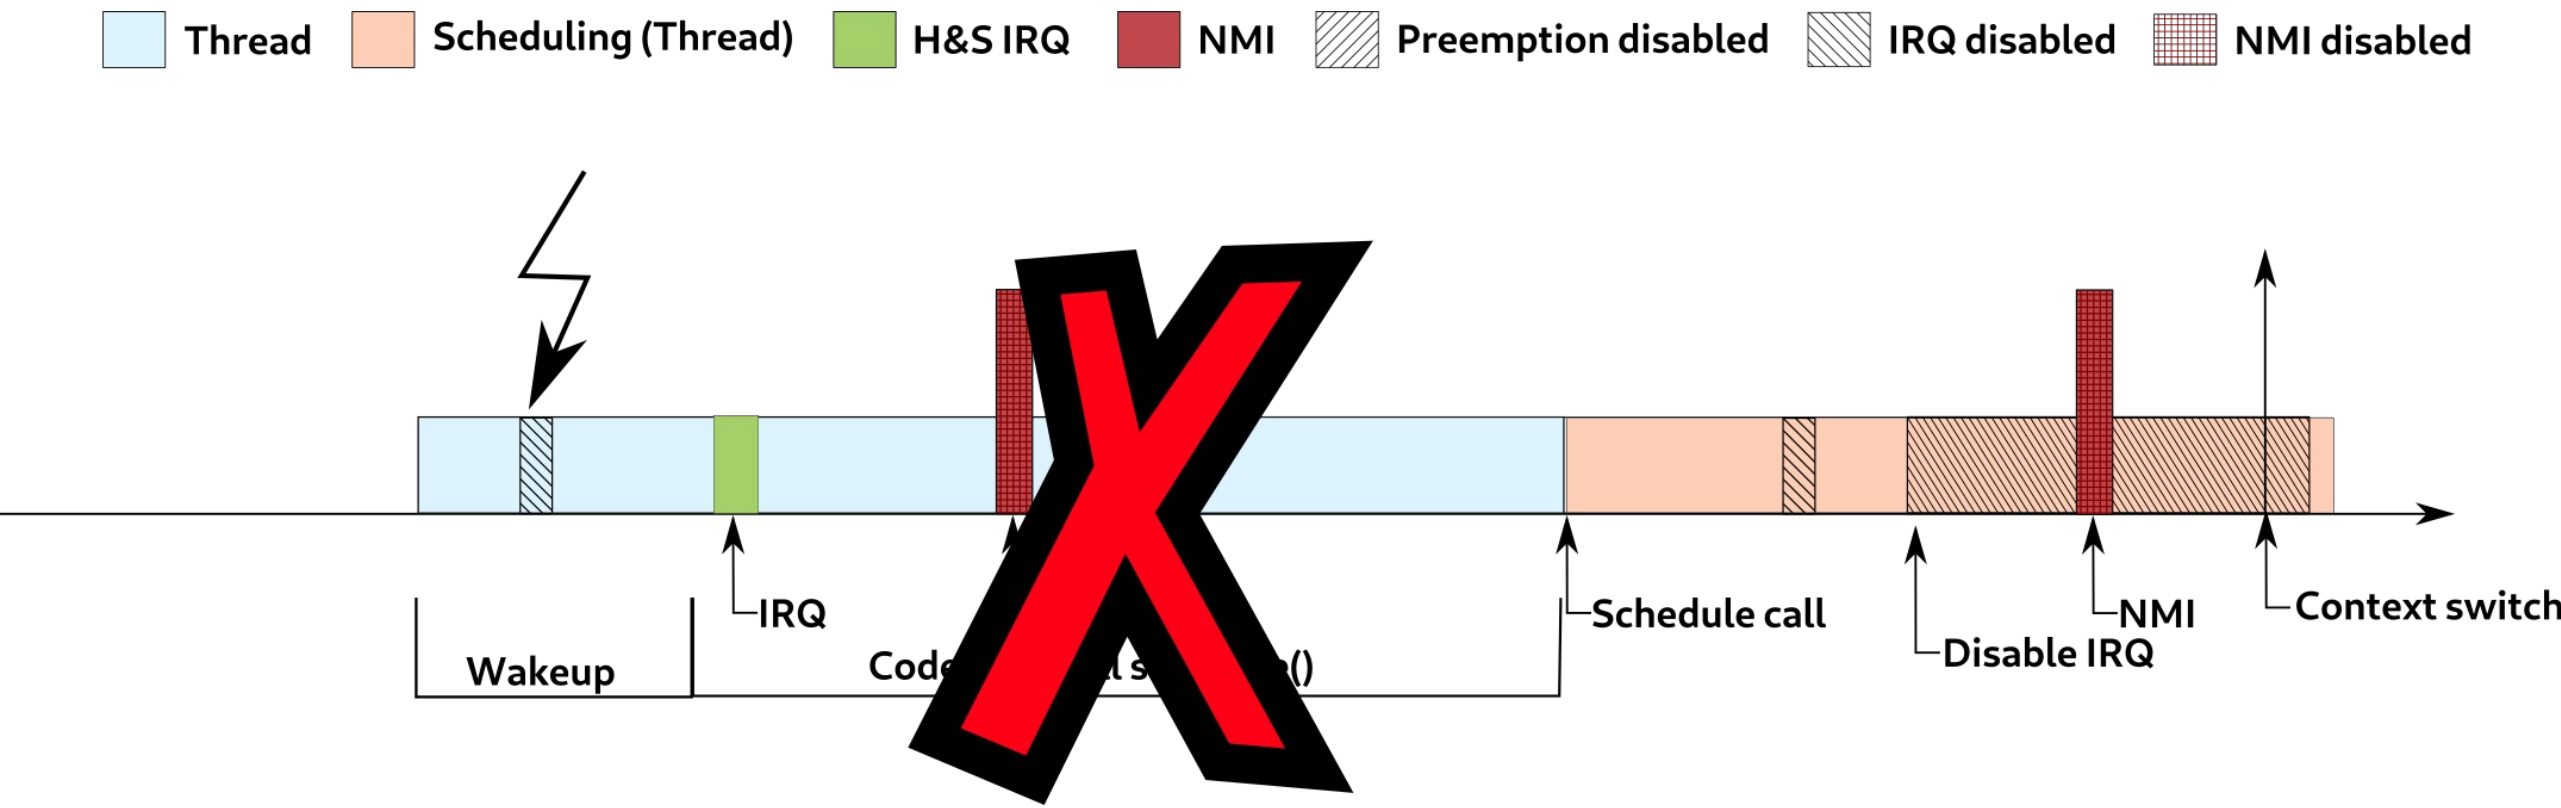

Red Hat

Non-preempt -> preempt

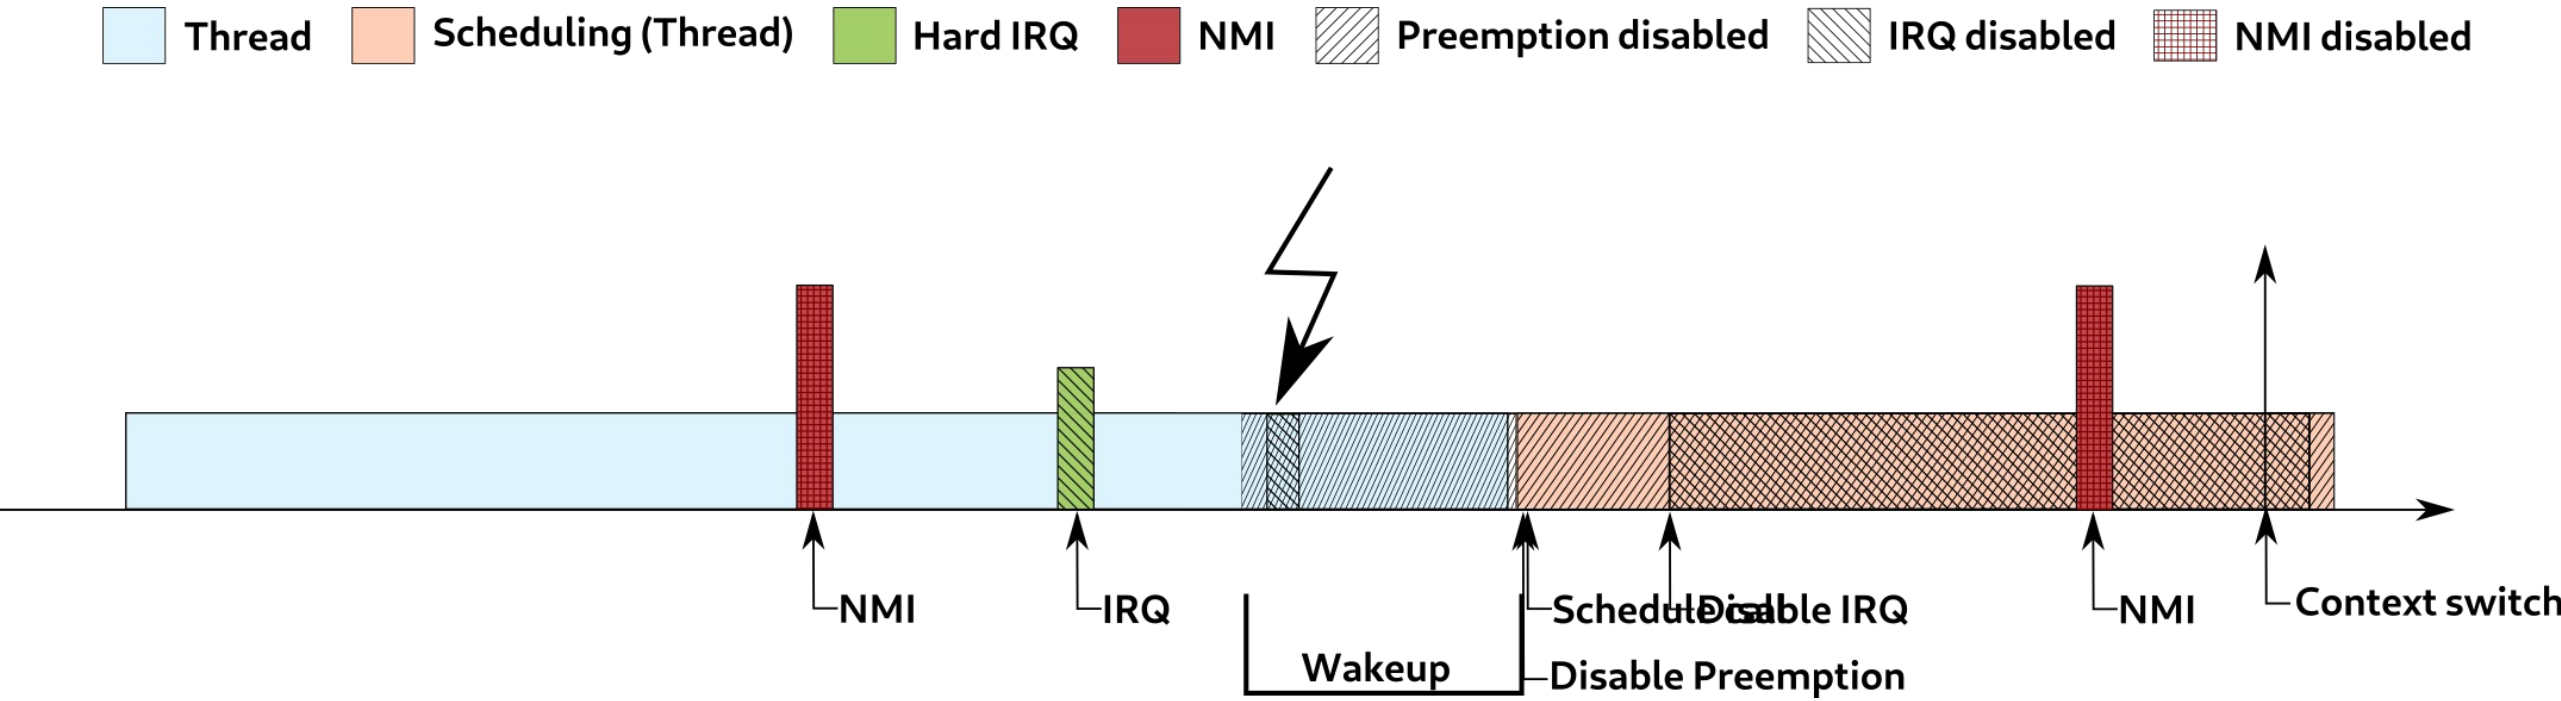

Latency is the output!

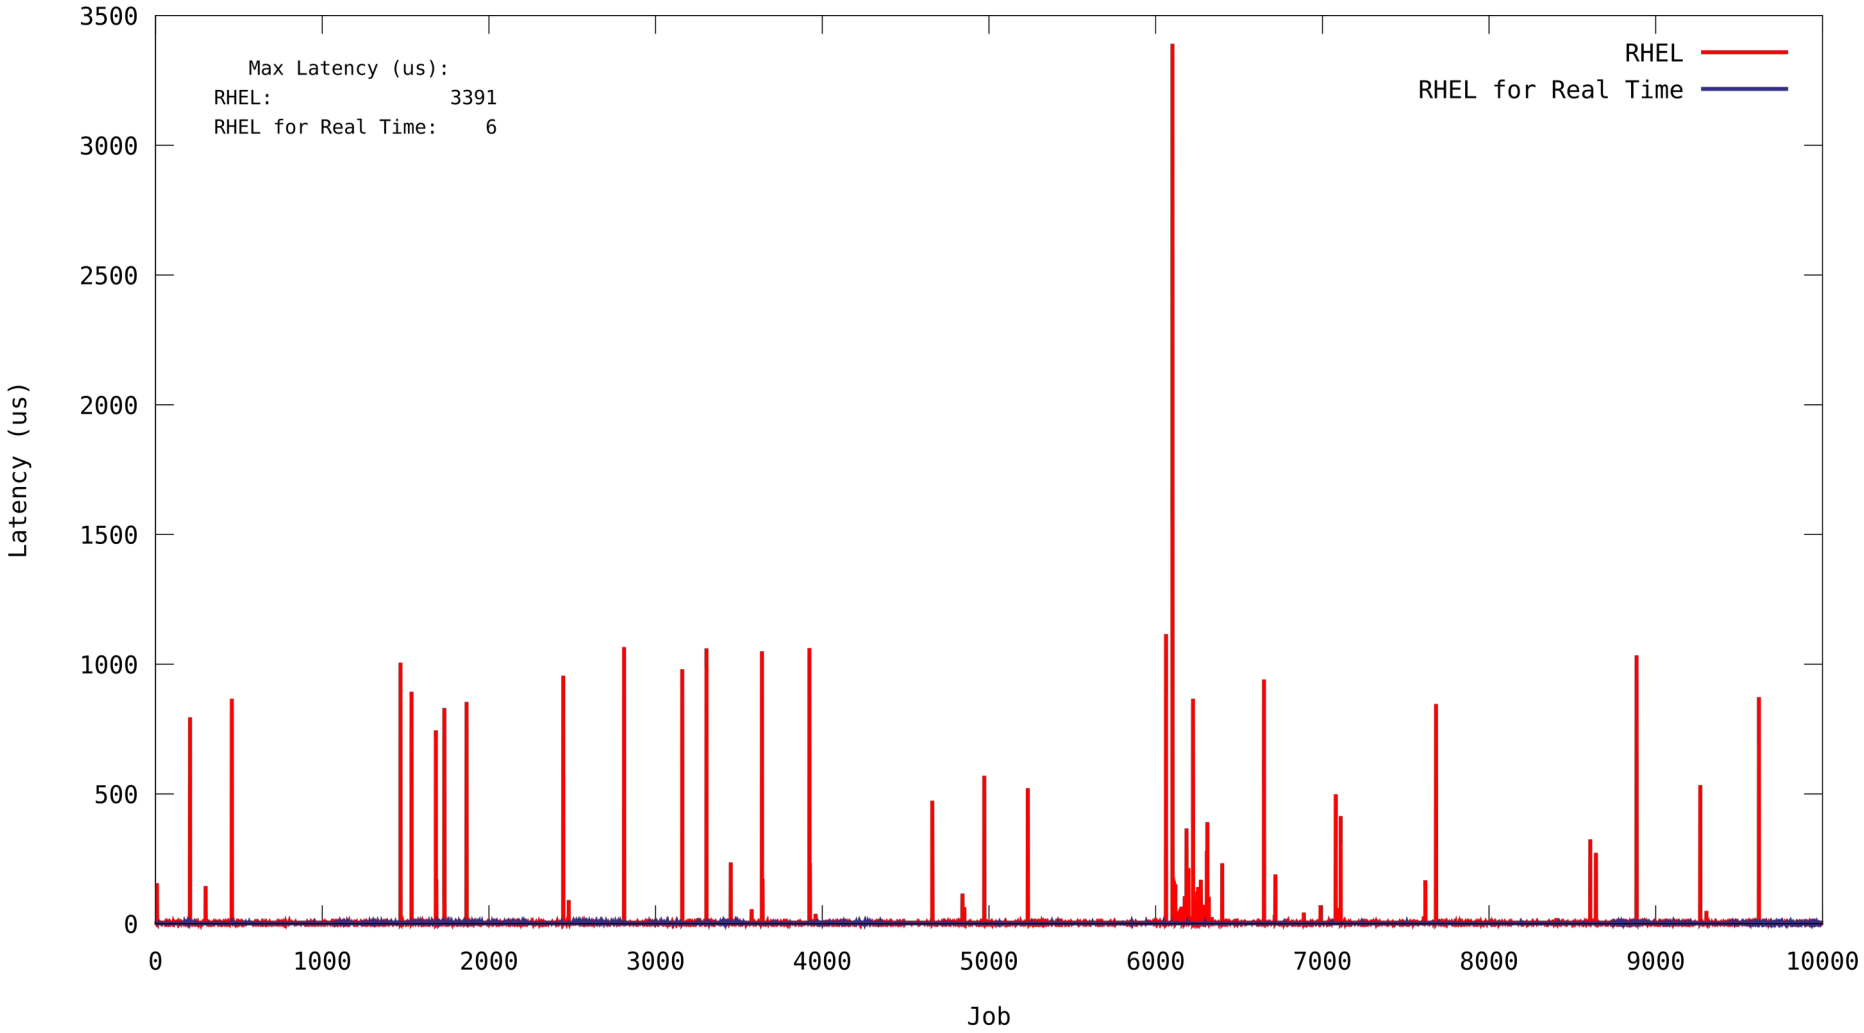

Red Hat

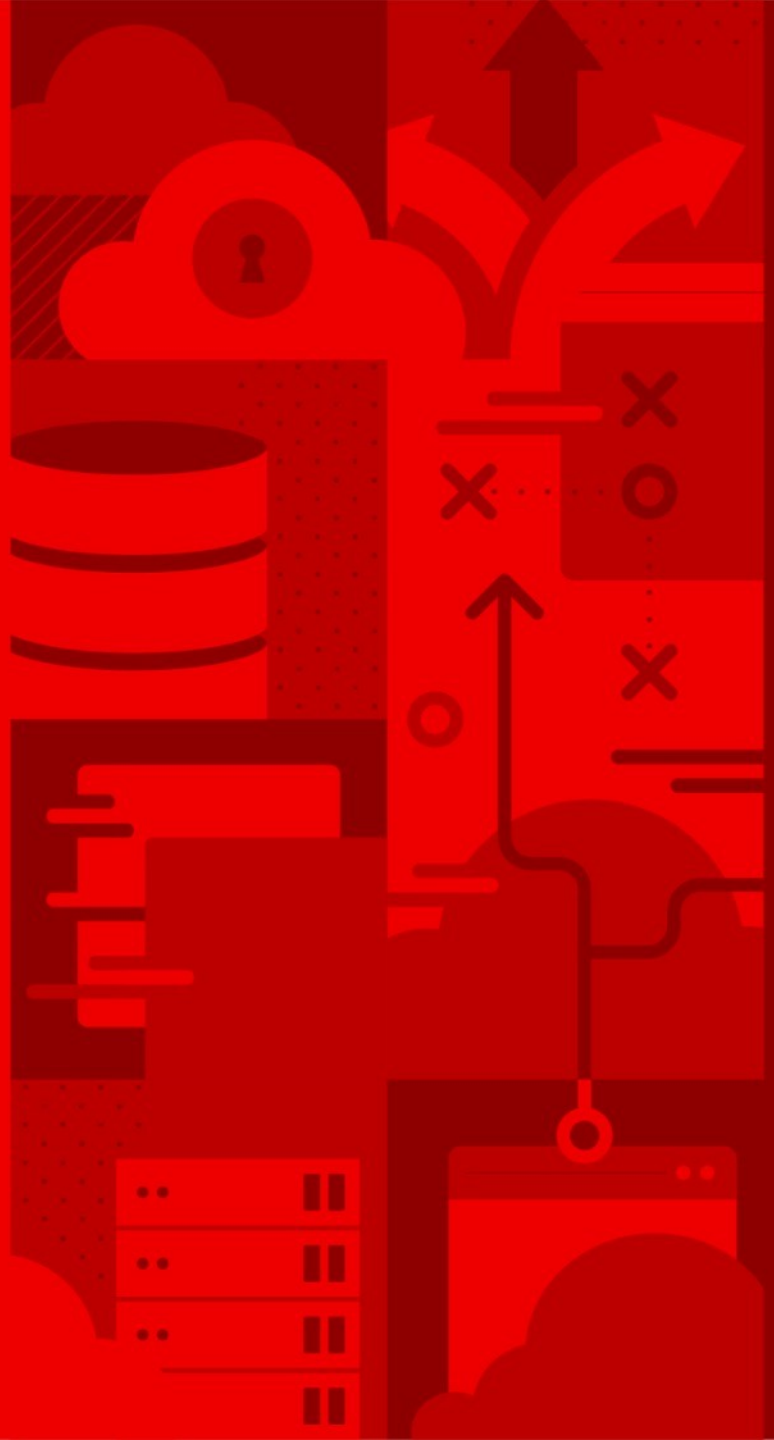

That is good!

Red Hat

In words

- IRQs become a scheduling problem
- The Preemptive mode:
  - Other than reducing the latency
  - Brings Linux closer to the theoretical models used on schedulers

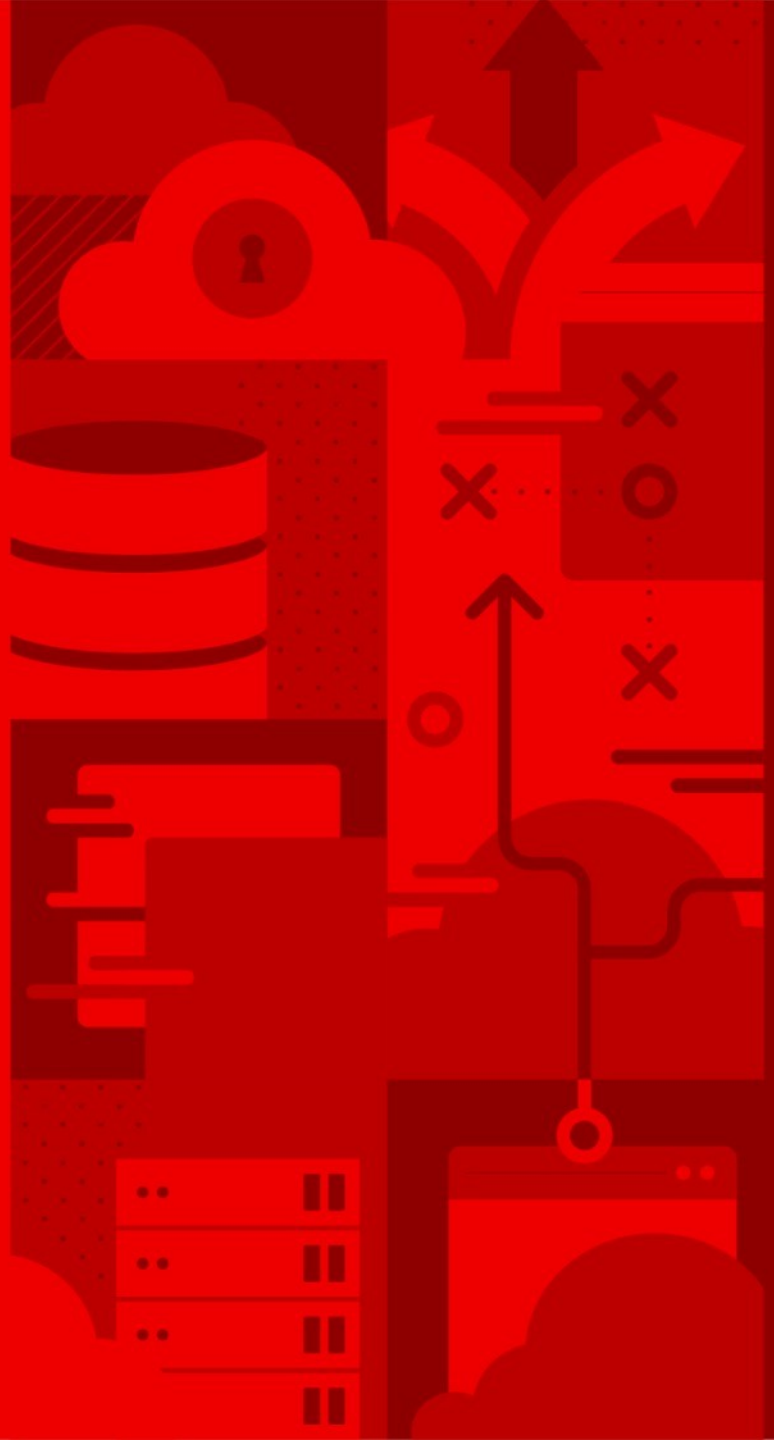

Another good thing:

Red Hat

RM | DL | Locking

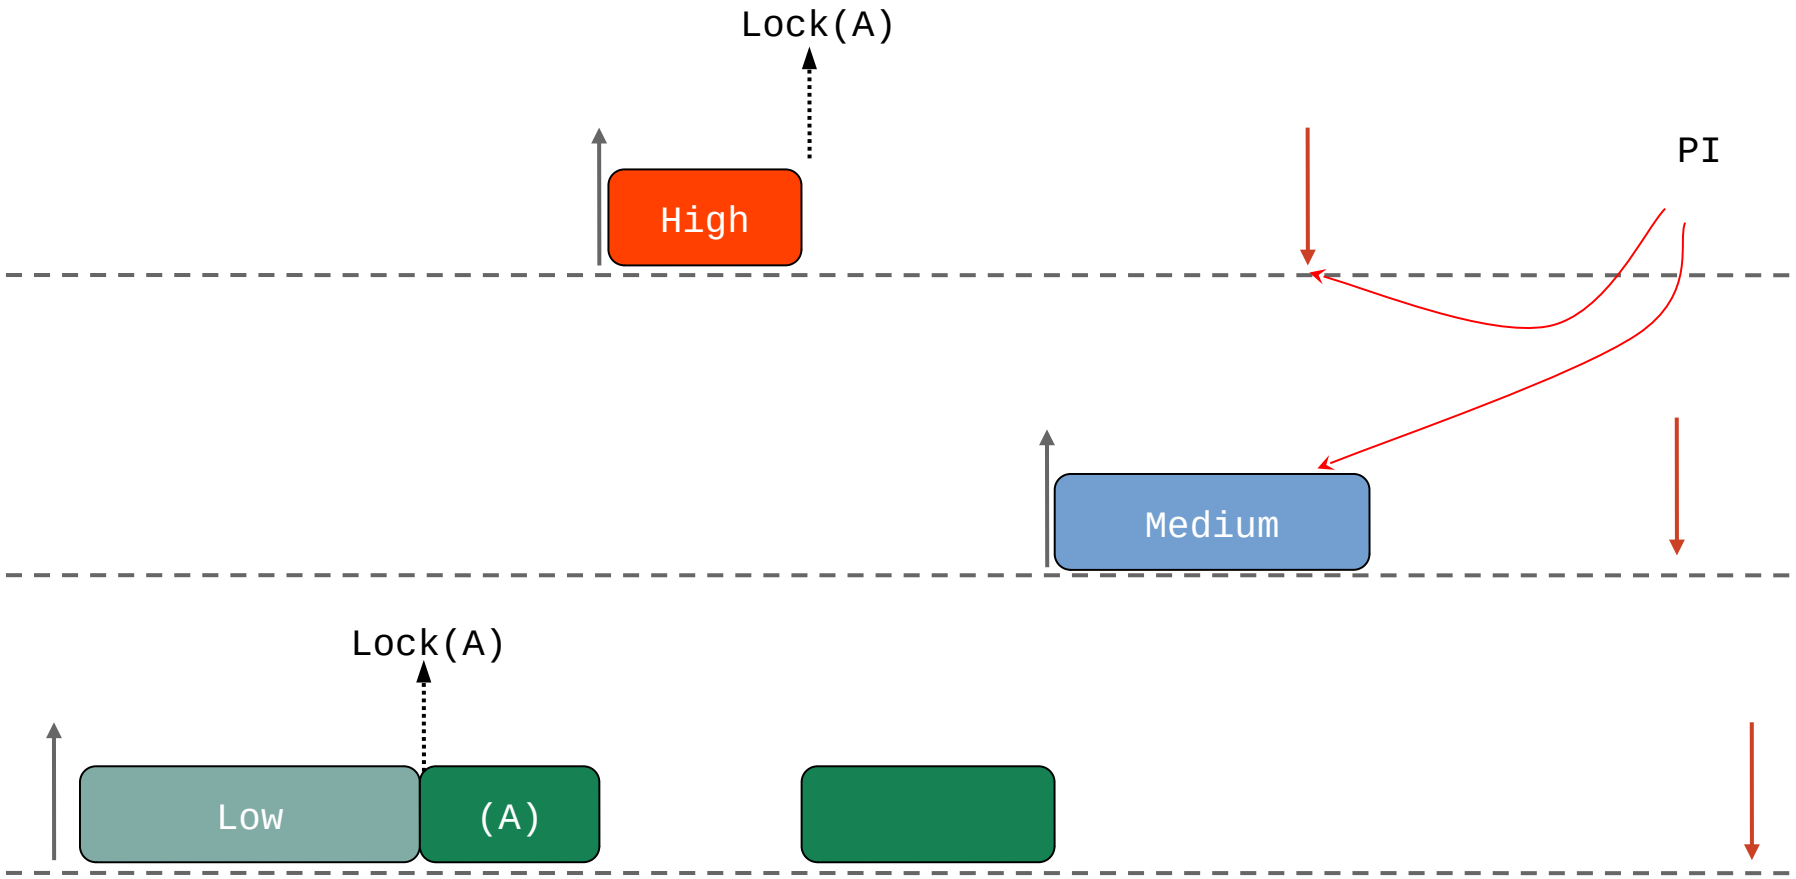

Red Hat

RM | DL | Locking

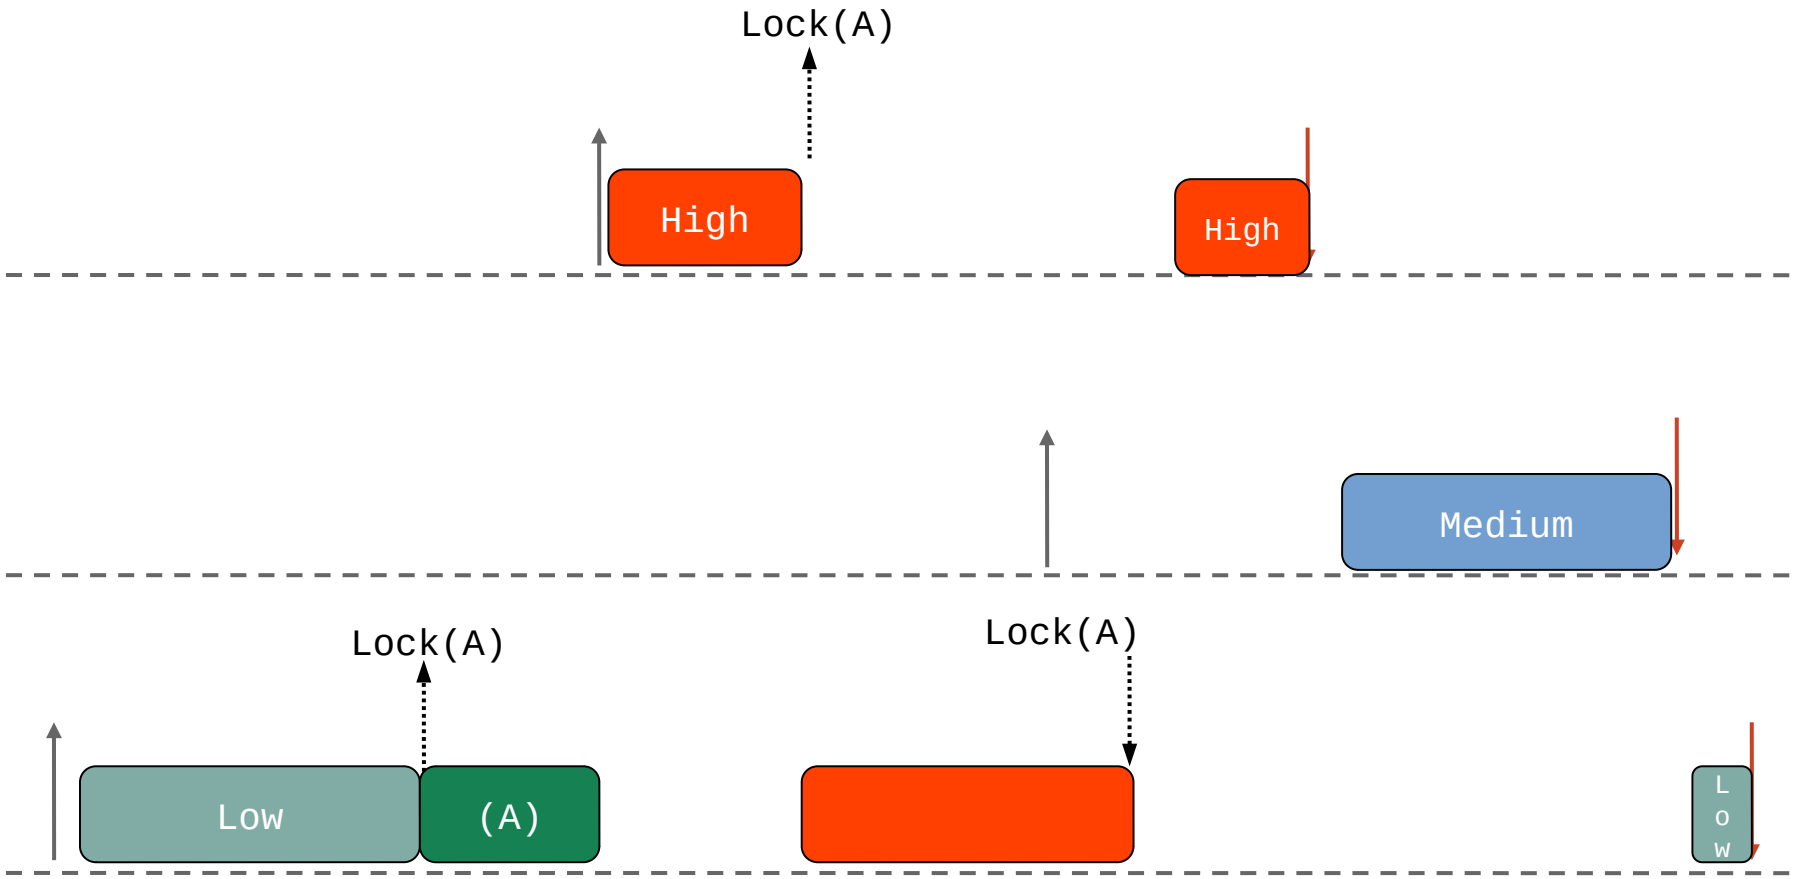

Red Hat

In words

- Informally, the **response time** of a task depends:
  - A little on the IRQ
  - A little on the Latency/scheduling
  - On the locks the thread depends on
  - On it's own execution time
  - On the interference of higher priority threads (which is OK)

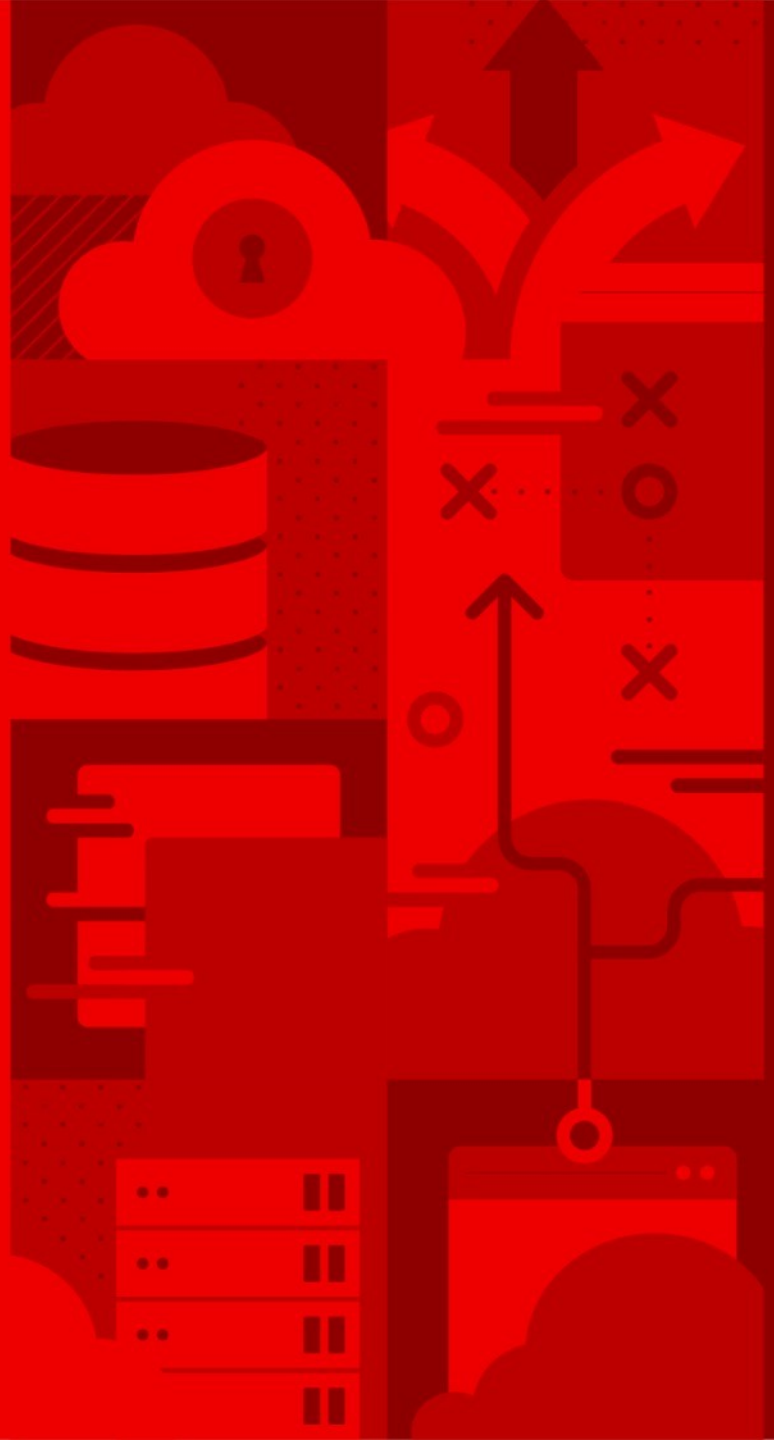

Empirically, turning Linux a possible RTOS!

Red Hat

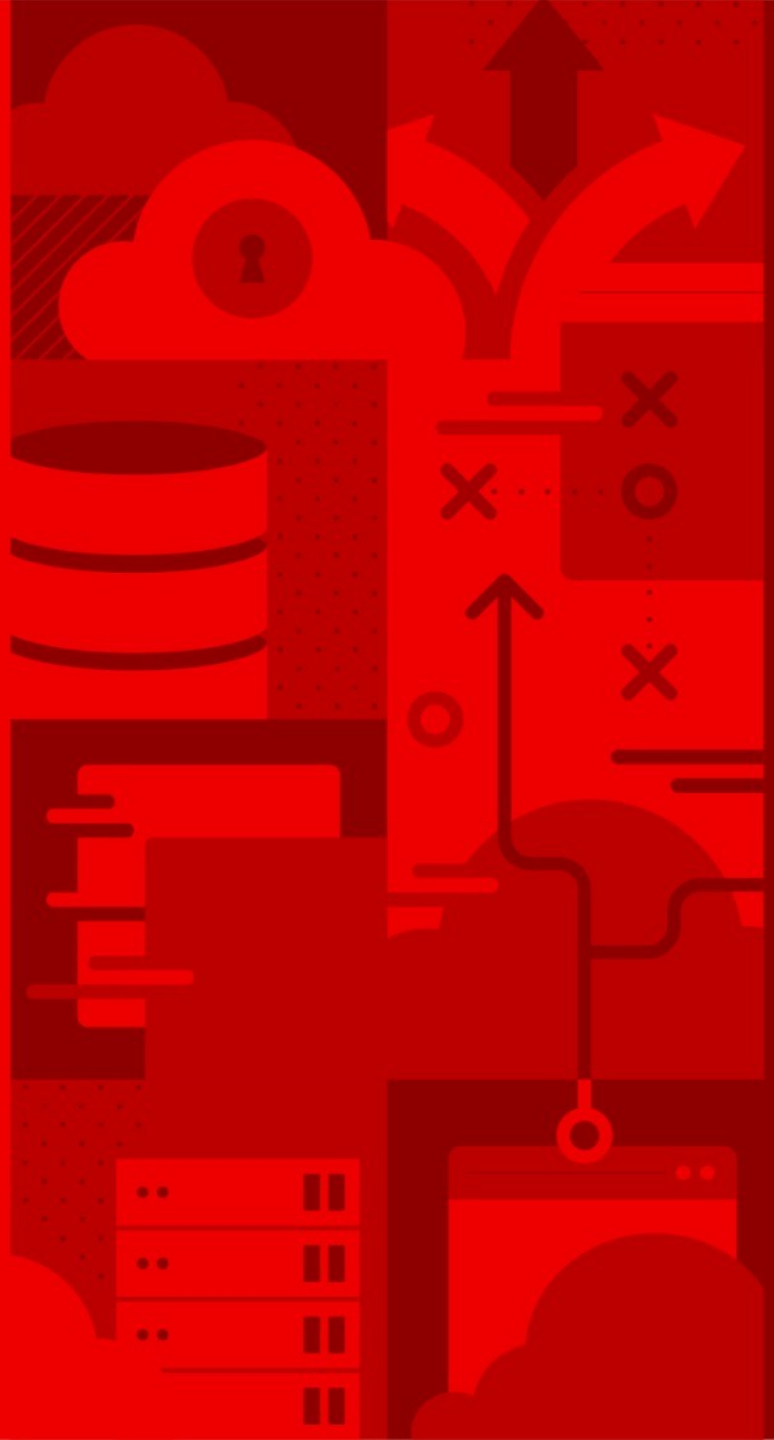

But can we say we are a Hard RTOS?

Red Hat

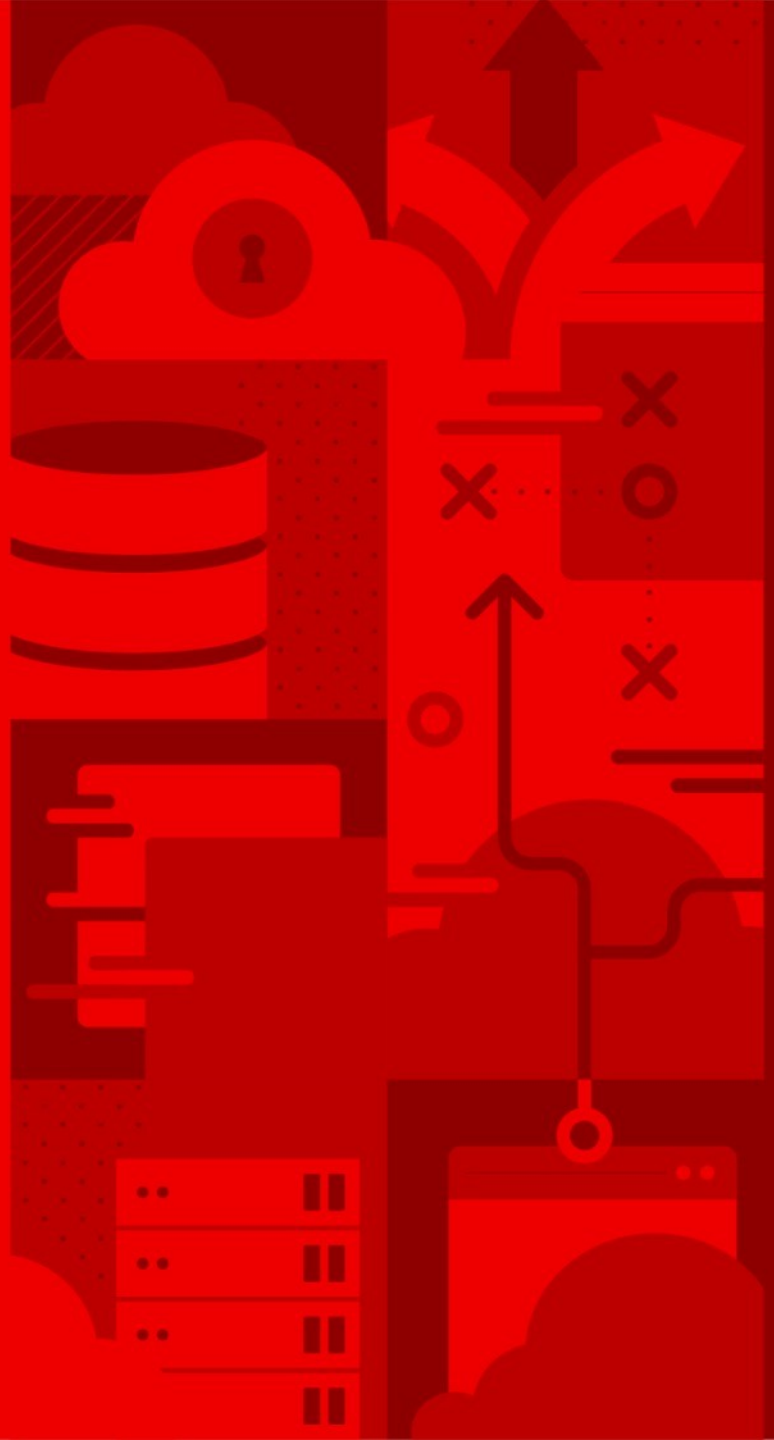

Let’s do an exercise!

Red Hat

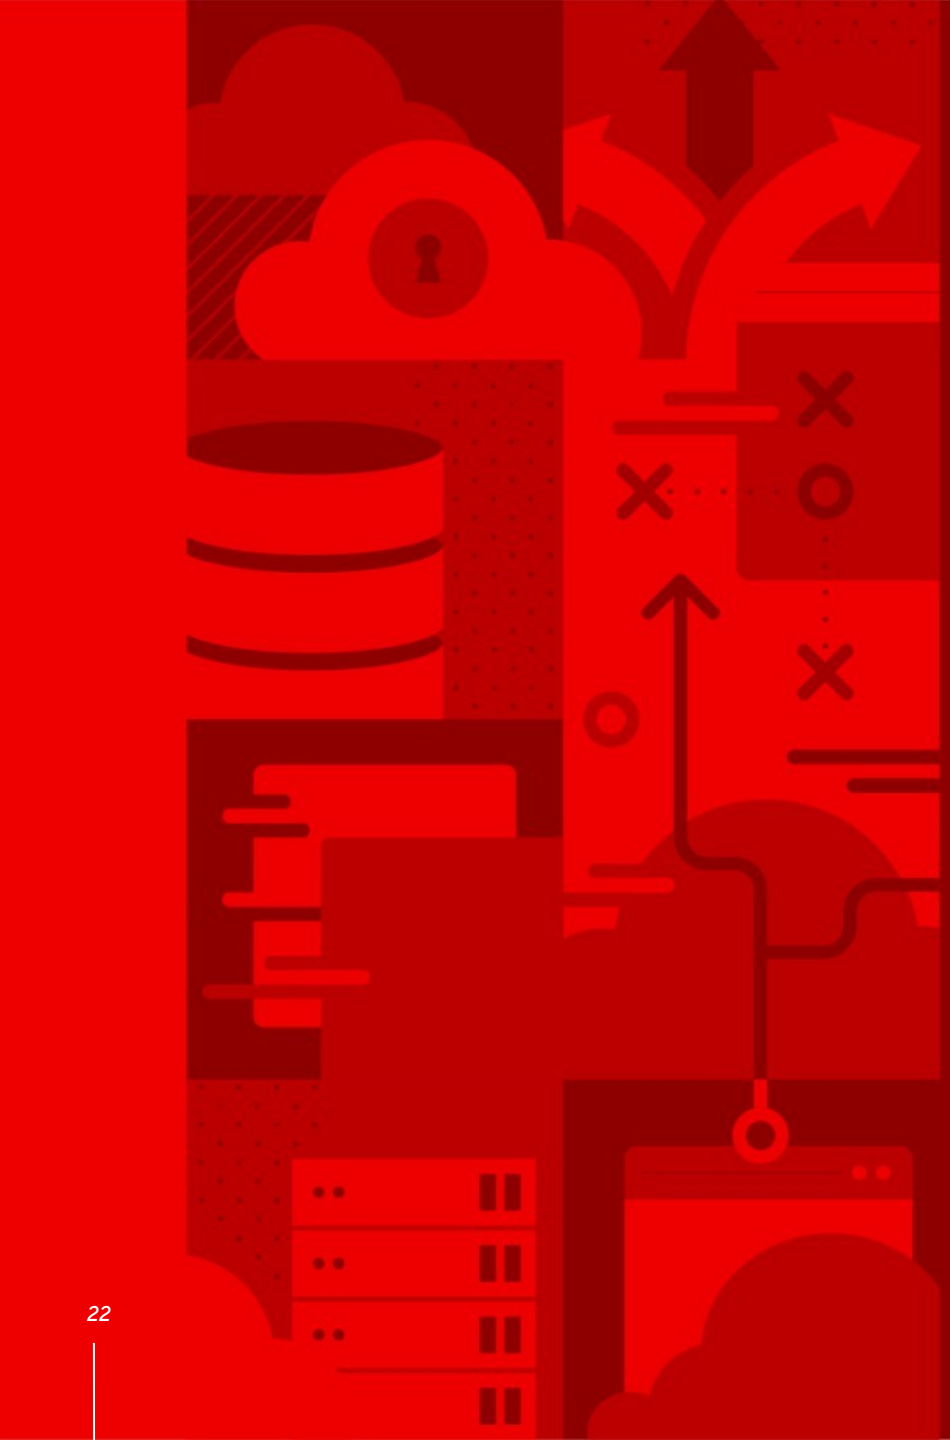

No deadline miss

Red Hat

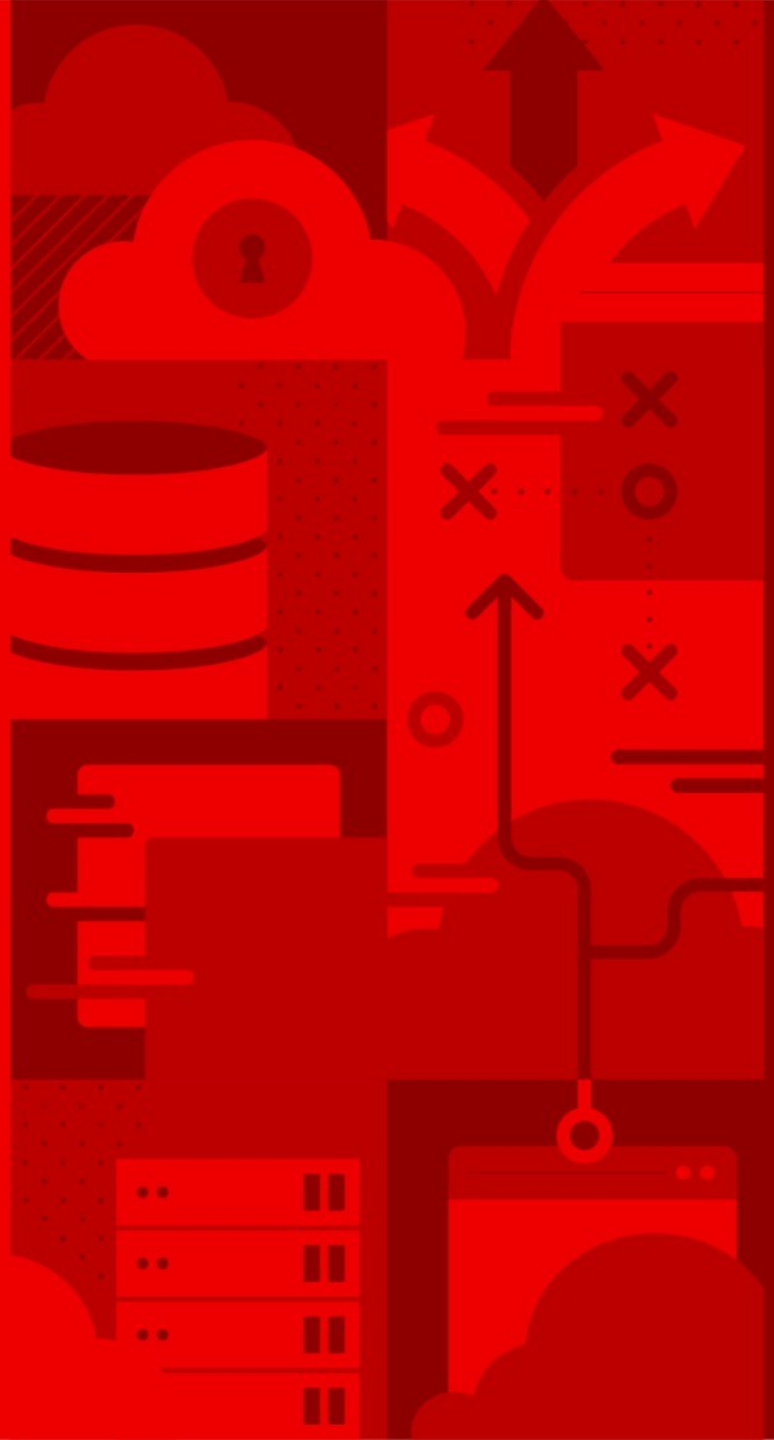

Mean knowing worst case scenarios

Red Hat

Worst Case Execution Time

- There are two classical methods
  - Static methods
    - The code is not executed
    - The control flows are analyzed
      - Execution time is “computed” based on model of the hardware.
  - Measurement methods:
    - Measures the code running in the real hardware.

We do measurements!

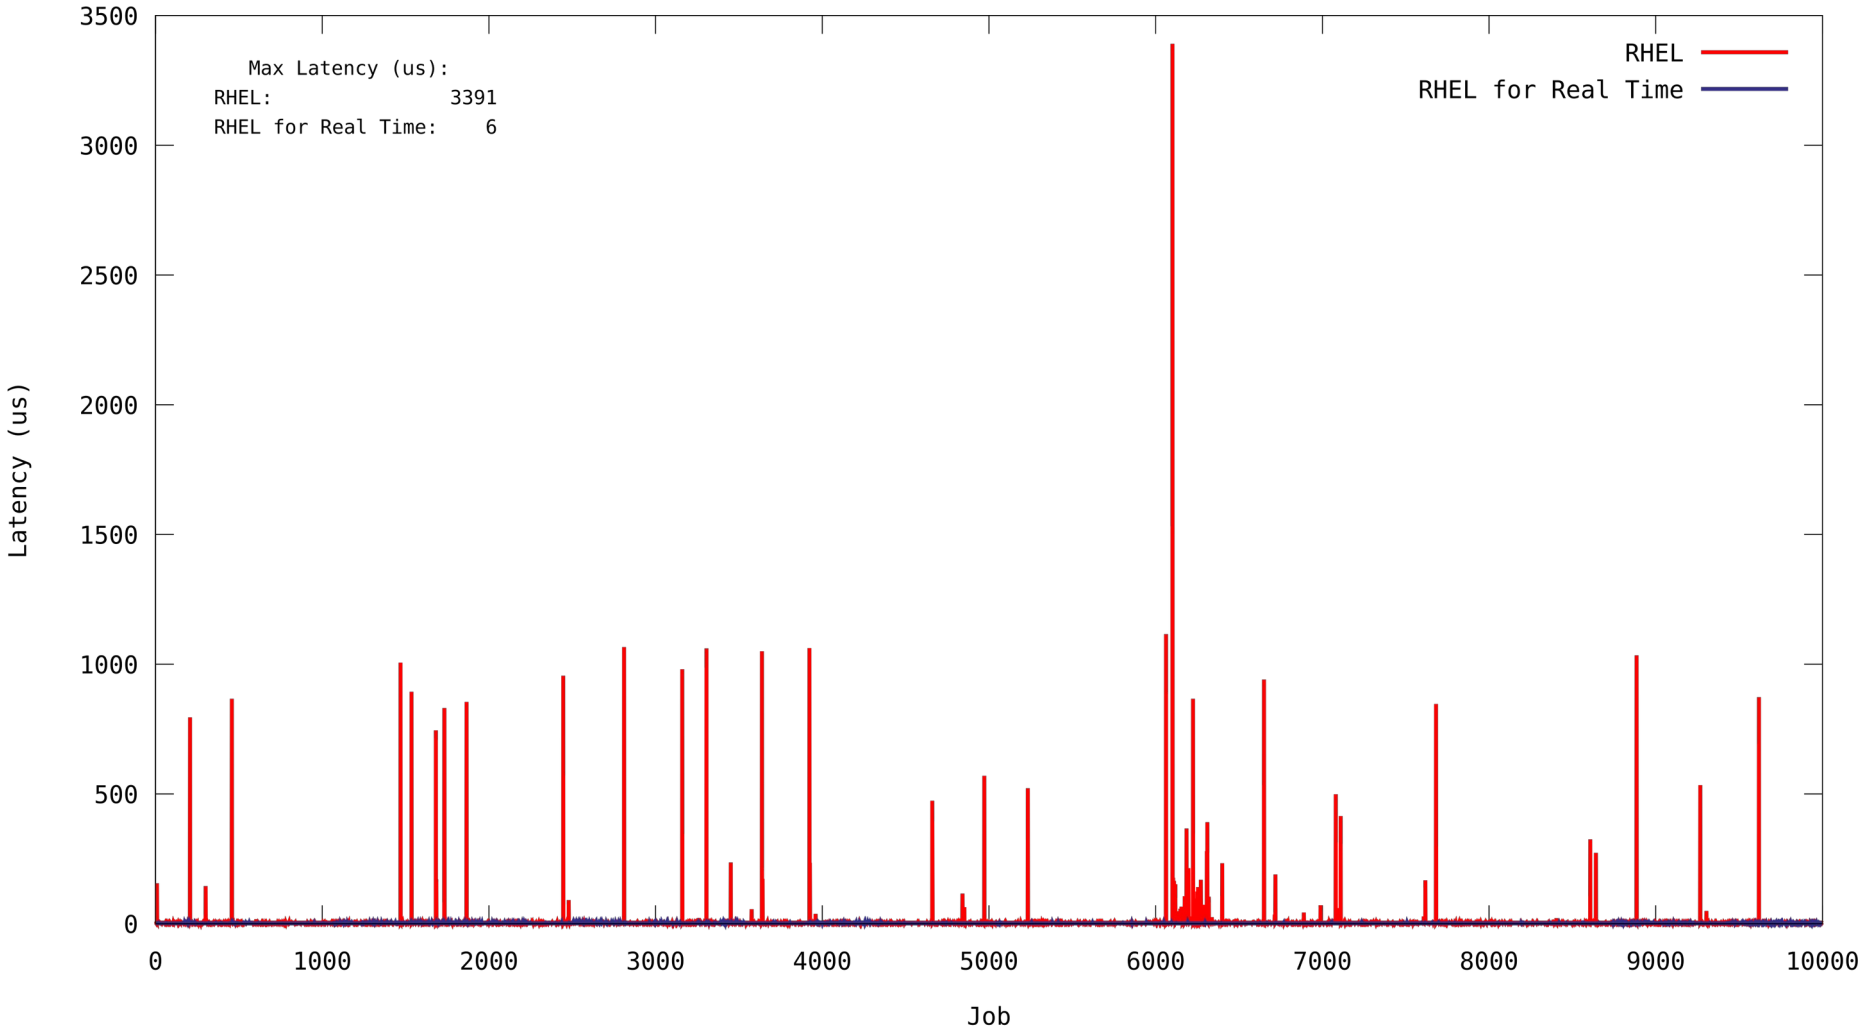

Red Hat

For example:

- DO-178C, **Software Considerations in Airborne Systems and Equipment Certification:**
  - *"Timing measurements by themselves cannot be used without an analysis demonstrating that the worst-case timing would be achieved*. In other words, testing alone is not adequate for demonstrating worst case execution times – some form of analysis is also required."

- From: Automating WCET analysis for DO-178B/C, Rapita Systems.

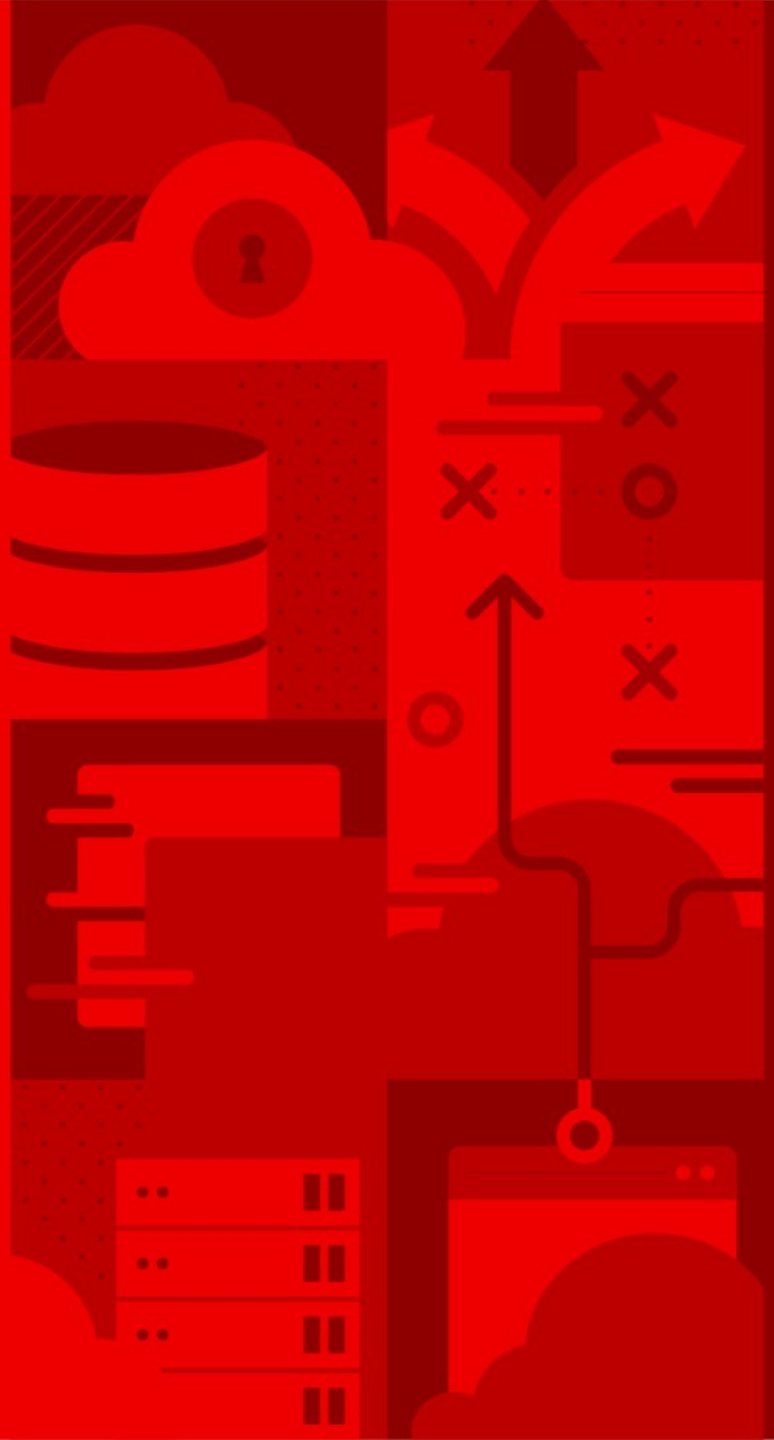

We need more analsys!

Red Hat

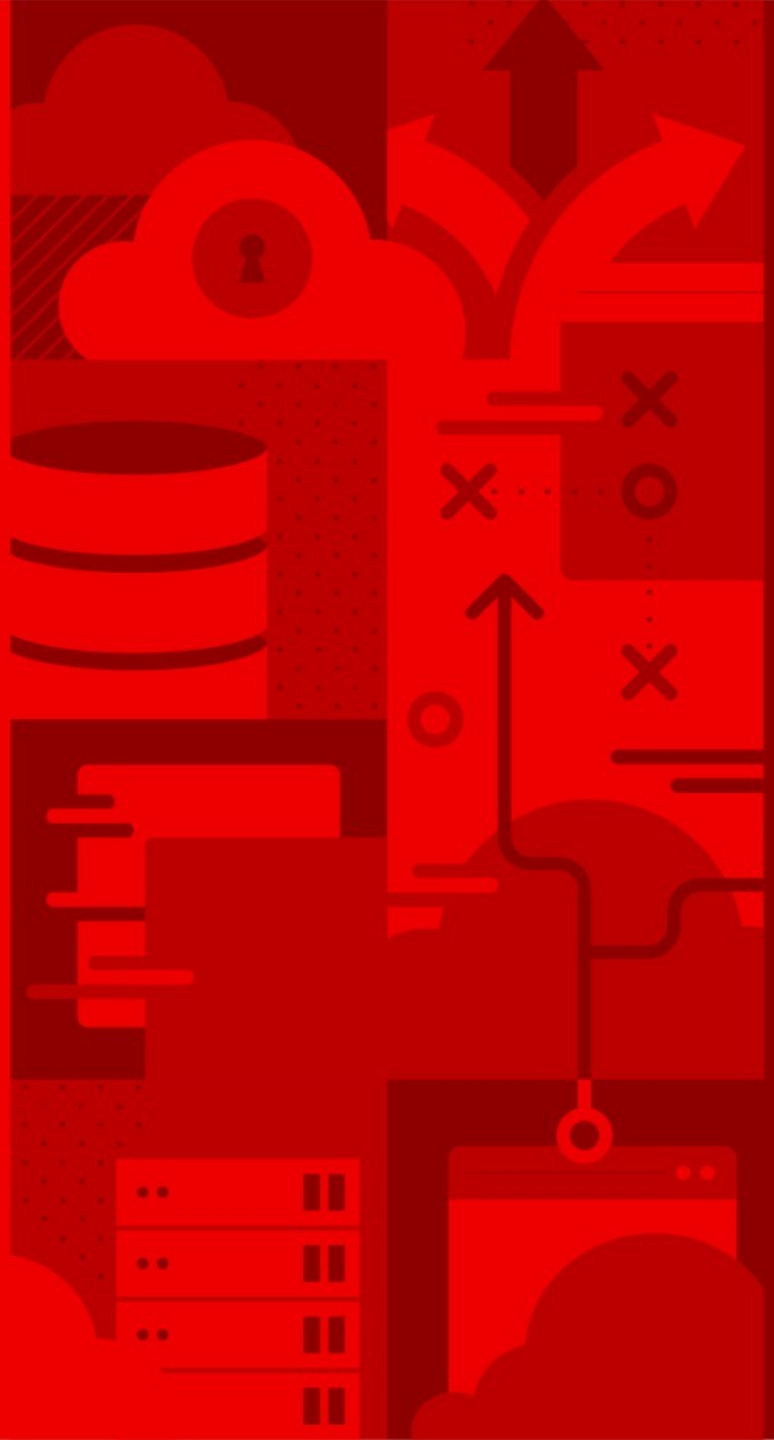

But we have some!

Red Hat

For sched deadline:

$$U = \sum U_i$$

$$U_i = \frac{C_i}{T_i}$$

$$is\ schedulable \Leftrightarrow U < 1$$

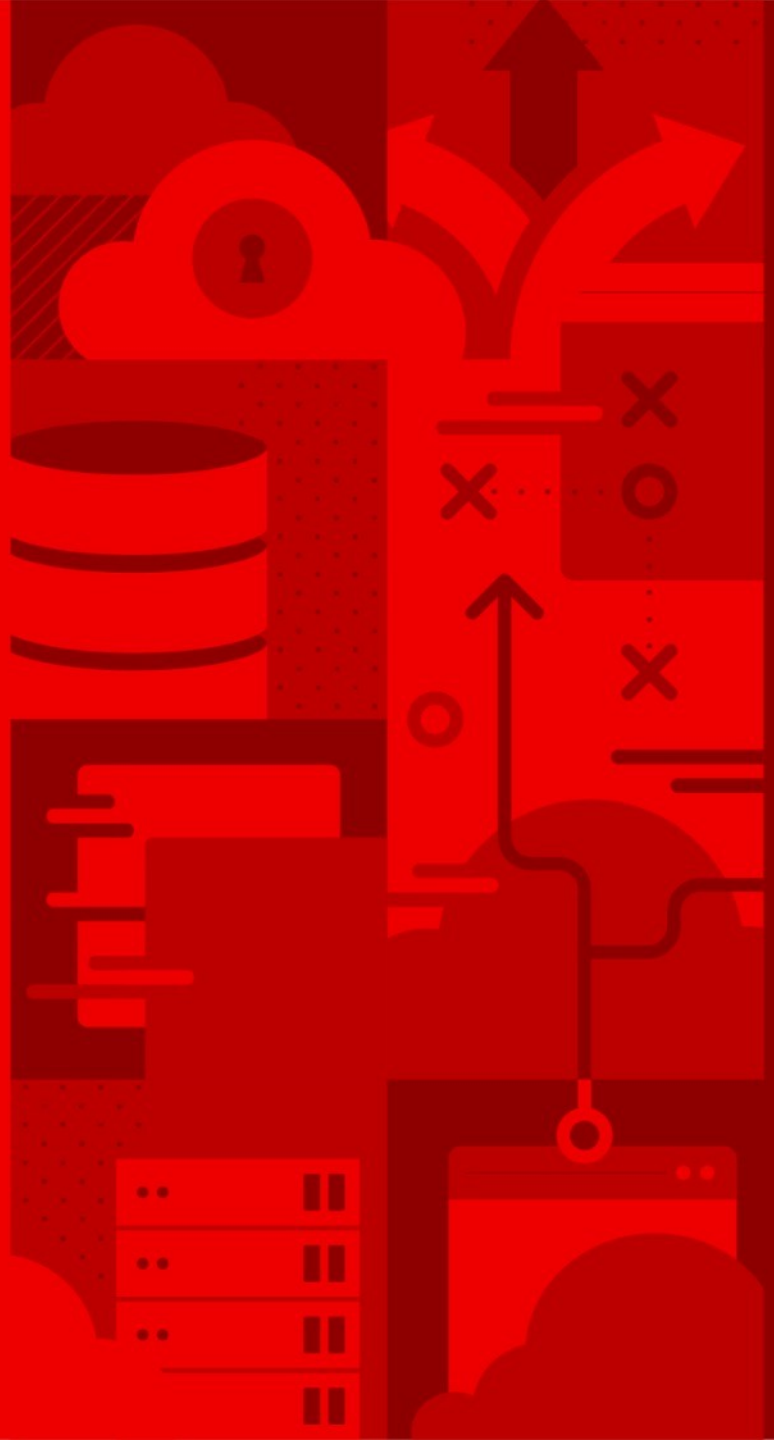

This is good!

Red Hat

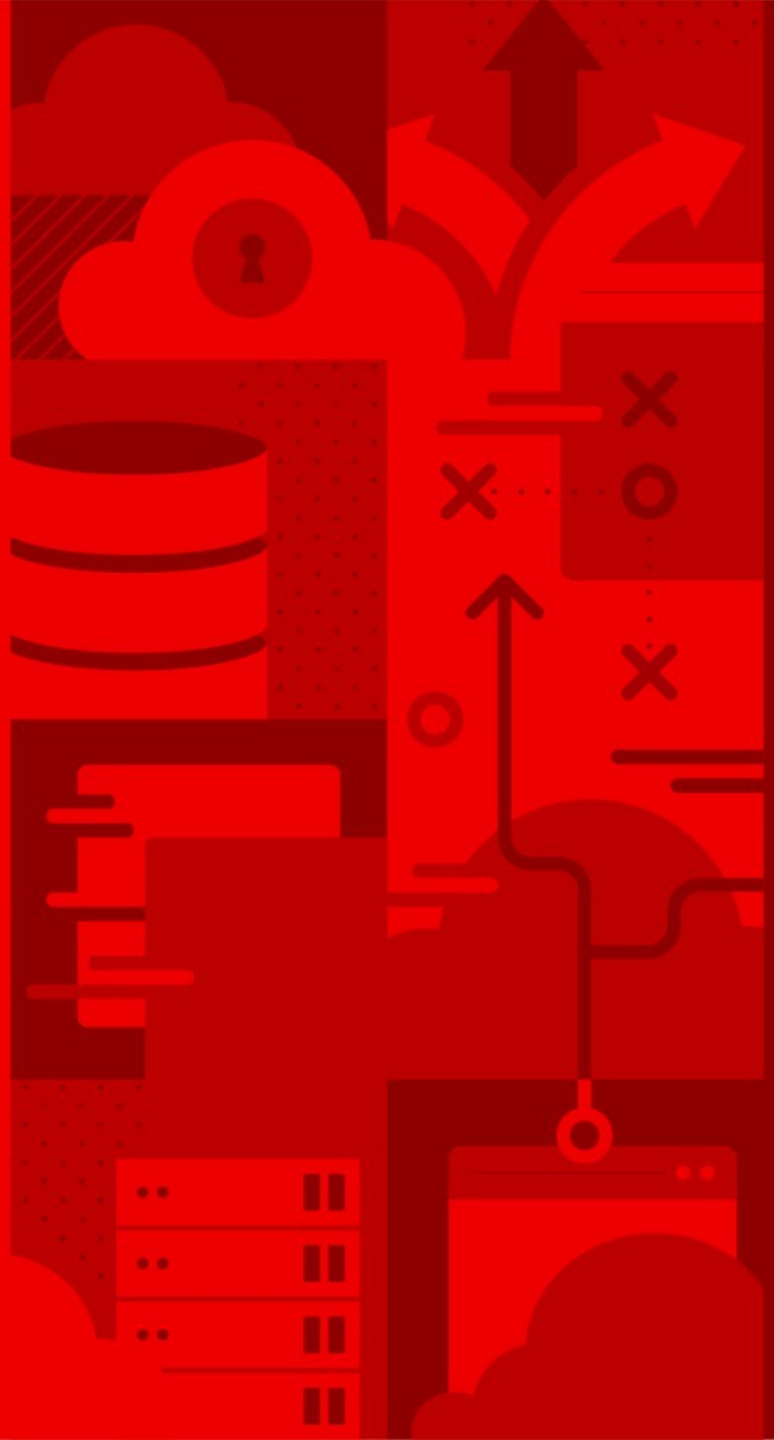

But this is an oversimplification

Red Hat

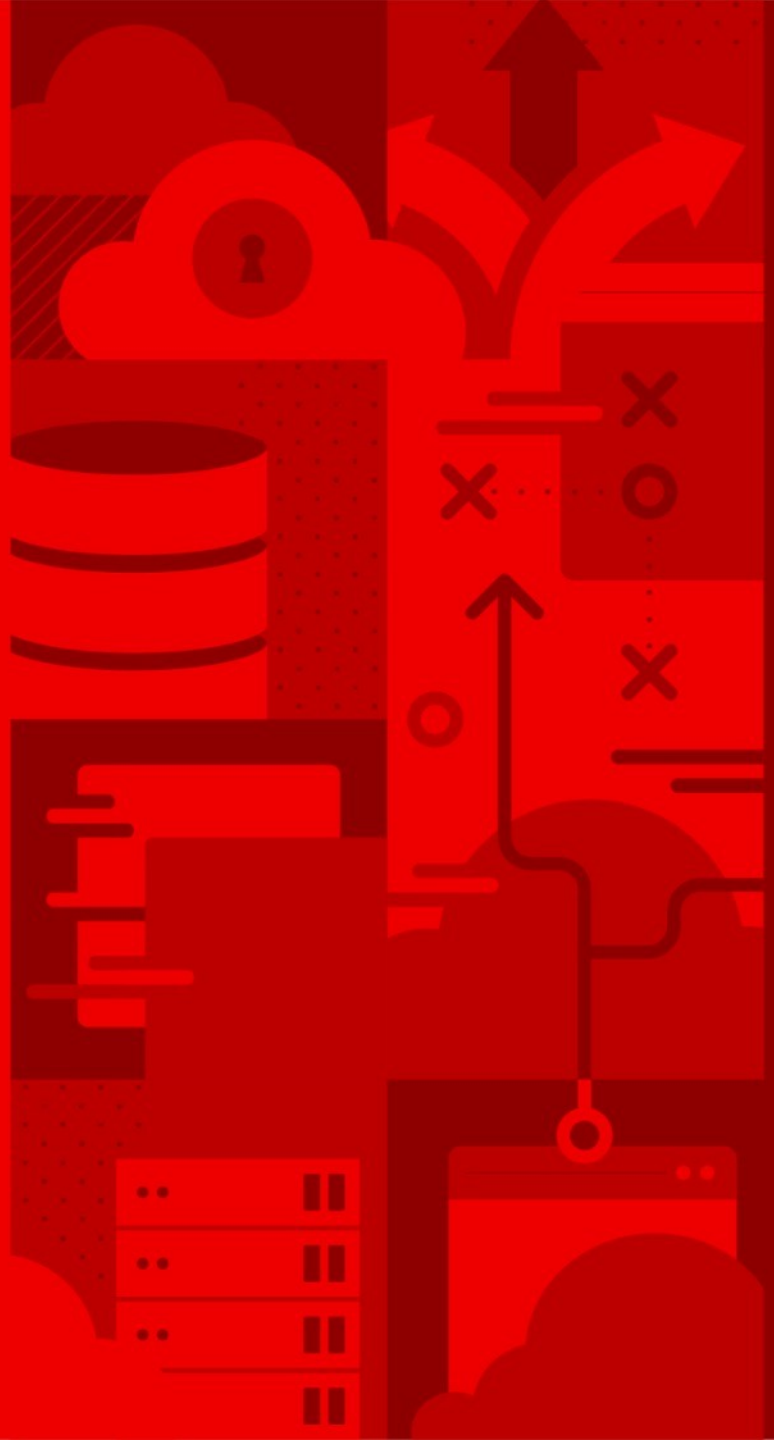

We need more fined gained analysis

Red Hat

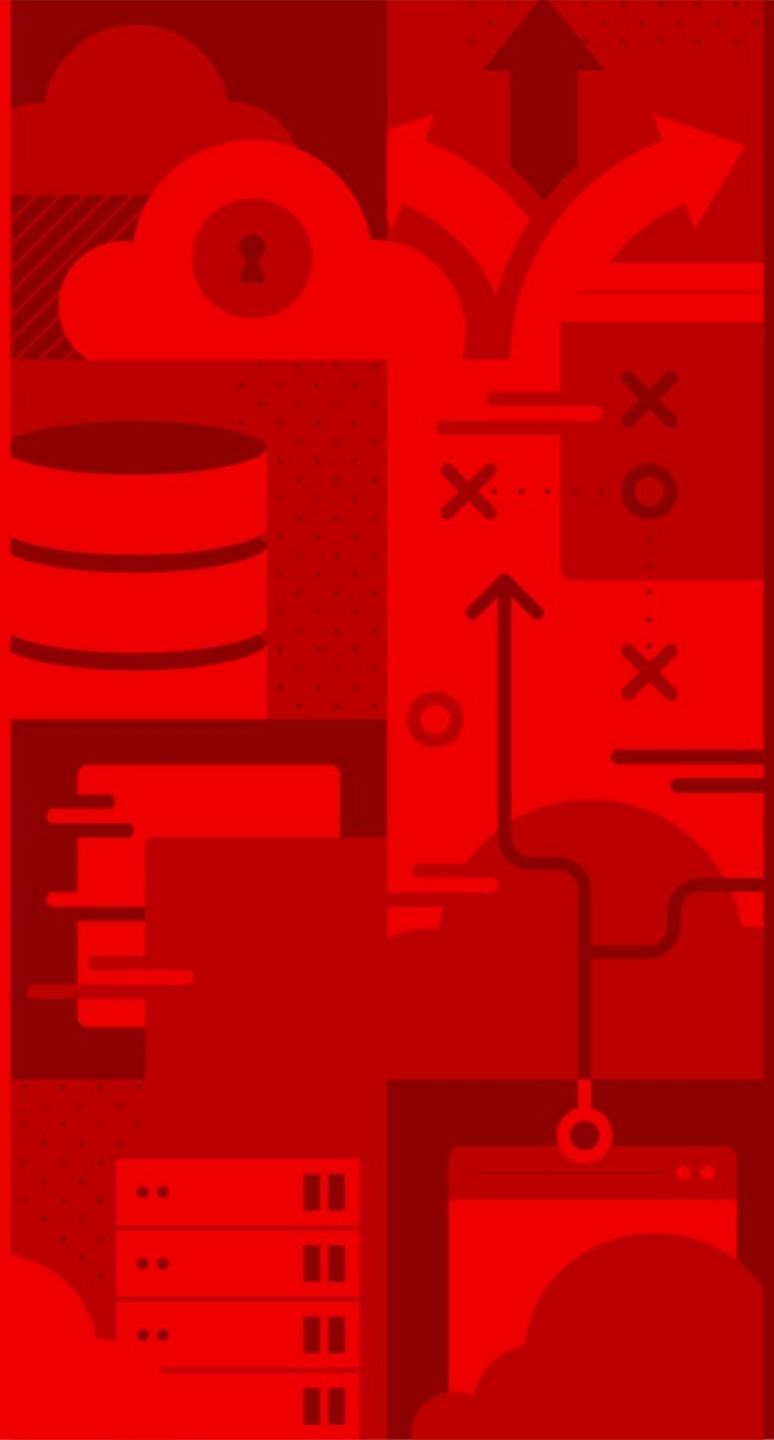

From the sched until the code!

Red Hat

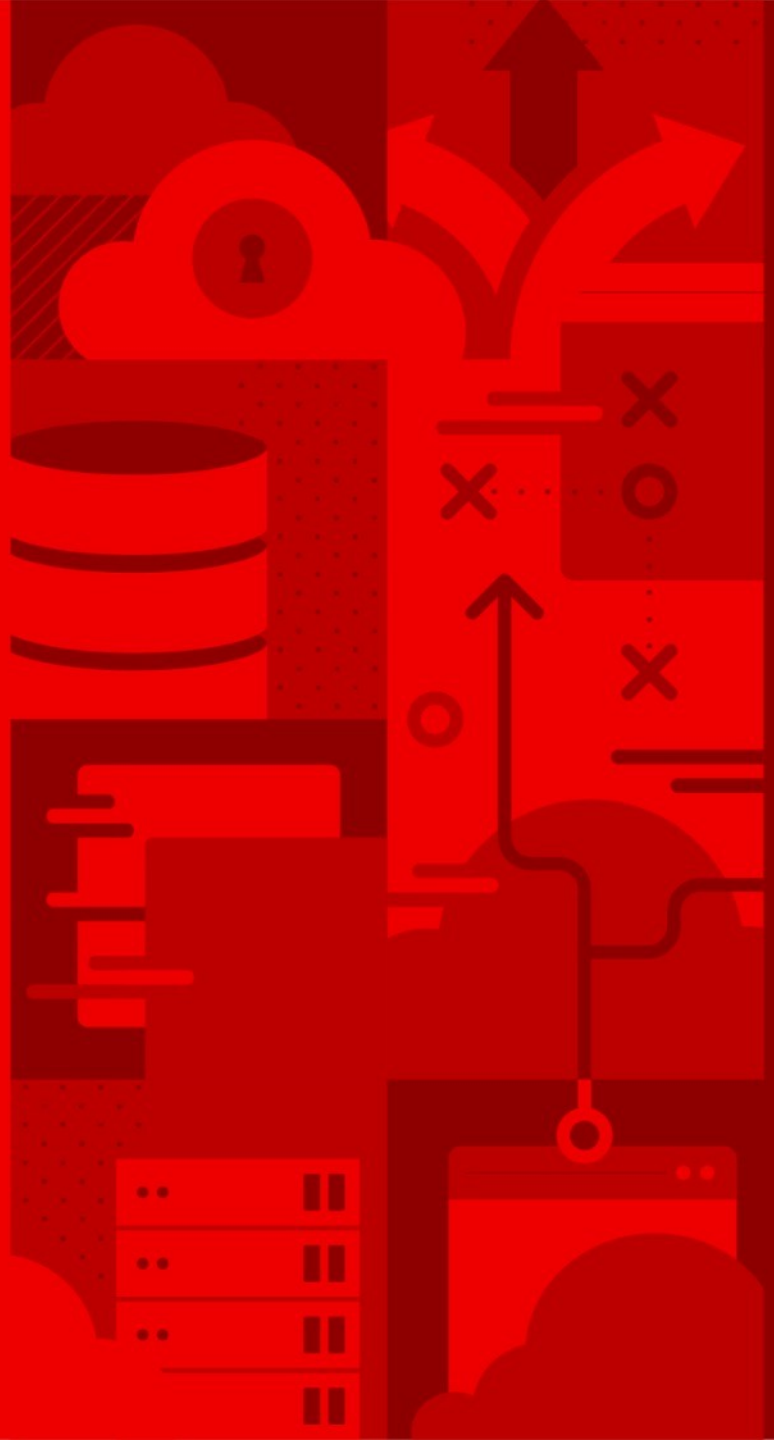

From the scheduler side

Red Hat

Sched

- Being aware of the delays caused by synchronization
- More features
  - Hierarchical
  - Arbitrary affinities
  - Better schedulability
  - We will see more this afternoon

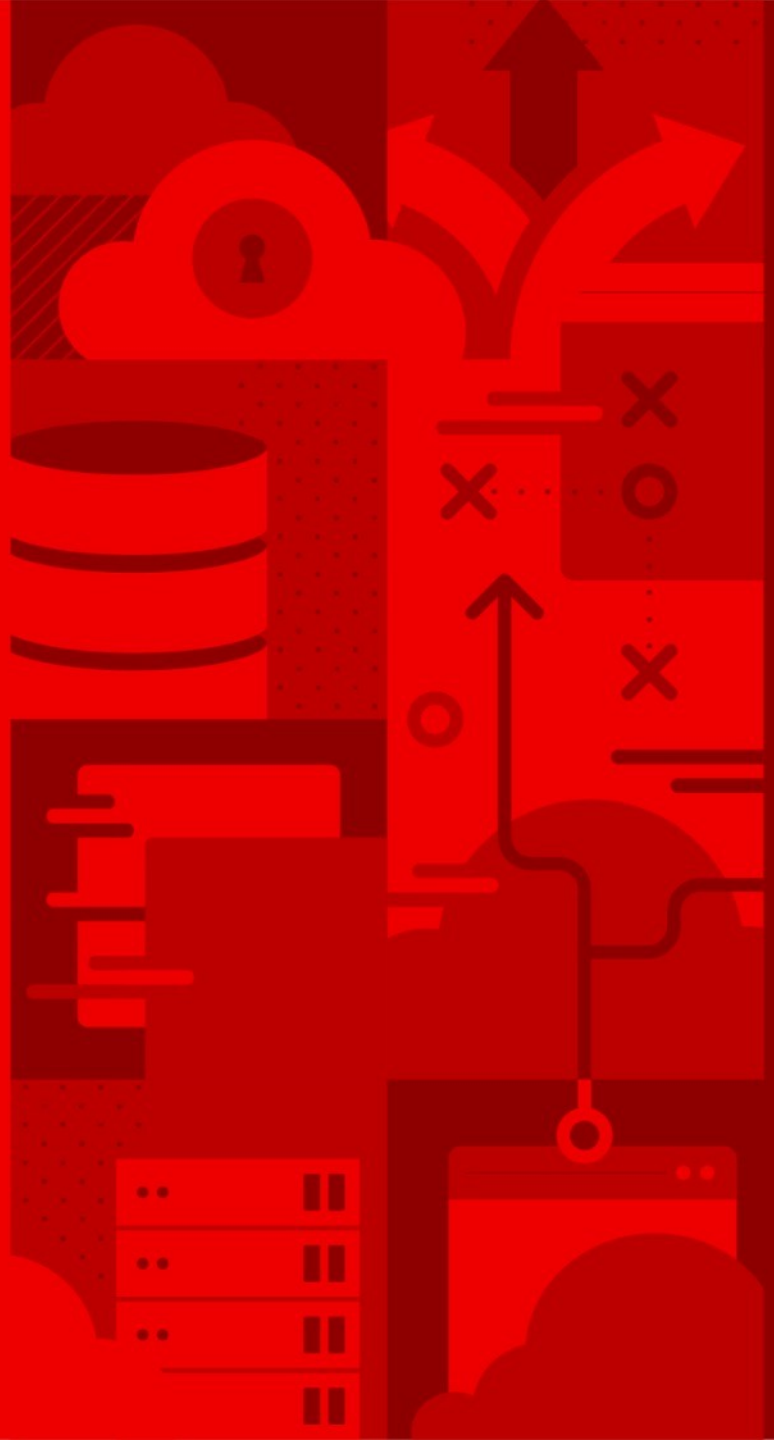

From the locking side

Red Hat

Locking

- We have problems with sched deadline
  - Proxy execution
- Other methods that do not have PI support
- Analysis from the locking (nested locking is a problem even in theory)

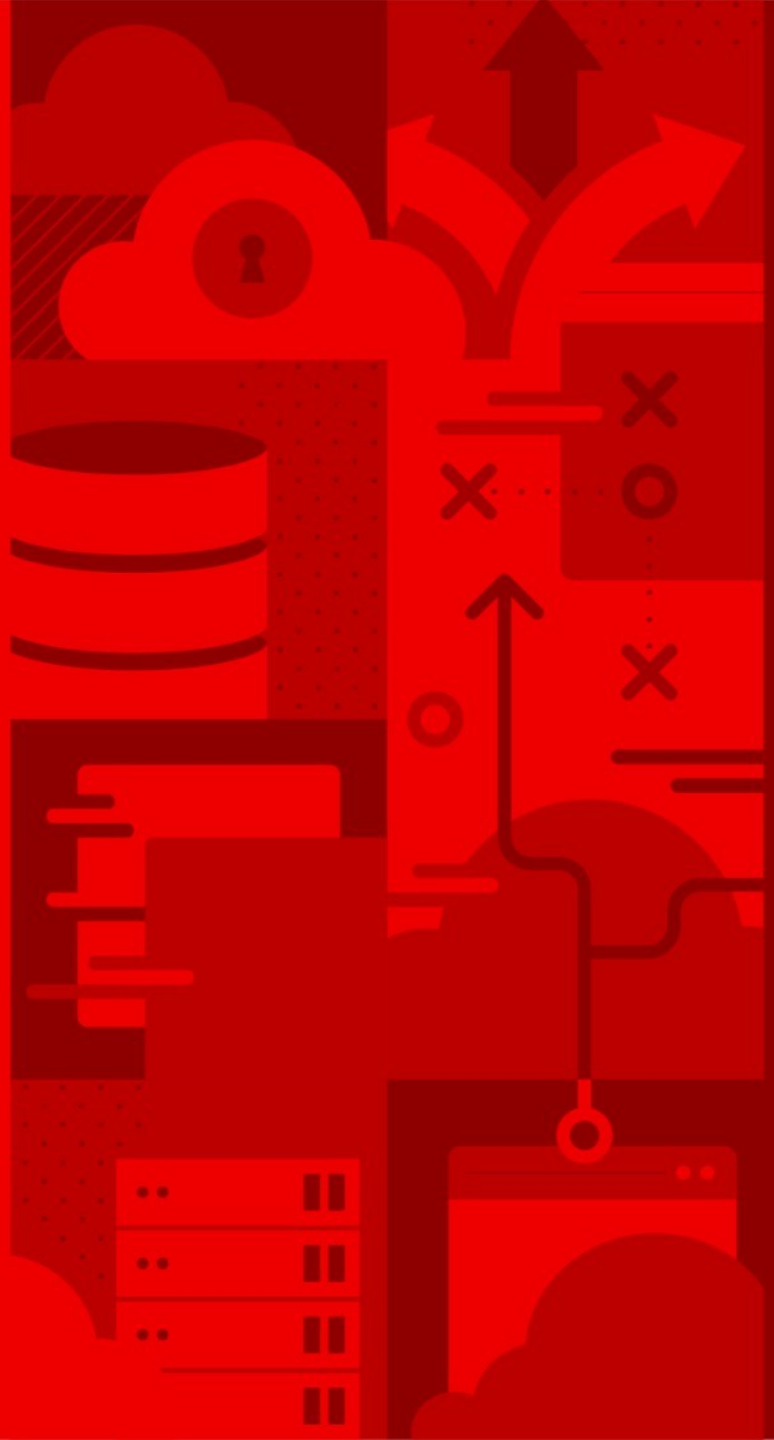

From the code side

Red Hat

WCET

- We currently just observe
- There are many people talking about pWCET
  - But there is a long way to go
  - The WCET companies in the field have pWCET but not for complex system like Linux

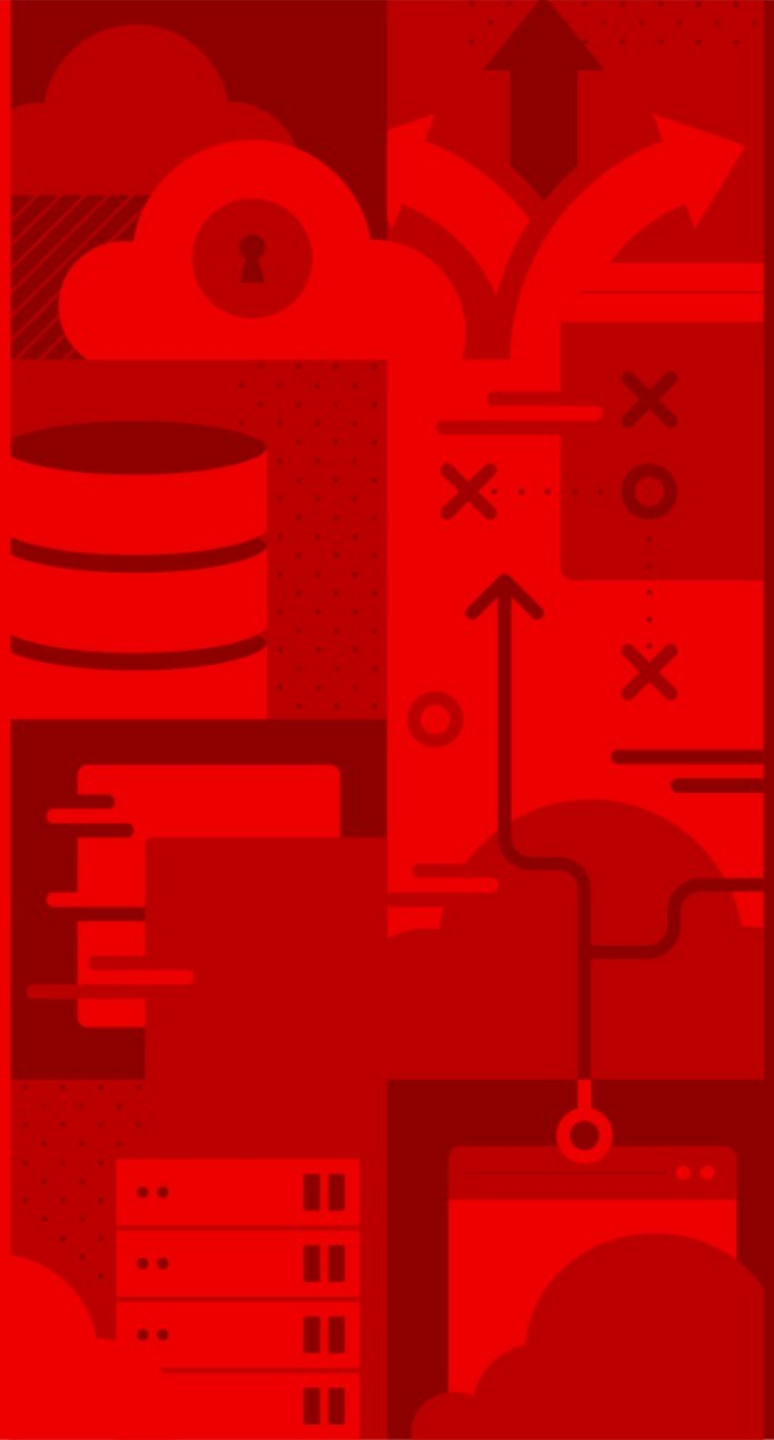

There is a lot of work to do!

Red Hat

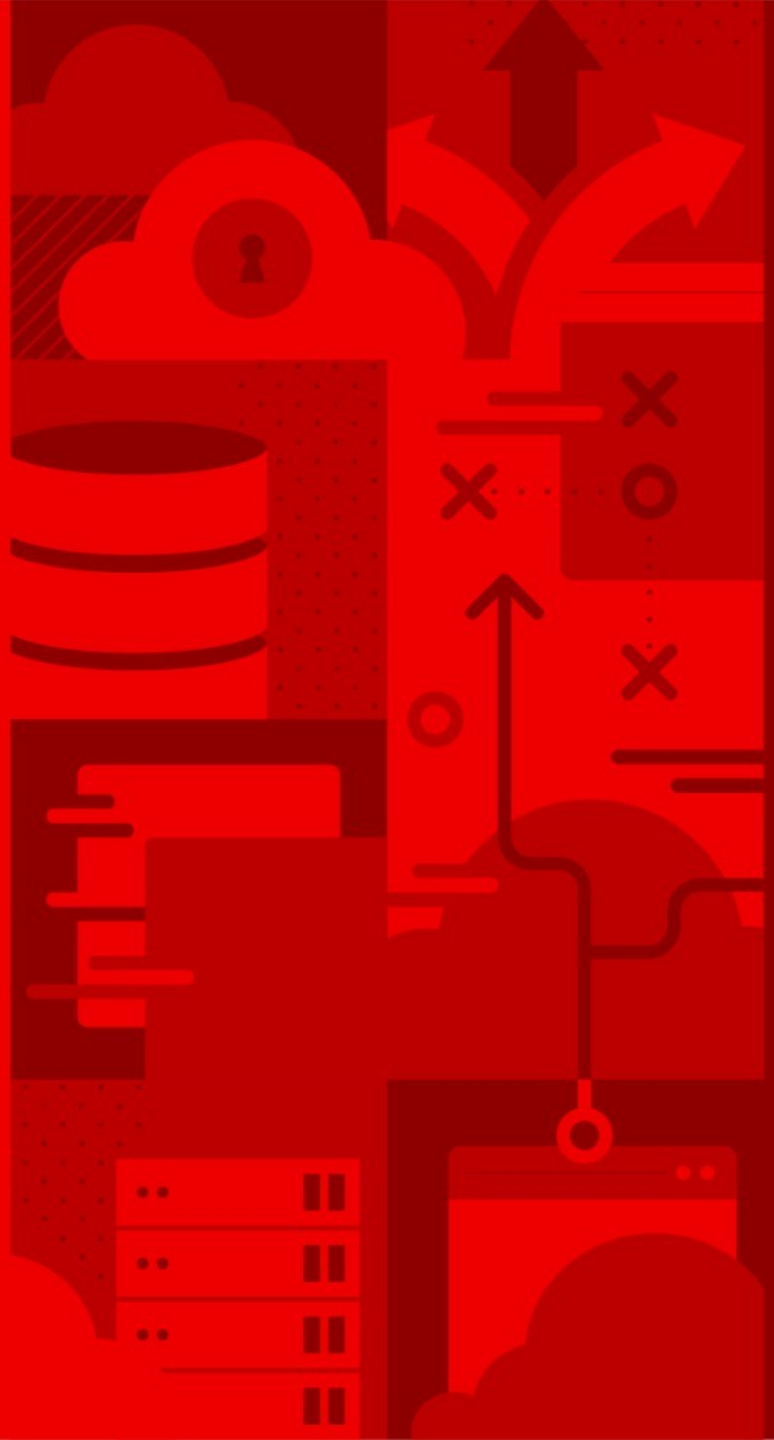

So Linux is not a Hard RTOS

Red Hat

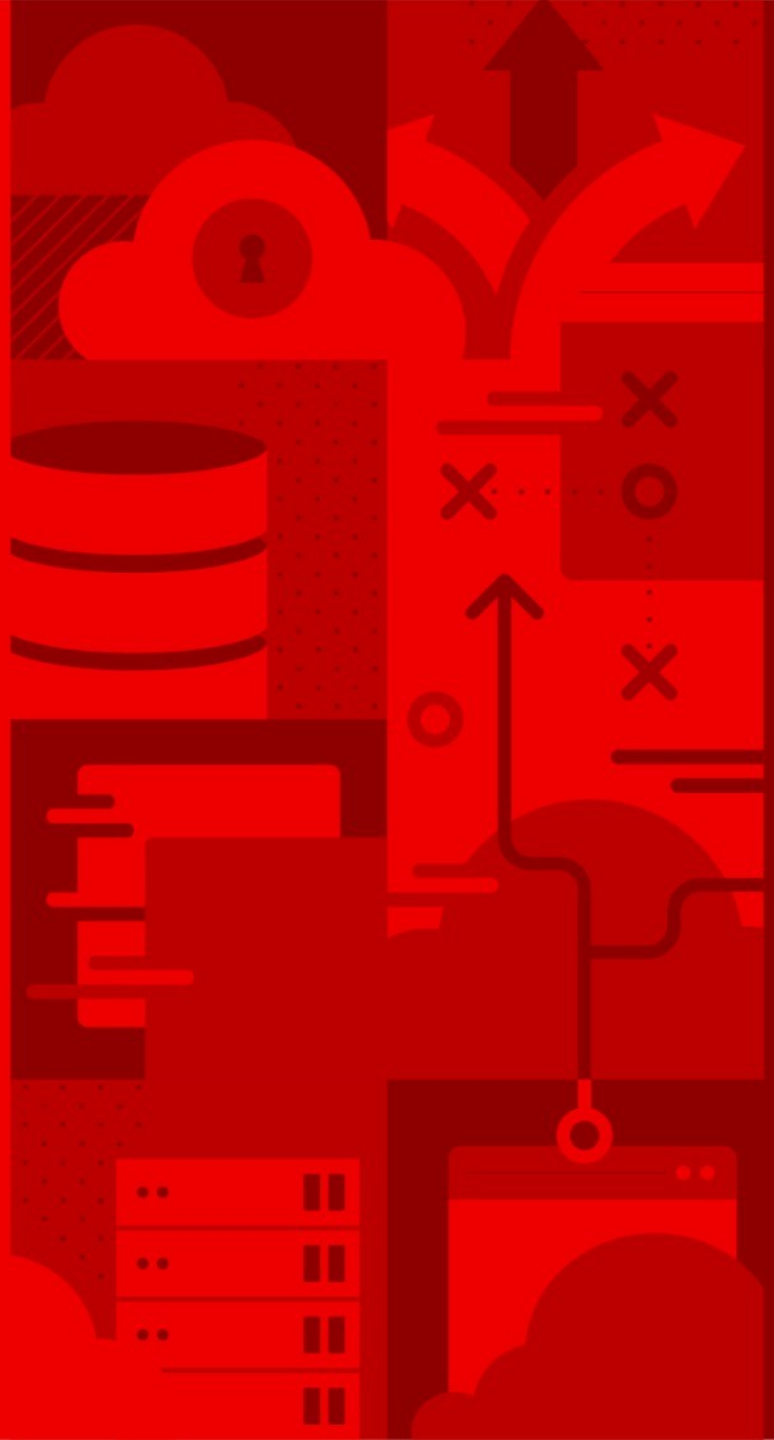

But we are touching the edges of RT theory

Red Hat

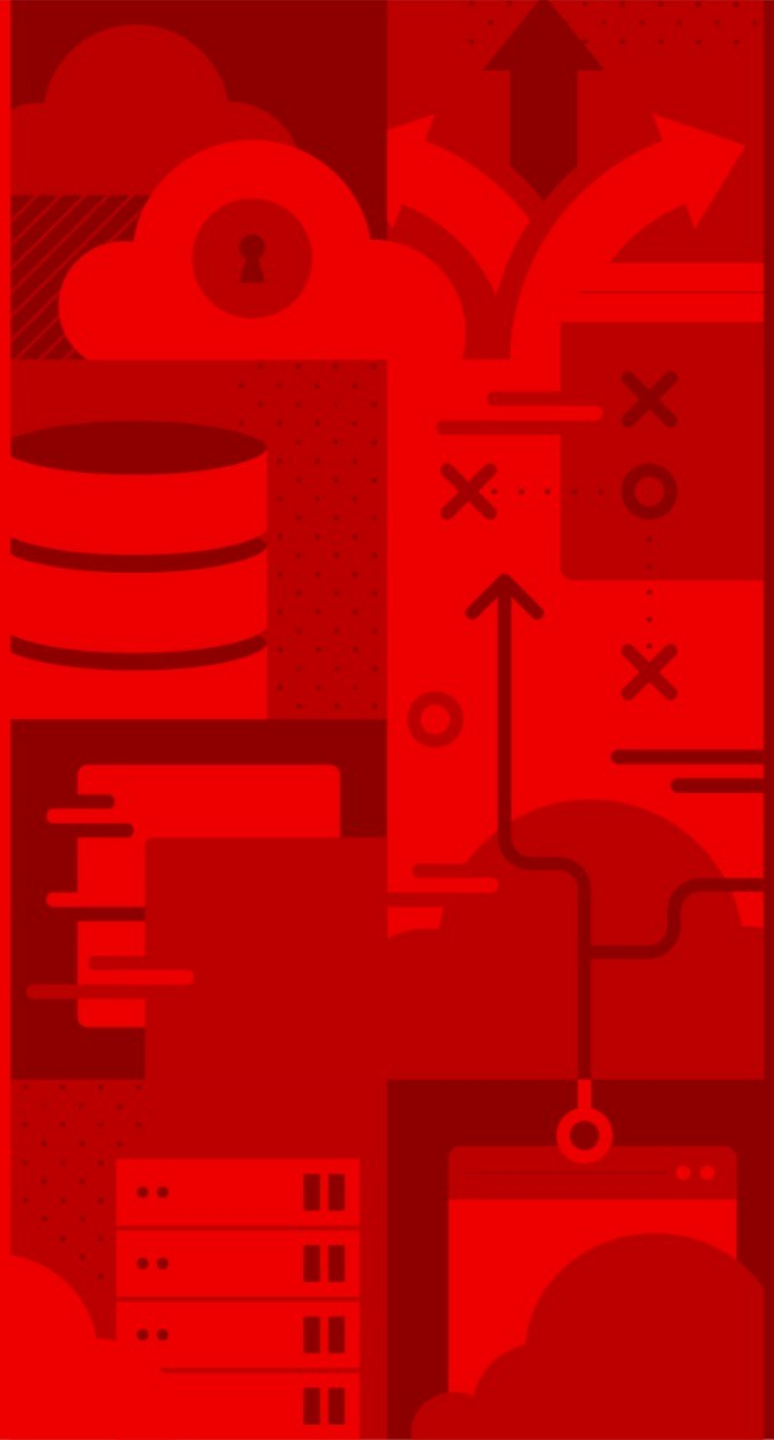

And the PREEMPT RT Enable that!

Red Hat

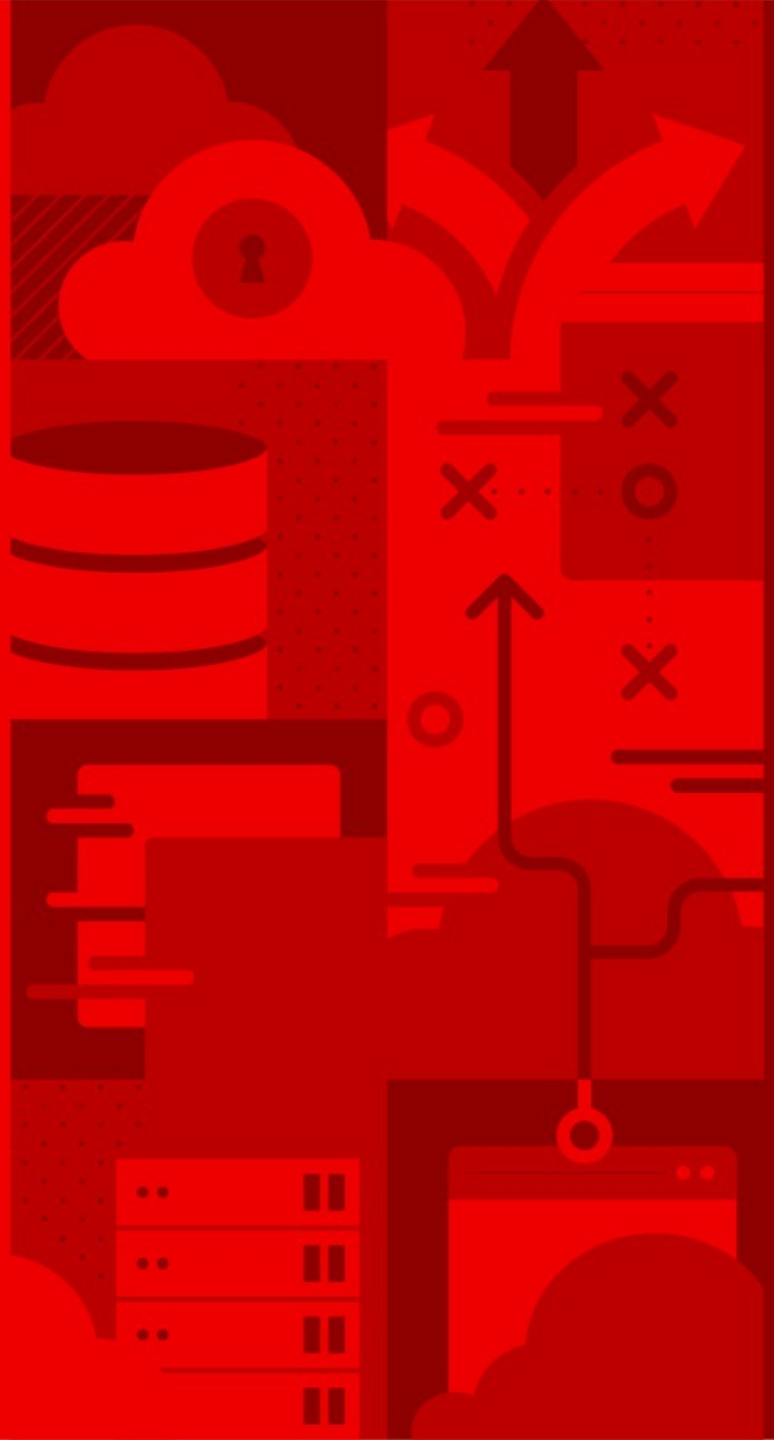

Job Security

Red Hat

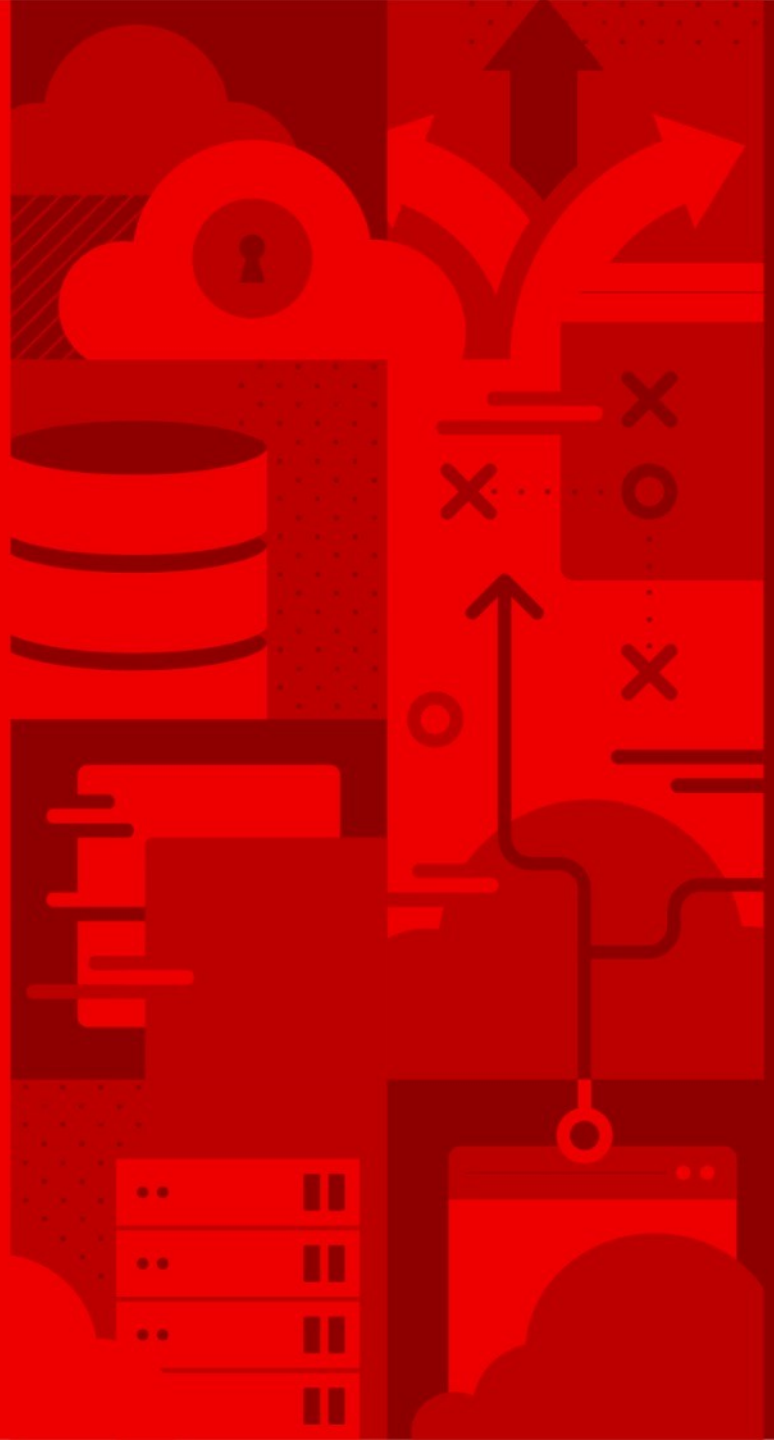

That is it

Red Hat

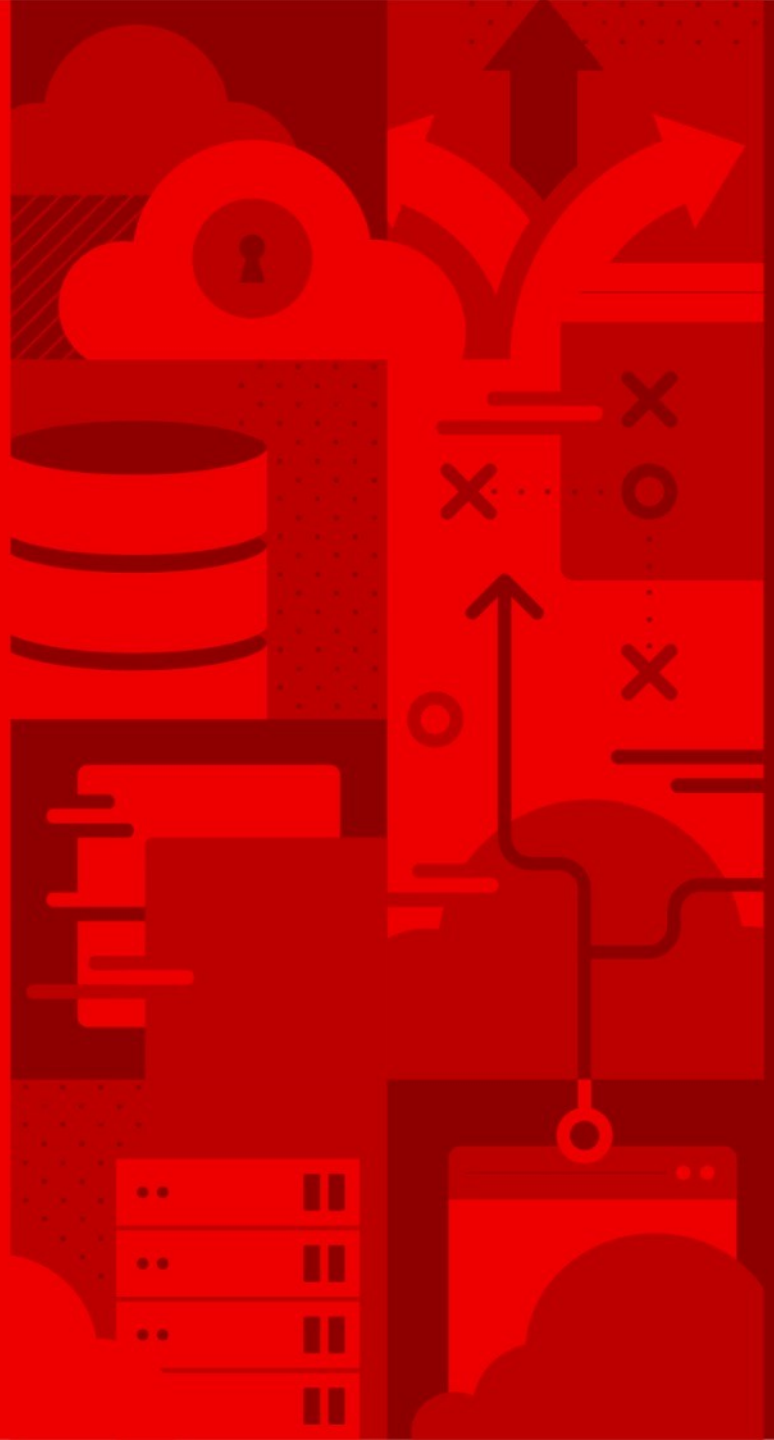

Thoughts?

Red Hat