



# Virtual Machine Boot-up Analysis

Abderrahmane Benbachir

May 5, 2017

École Polytechnique de Montréal

Laboratoire **DORSAL**

# Agenda

---

**Introduction**

**Tracing Through Para-virtualization Layers**

**Investigating Boot-up Issues**

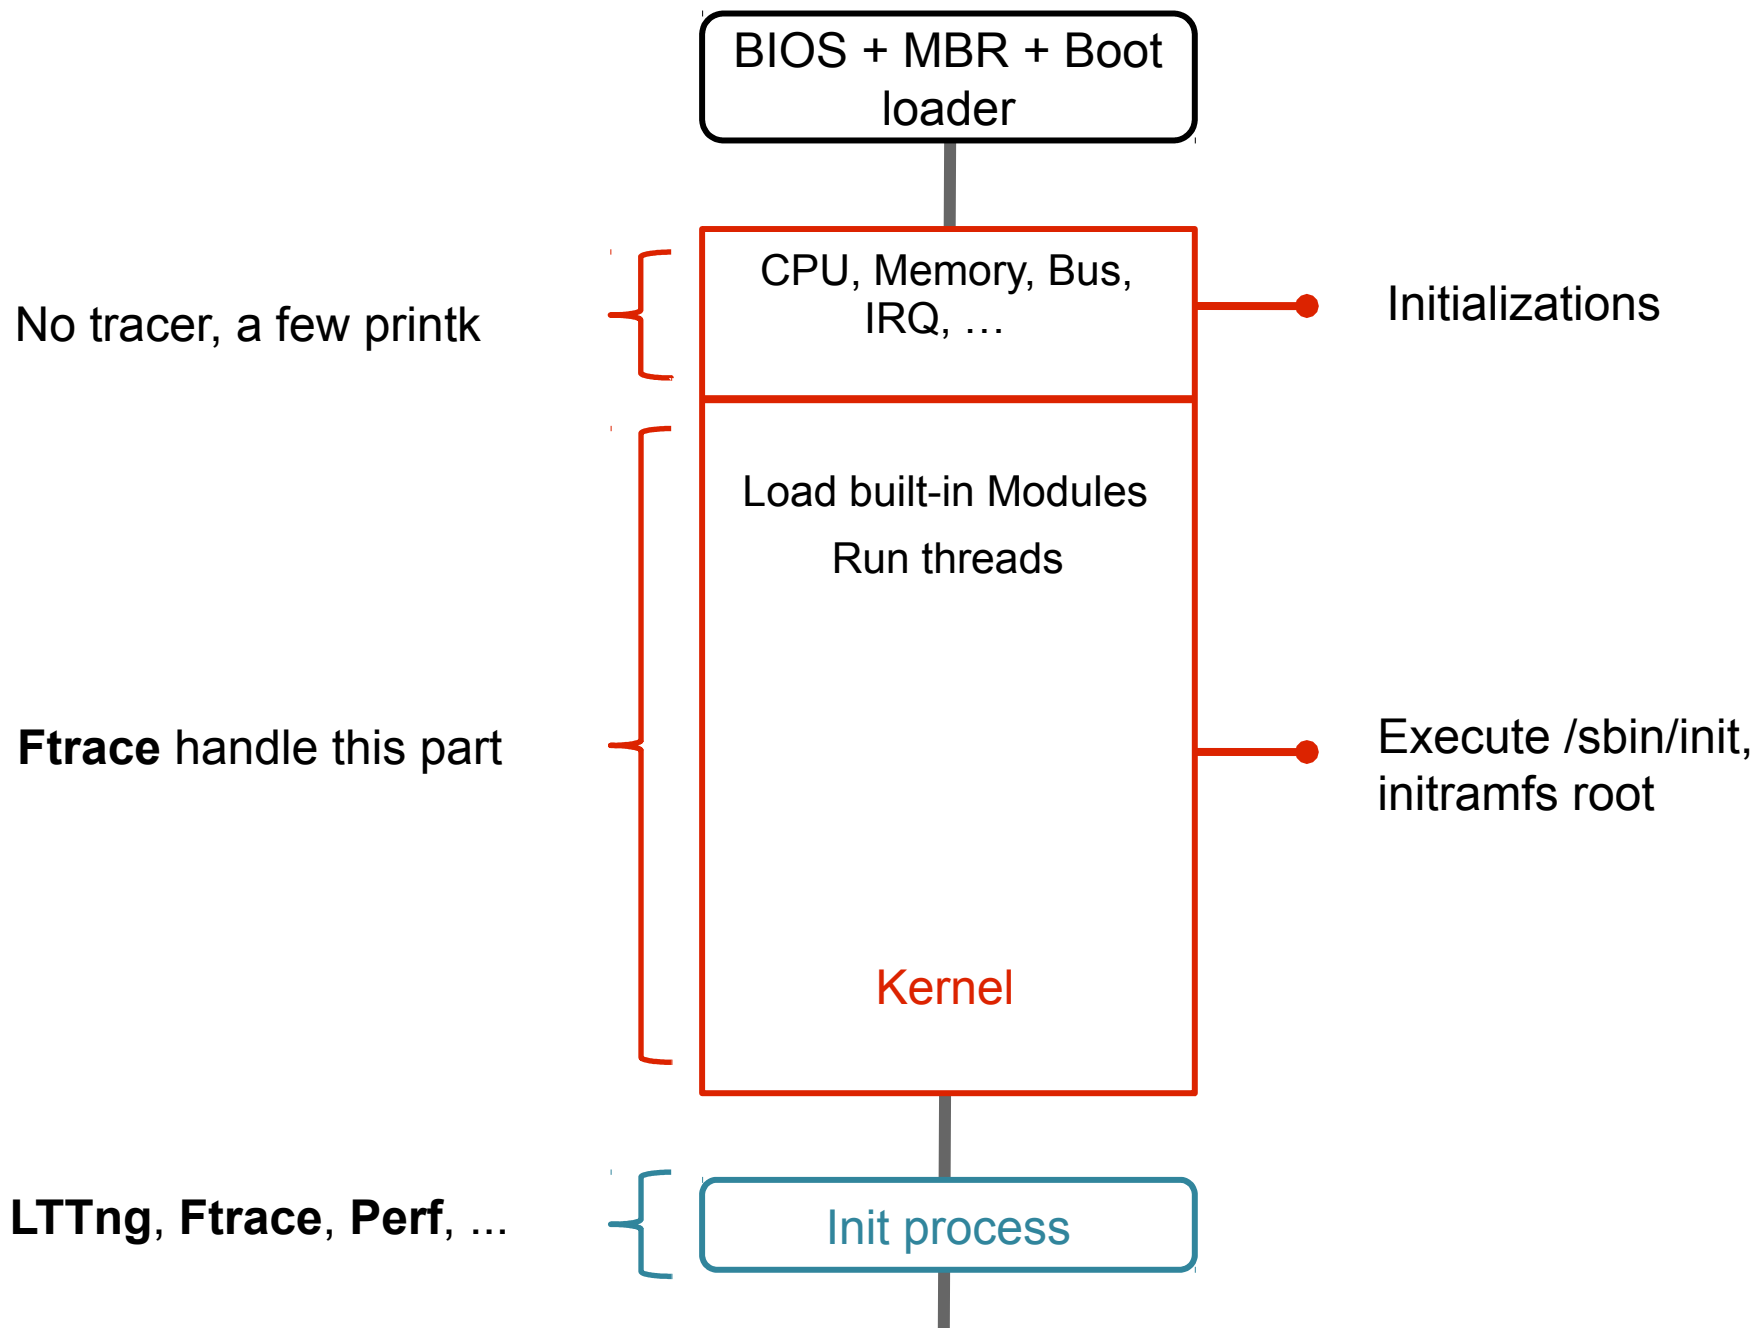
**Questions**

# Introduction

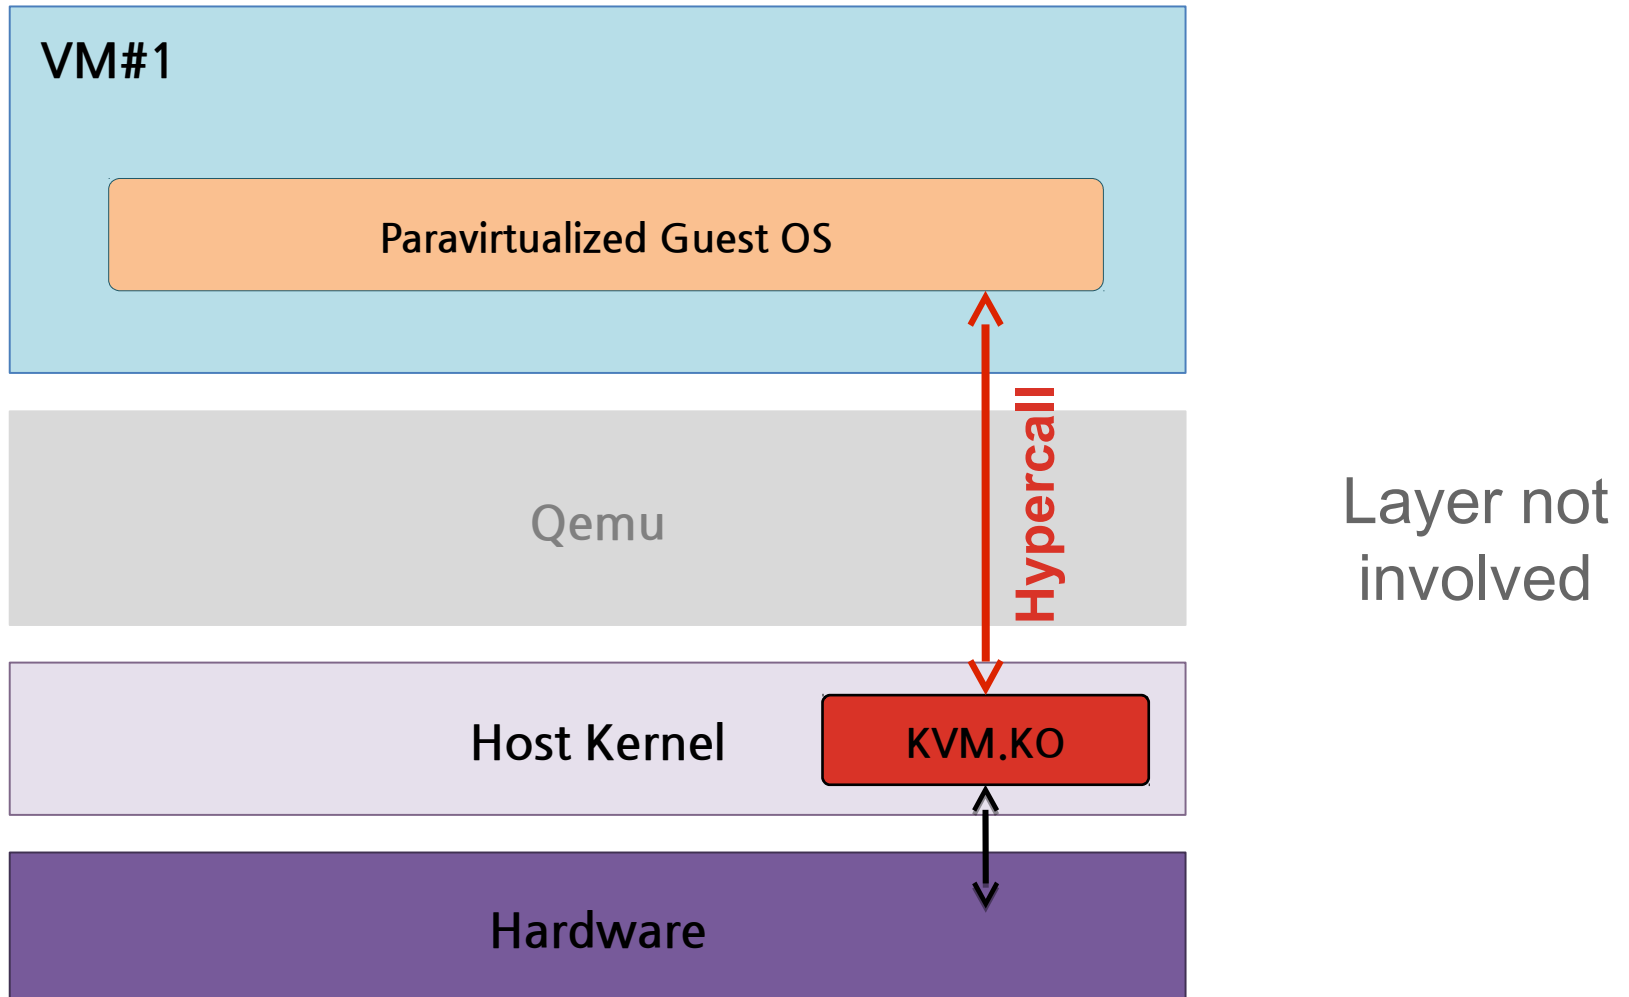
---



# Boot-up process

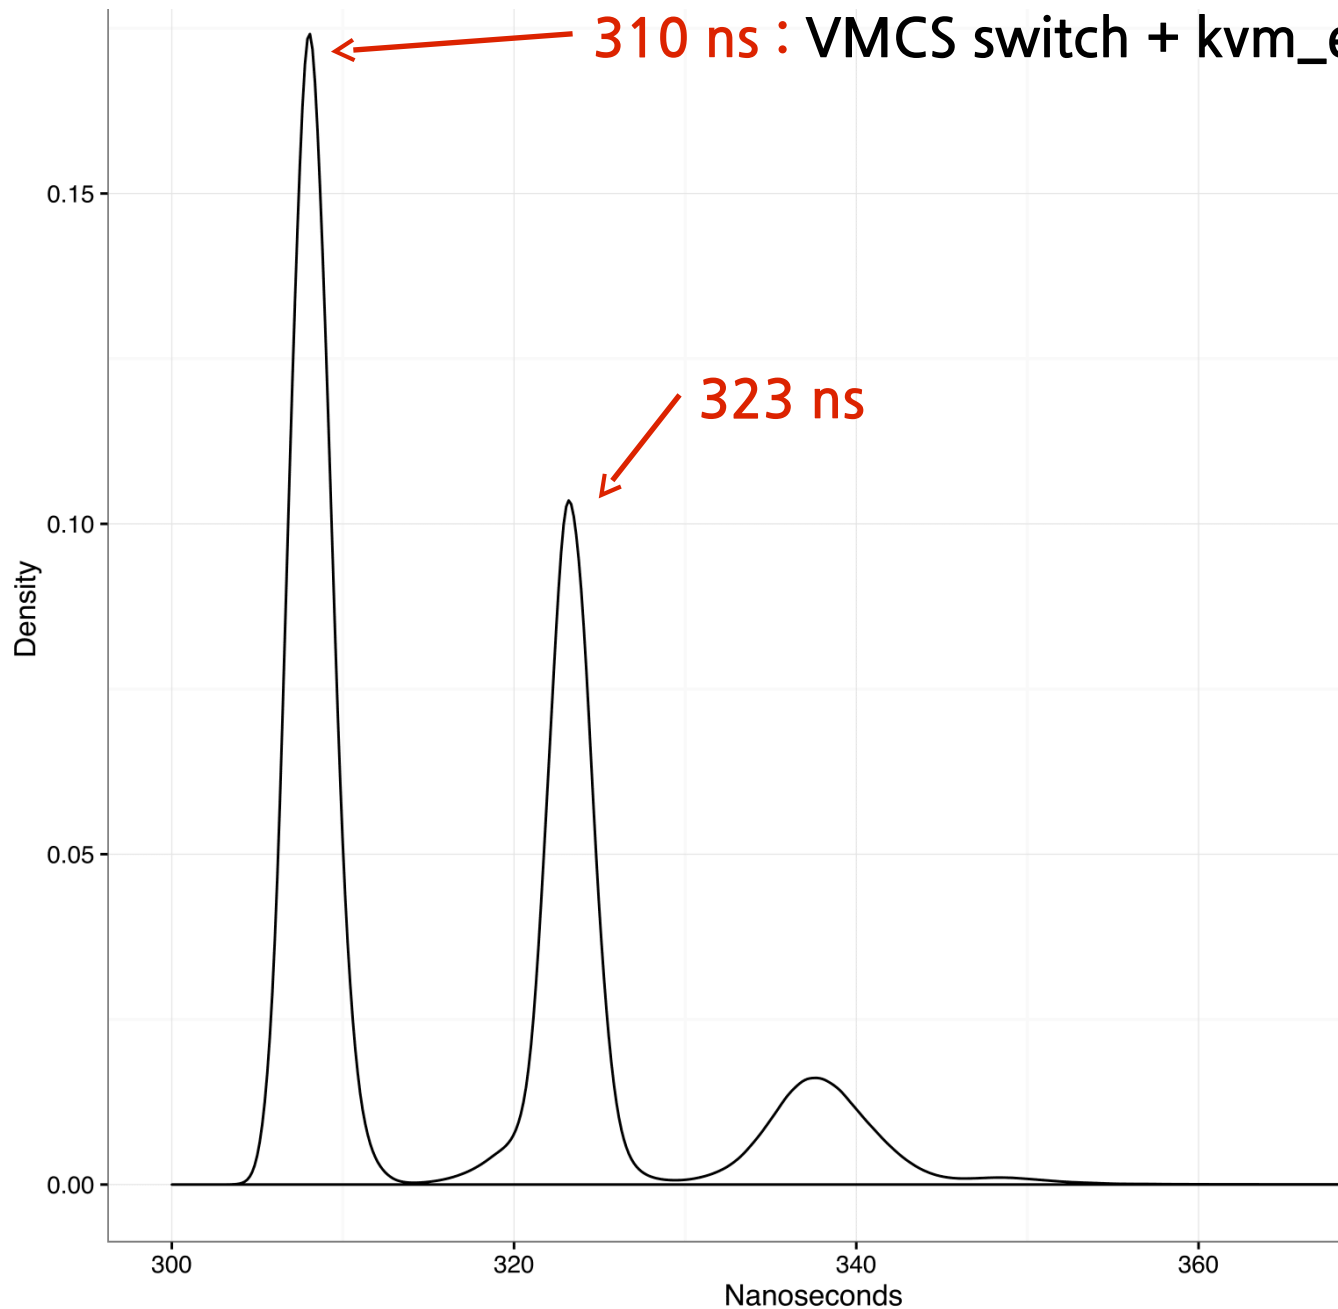


# Tracing through para-virtualization layers



# Micro-Benchmarks

## Hypercall Overhead

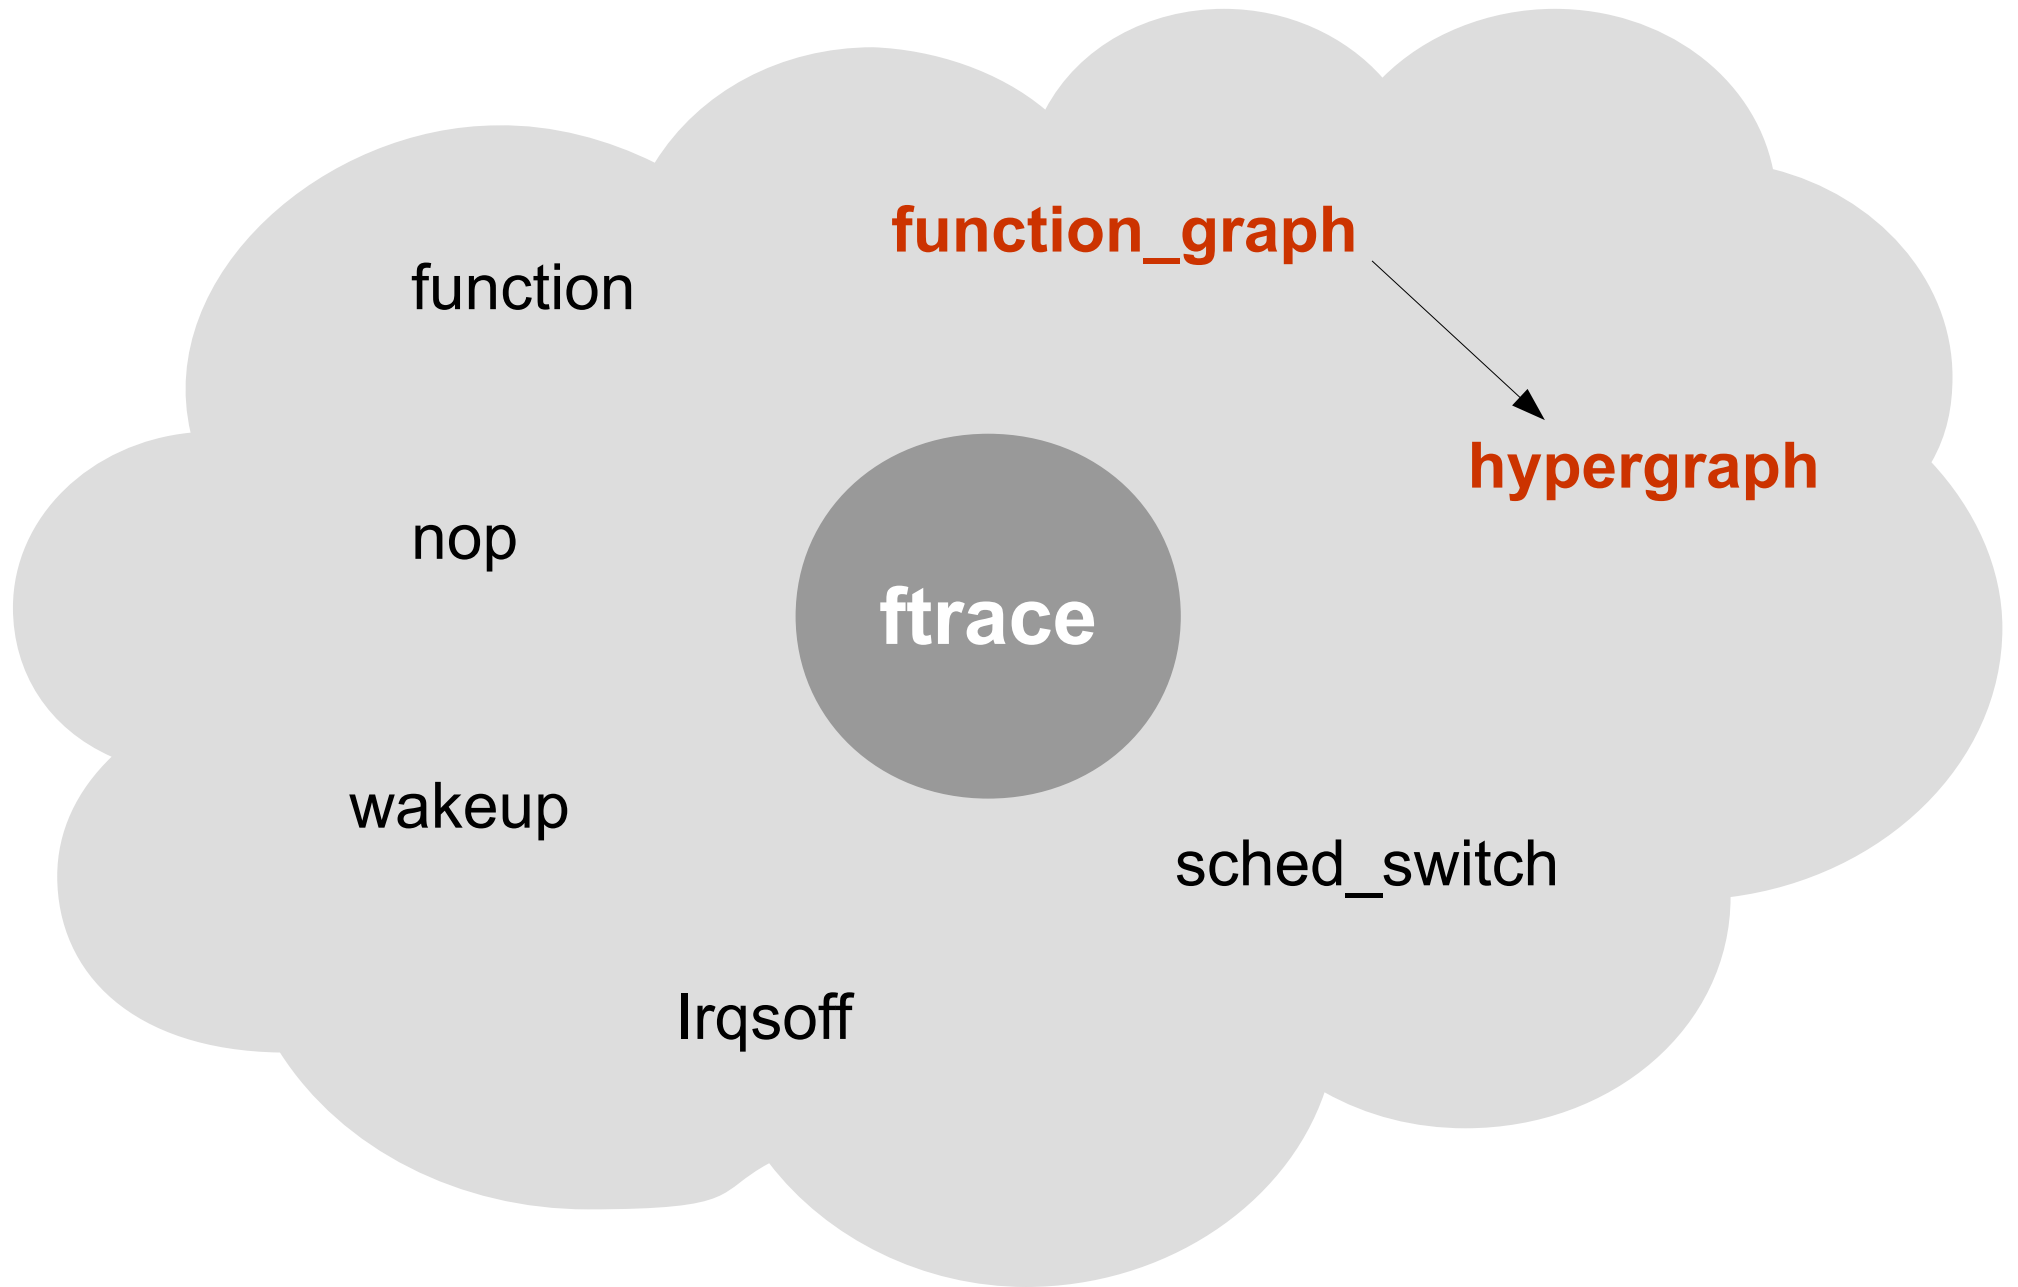


Mean	340 ns
Median	310 ns
Std deviation	833 ns

Hardware : Intel(R) Core(TM) i7-6700K CPU @ 4.00GHz  
(x86\_64)

# Implementation

---





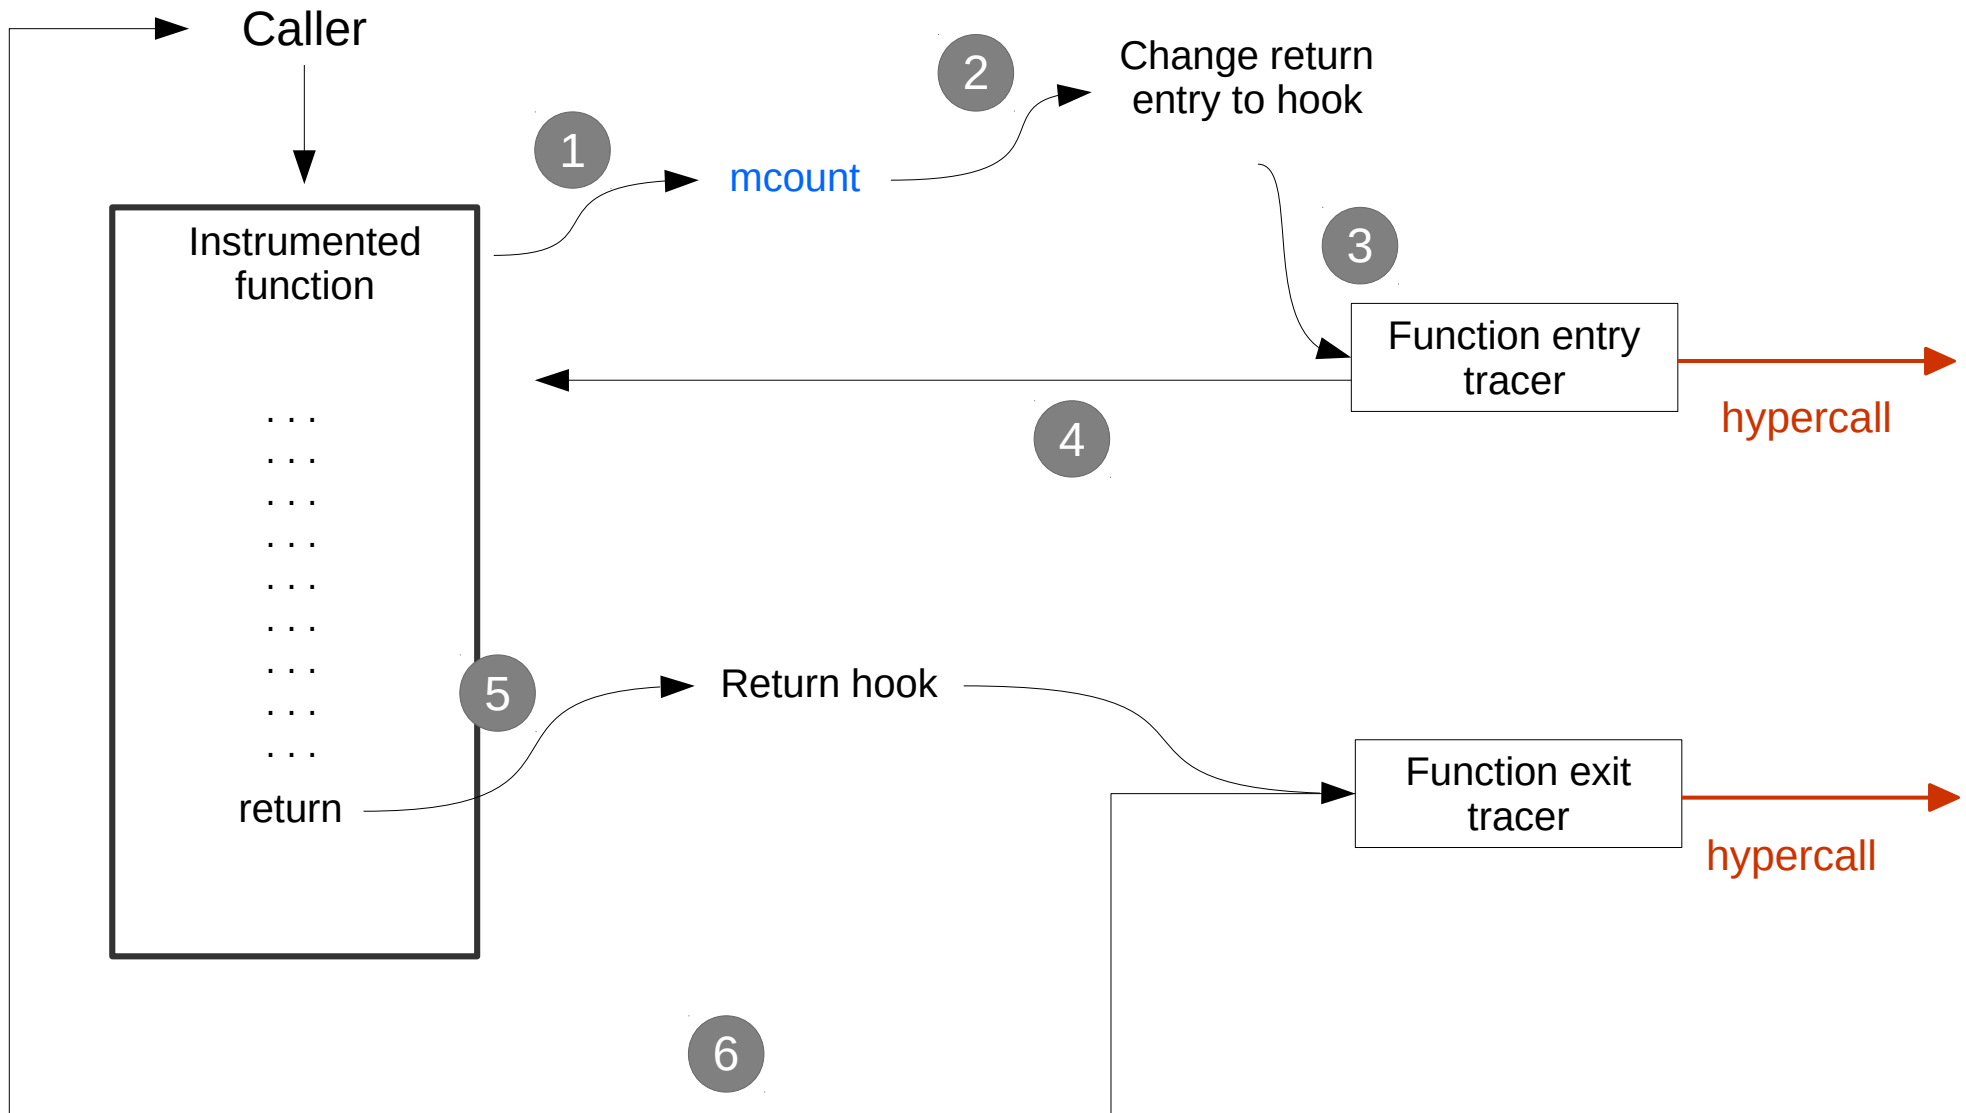
# Implementation

---


## Tracing through Hypercalls

- Hook to **Ftrace** “function graph” probe for entry & exit
- Send function address
- Tracing only the hypervisor
- Use only host Timestamp
  - > No synchronization is required
- Resolve addresses using guest kernel symbols

# Ftrace : Mcount and hypercall



# Depth Overhead Analysis




# Investigating Boot-up issues


Kernel Crash/Panic  
Latencies

# Boot-up issues

---



# Boot-up issues




Hypercall overhead

Entry

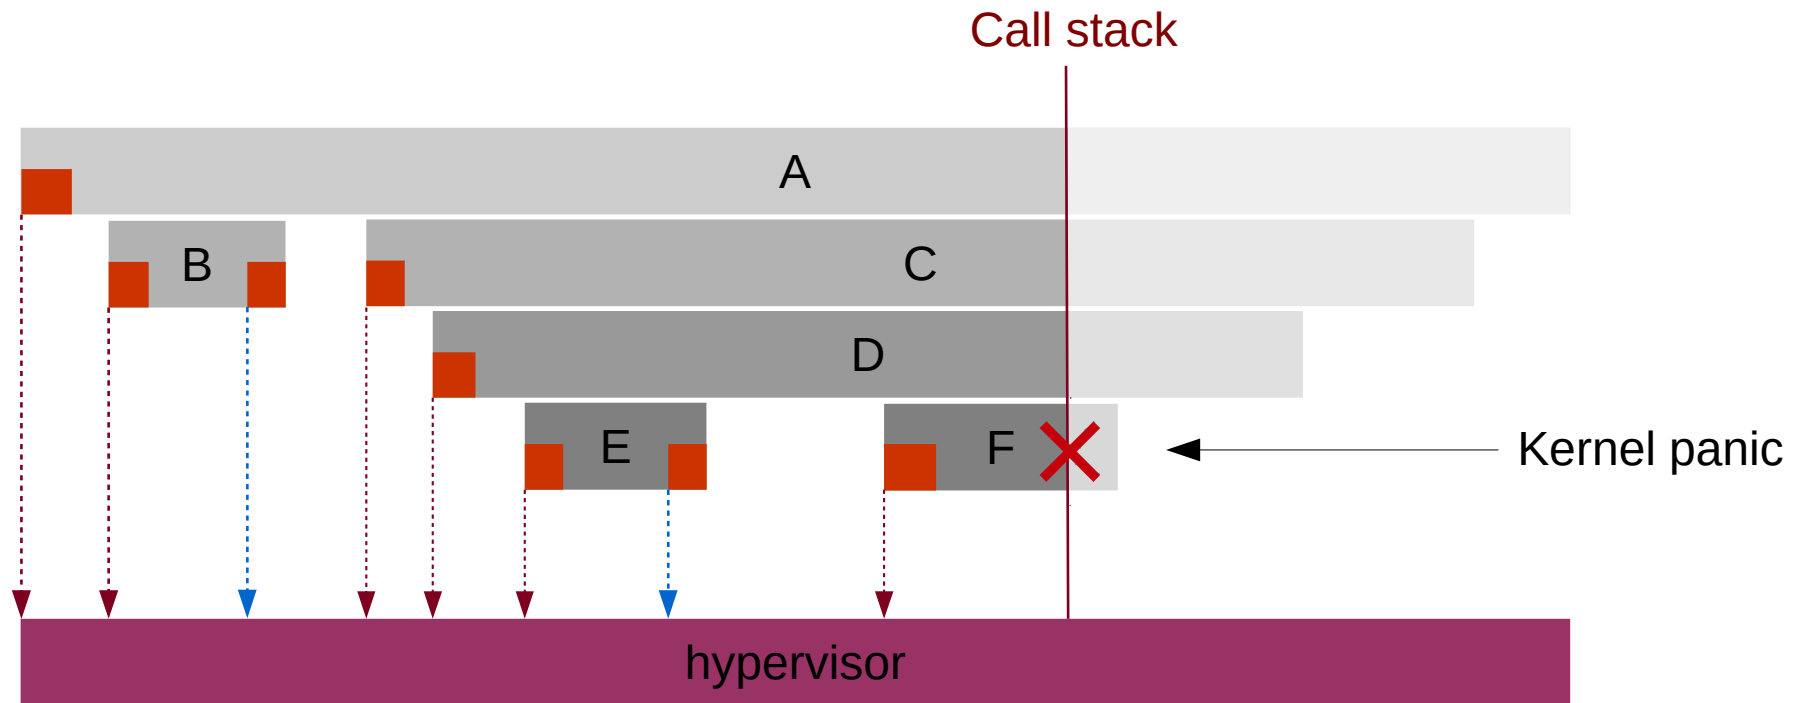
Exit

Received sequence : A-B-C-D-E-F



# Boot-up issues

Kernel crash : Oops, panic, ... etc.





Hypercall overhead

Entry

Exit

Received sequence : A-B-C-D-E-F



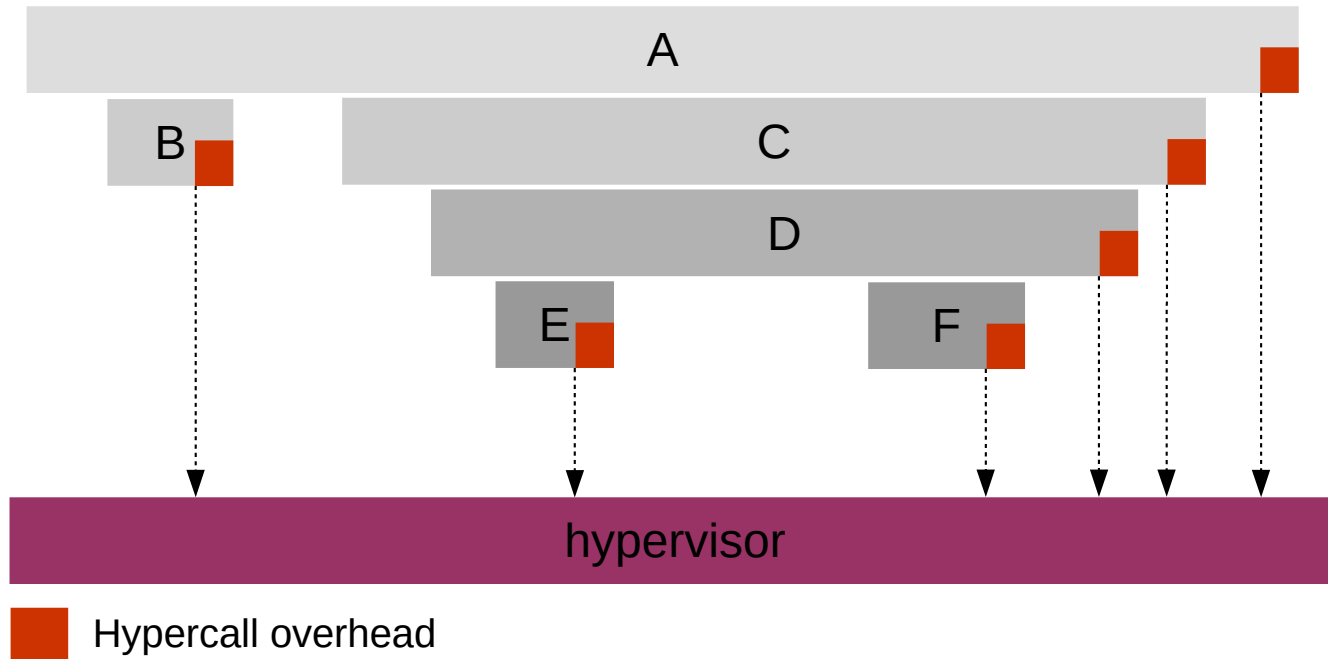


- Reduce by half the trace events
- Better for Performance Analysis




# Boot-up issues

## Post-order tree traversal

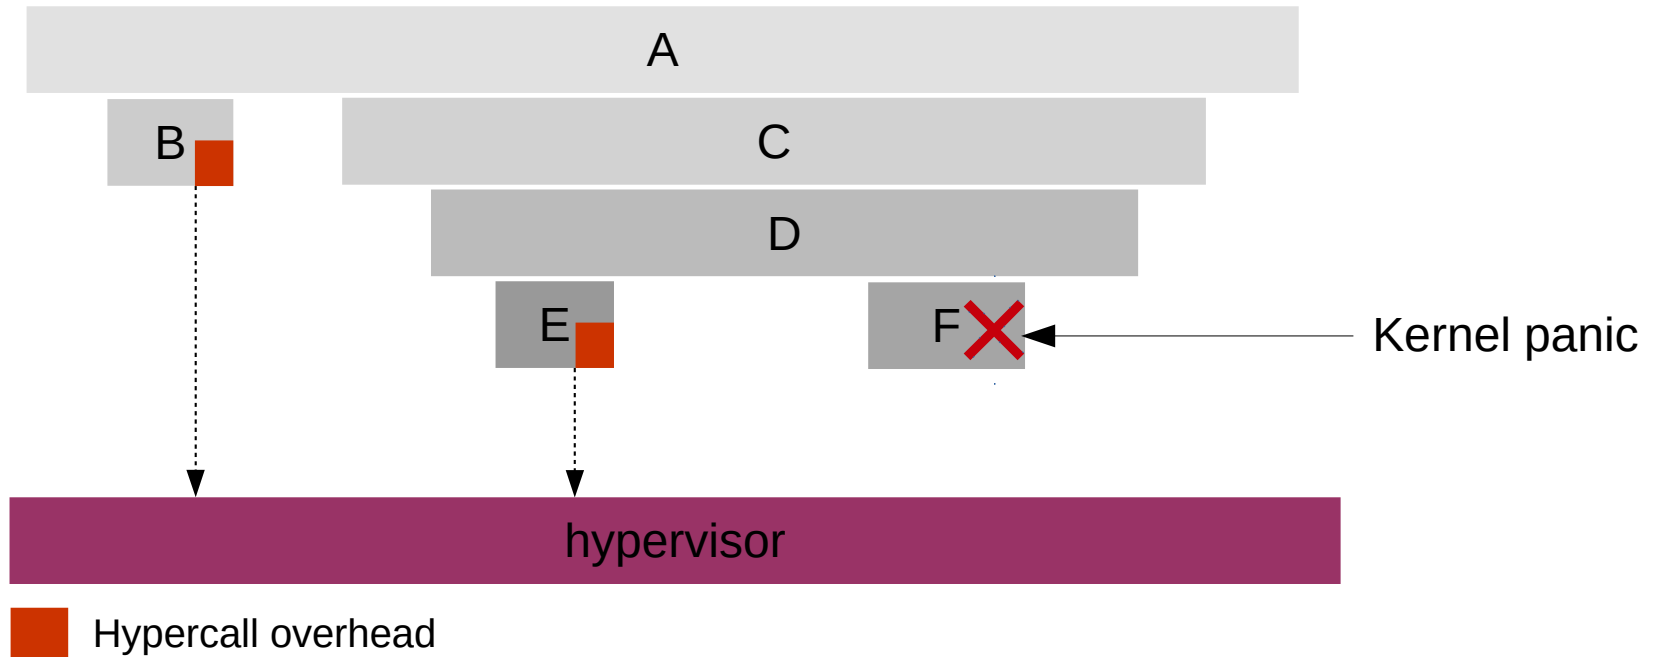


Received sequence : B-E-F-D-C-A




# Boot-up issues

Kernel crash : Oops, panic, ... etc.



Received sequence : B-E



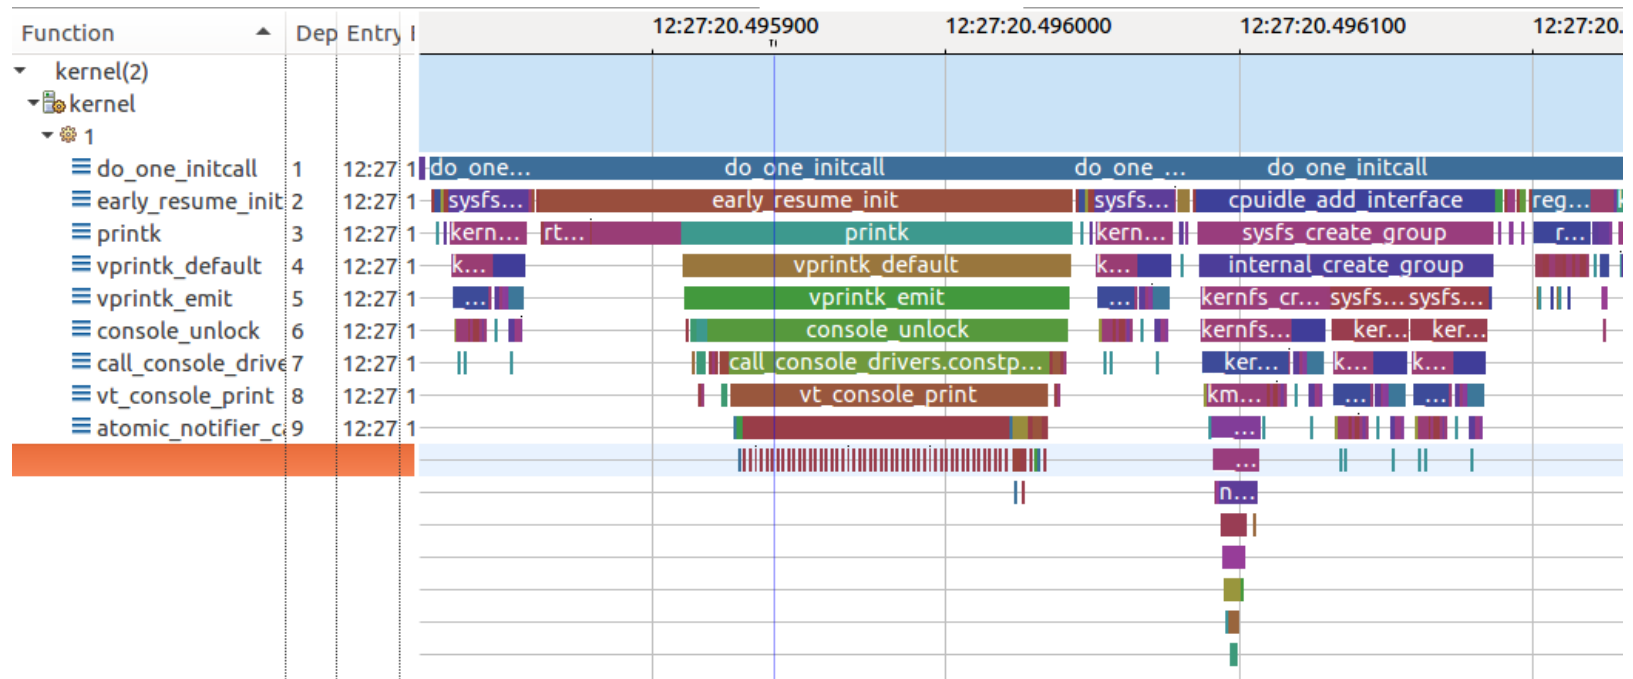
→ Not good for debugging kernel crash



# Visualization



## Graph of Guest Kernel space




*Thanks to Geneviève*



# Questions

[abderrahmane.benbachir@polymtl.ca](mailto:abderrahmane.benbachir@polymtl.ca)

<https://github.com/abenbachir>



# References

---

Hypercall Implementation : <https://gist.github.com/abenbachir/344822b5ba9fc5ac384cdec3f087e018>

QEMU Hypertrace Patches: <http://patchwork.ozlabs.org/project/qemu-devel/list/?state=&q=Hypertrace&archive>

Trampoline: <https://www.kernel.org/doc/ols/2009/ols2009-pages-47-54.pdf>

Callstack xml analysis: <https://gist.github.com/abenbachir/e813790f183945b6f74dc74ecee57c75>

