

# Update on active path from kernel trace

Progress Report Meeting  
Dec. 10<sup>th</sup> 2015

Francis Giraldeau  
francis.giraldeau@polymtl.ca

Under the direction of Michel Dagenais  
DORSAL Lab, École Polytechnique de Montréal

Base template

django-httpd.phd.vm/wkdb/polls/

- question 1
- question 2

Elements Network Sources Timeline Profiles Resources Audits Console

Preserve log

Name Path	M...	St... Text	Type	Initiator	Size Cont...	Time Later	Timeline	
polls/ /wkdb	GET	200 OK	te...	Other	49... 298 E	23... 22 m:	22 ms	40 ms

Blocking 1.448 ms

DNS Lookup 0.354 ms

Connecting 0.274 ms

Sending 0.060 ms

Waiting 20.220 ms

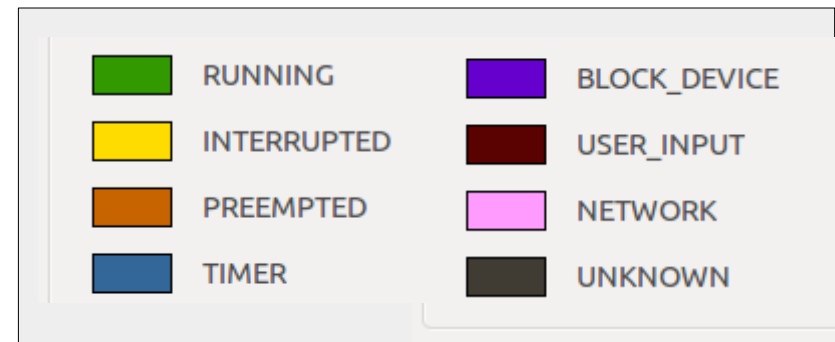
Receiving 0.718 ms

1 requests | 493 B transferred | 23 ms (load: 56 ms, DOMContentLoaded: 57 ms)

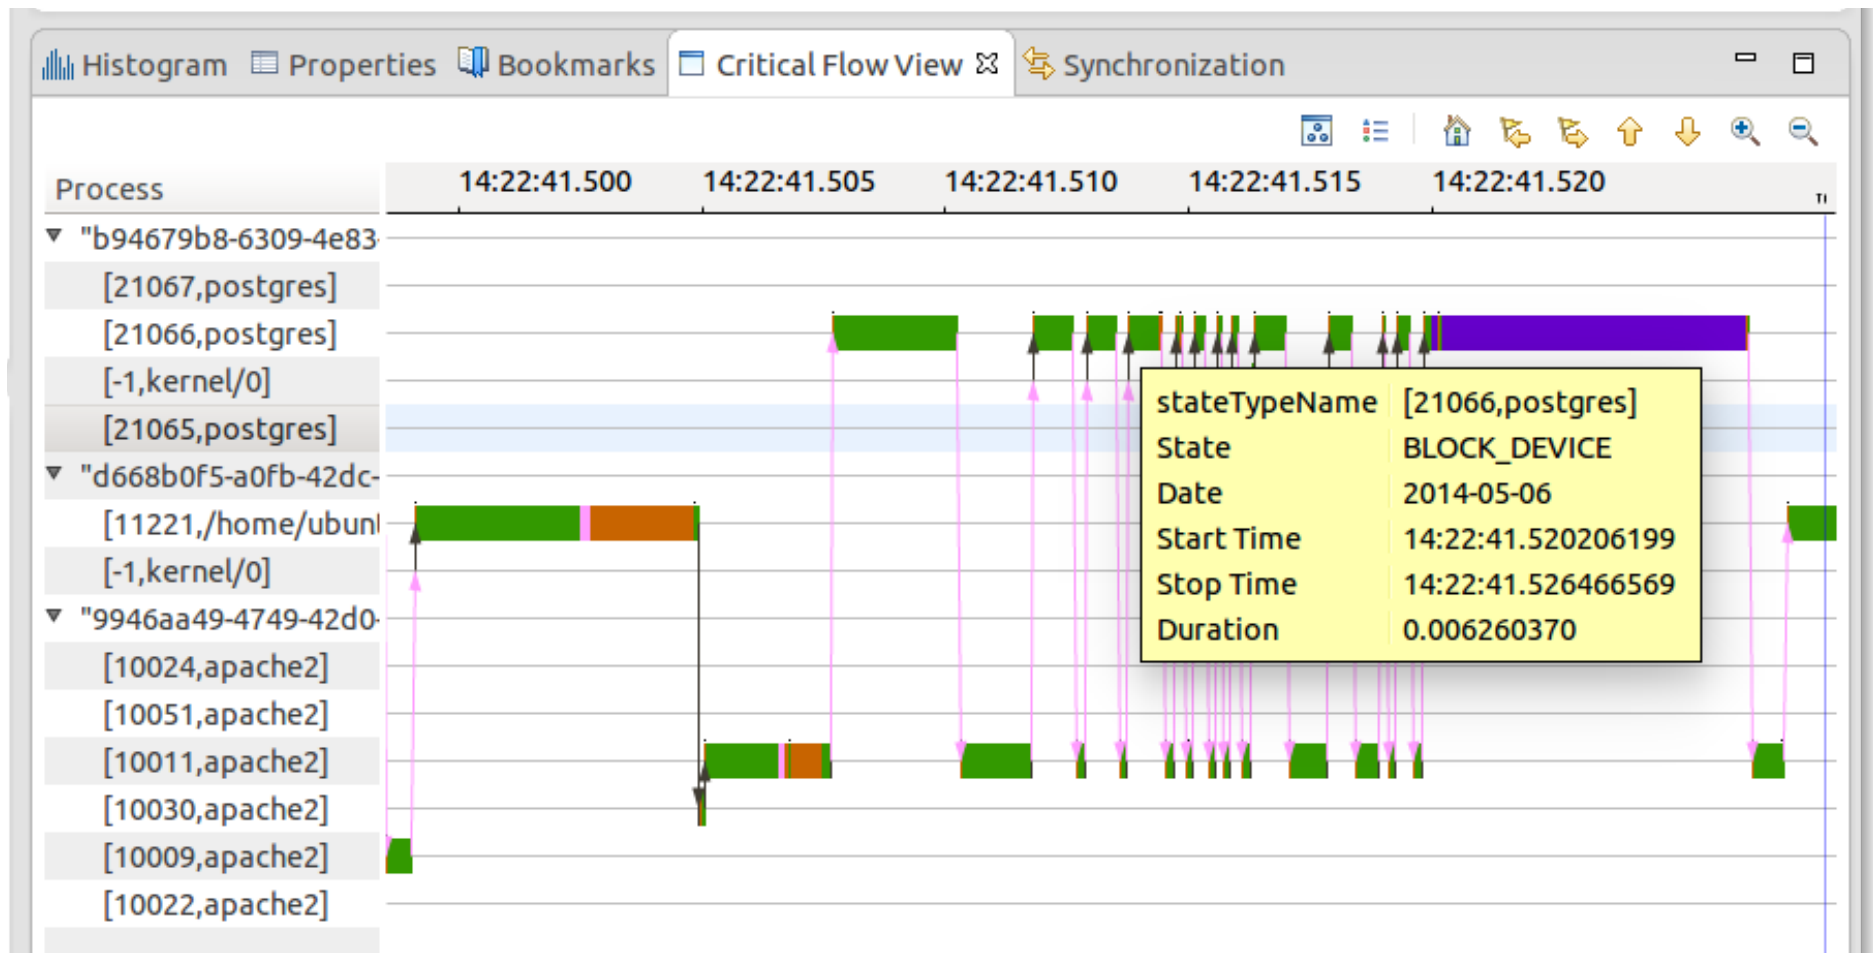
What does the server?

# Django Polls App

mechanize + Apache WSGI +  
postgresql




## 2) POST form



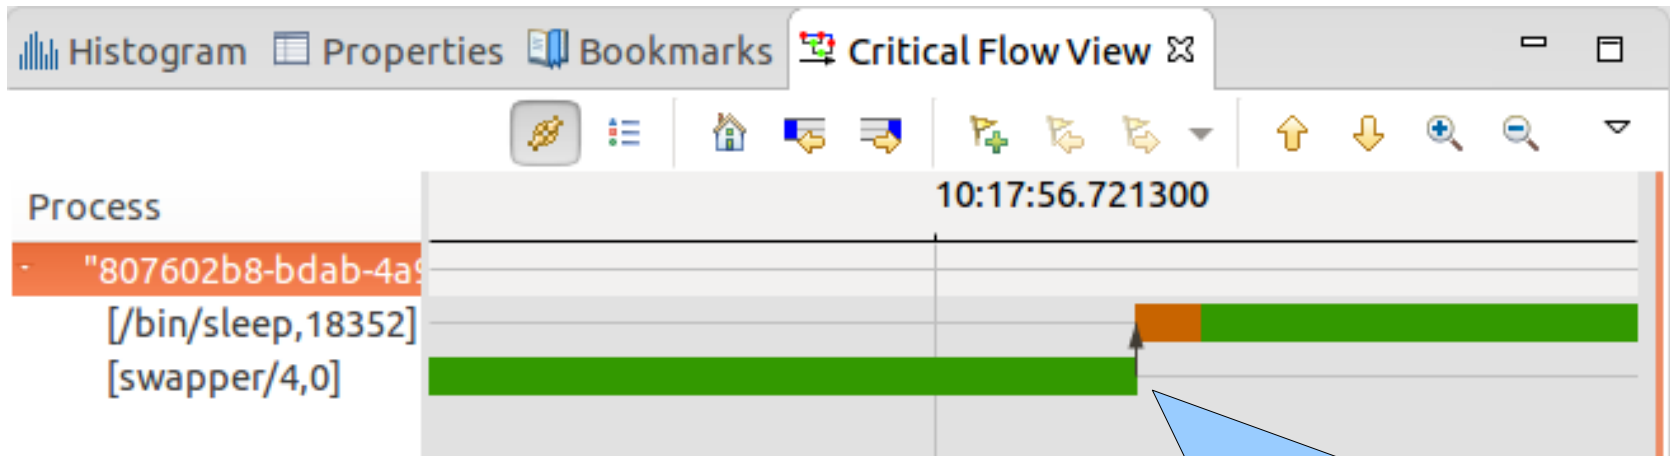
# Device wake-up




# Task wake-up




# Wrong active path



No task waits for the idle thread swapper



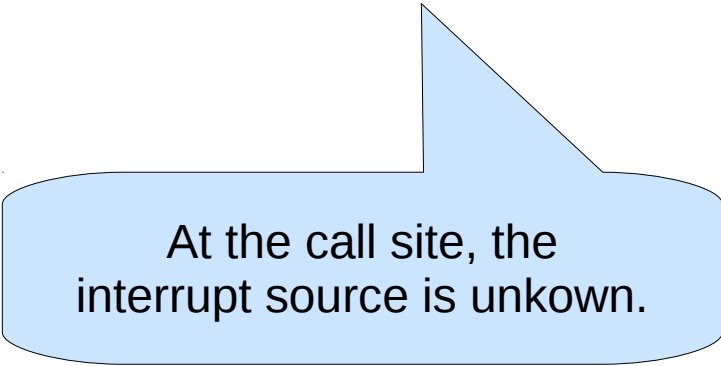
# Inter-process interrupt (IPI)



\*Advanced Programmable  
Interrupt Controller

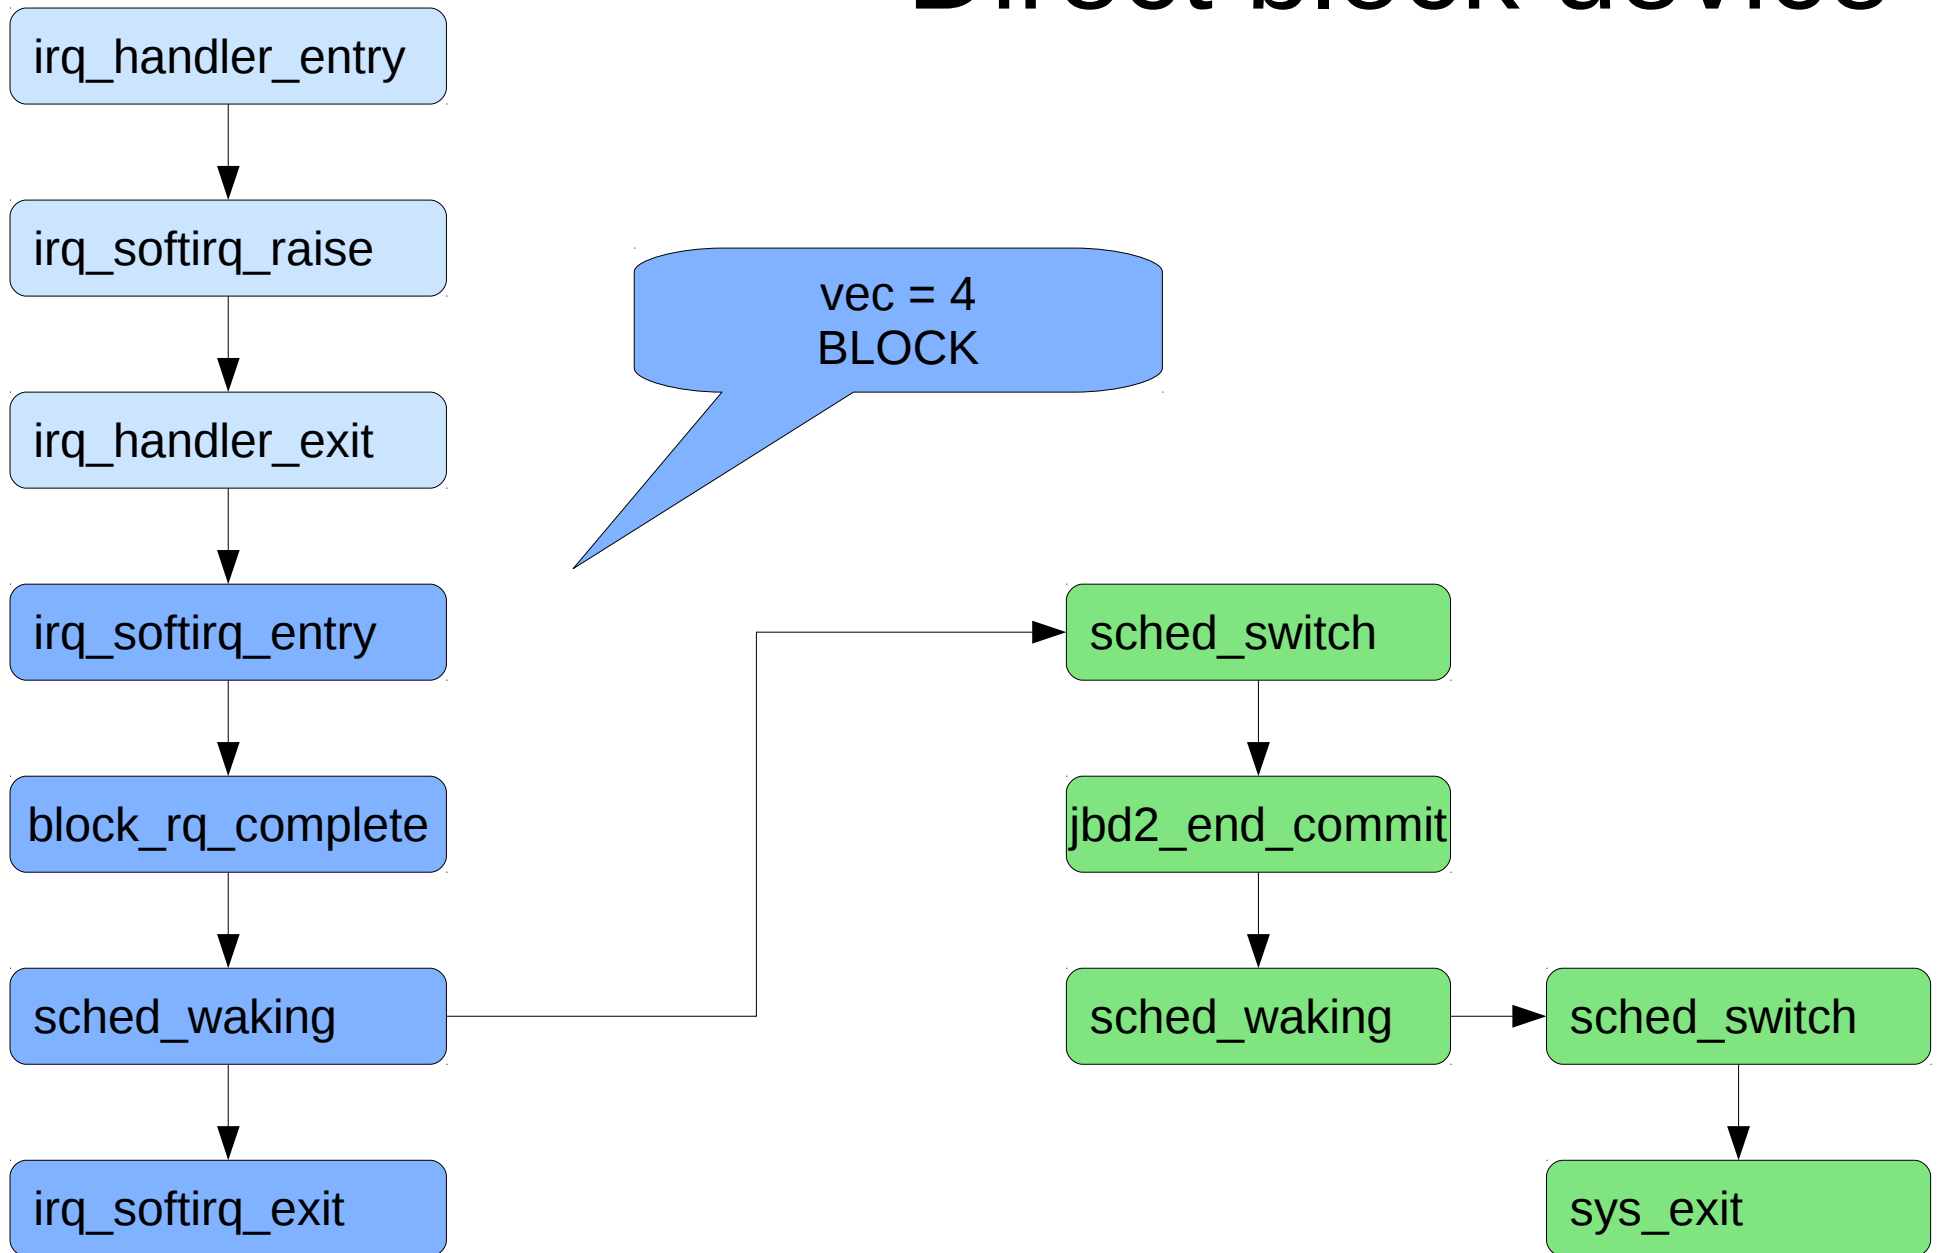
# Event sequence

- Source CPU
- Interrupt occurs
- Enqueue work item to per CPU list (fn, data)
- Raise IPI on target CPU
- Target CPU
- IPI occurs
- Dequeue work item
- Call fn(data)



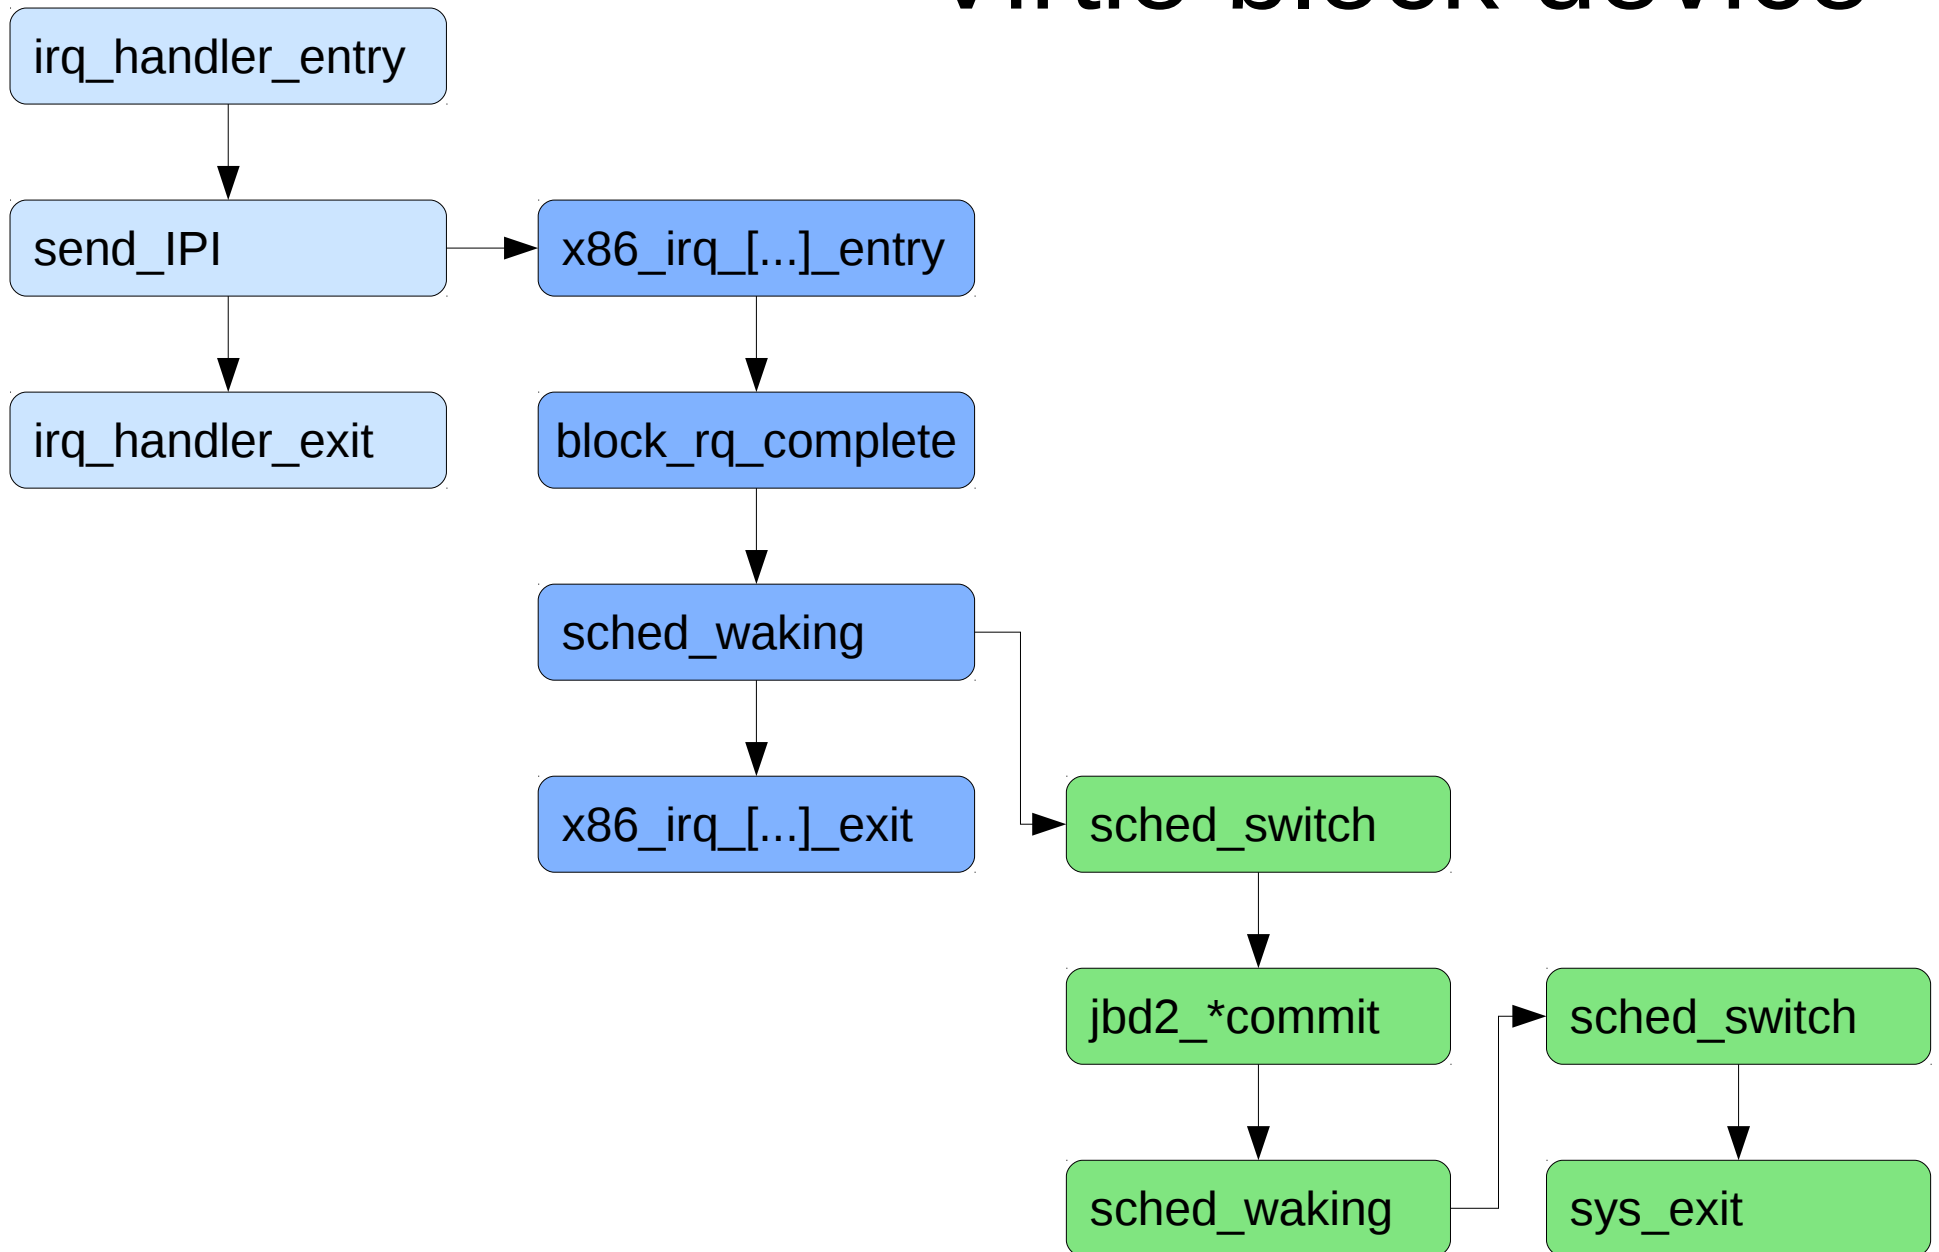
At the call site, the interrupt source is unknown.

# Direct block device

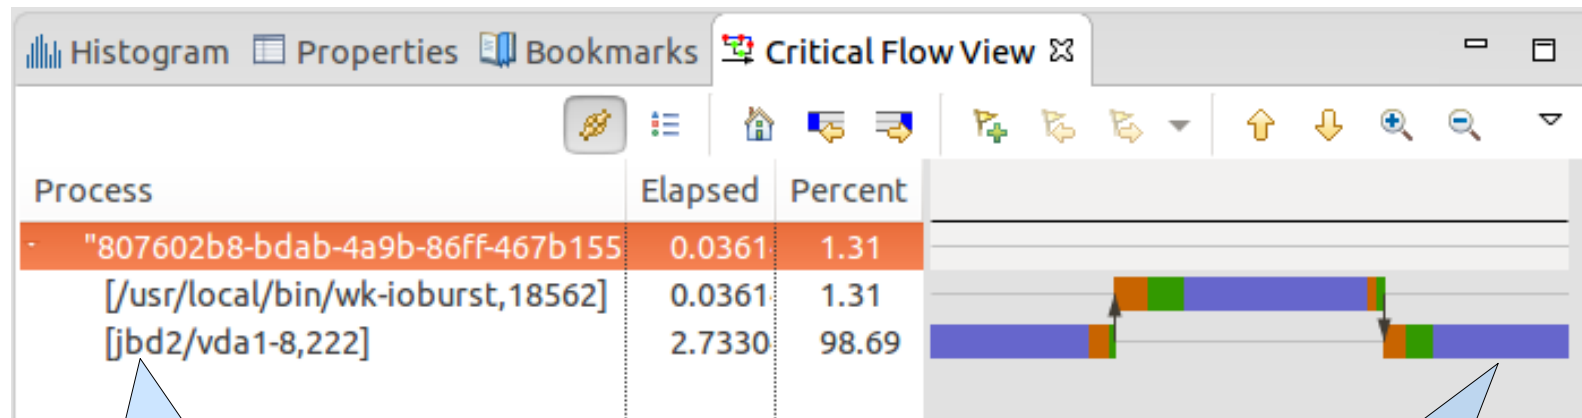




# Virtio block device



# Virtio block device



Swapper is no more  
on the active path

New IPI color  
Actual IPI source is unknown

# Conclusion

- New sched\_waking available with linux  $\geq 4.3$  and LTTng  $\geq 2.8$ 
  - Patch backported to linux RT 4.1
- New IPI entry/exit events are required
- Virtio block device uses IPI under the hood

# Download the paper on IEEE Xplore



IEEE TRANSACTIONS ON  
PARALLEL AND  
DISTRIBUTED SYSTEMS

**OPEN**  **ACCESS**



<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7294678>

Thanks to Professor Michel Dagenais and our partners  
EfficiOS and Ericsson.

Home page: <http://step.polymtl.ca/~fgiraldeau/>

Software:

<http://secretaire.dorsal.polymtl.ca/~fgiraldeau/workload-kit/>

<http://secretaire.dorsal.polymtl.ca/~fgiraldeau/traceset/>

<https://github.com/giraldeau>

