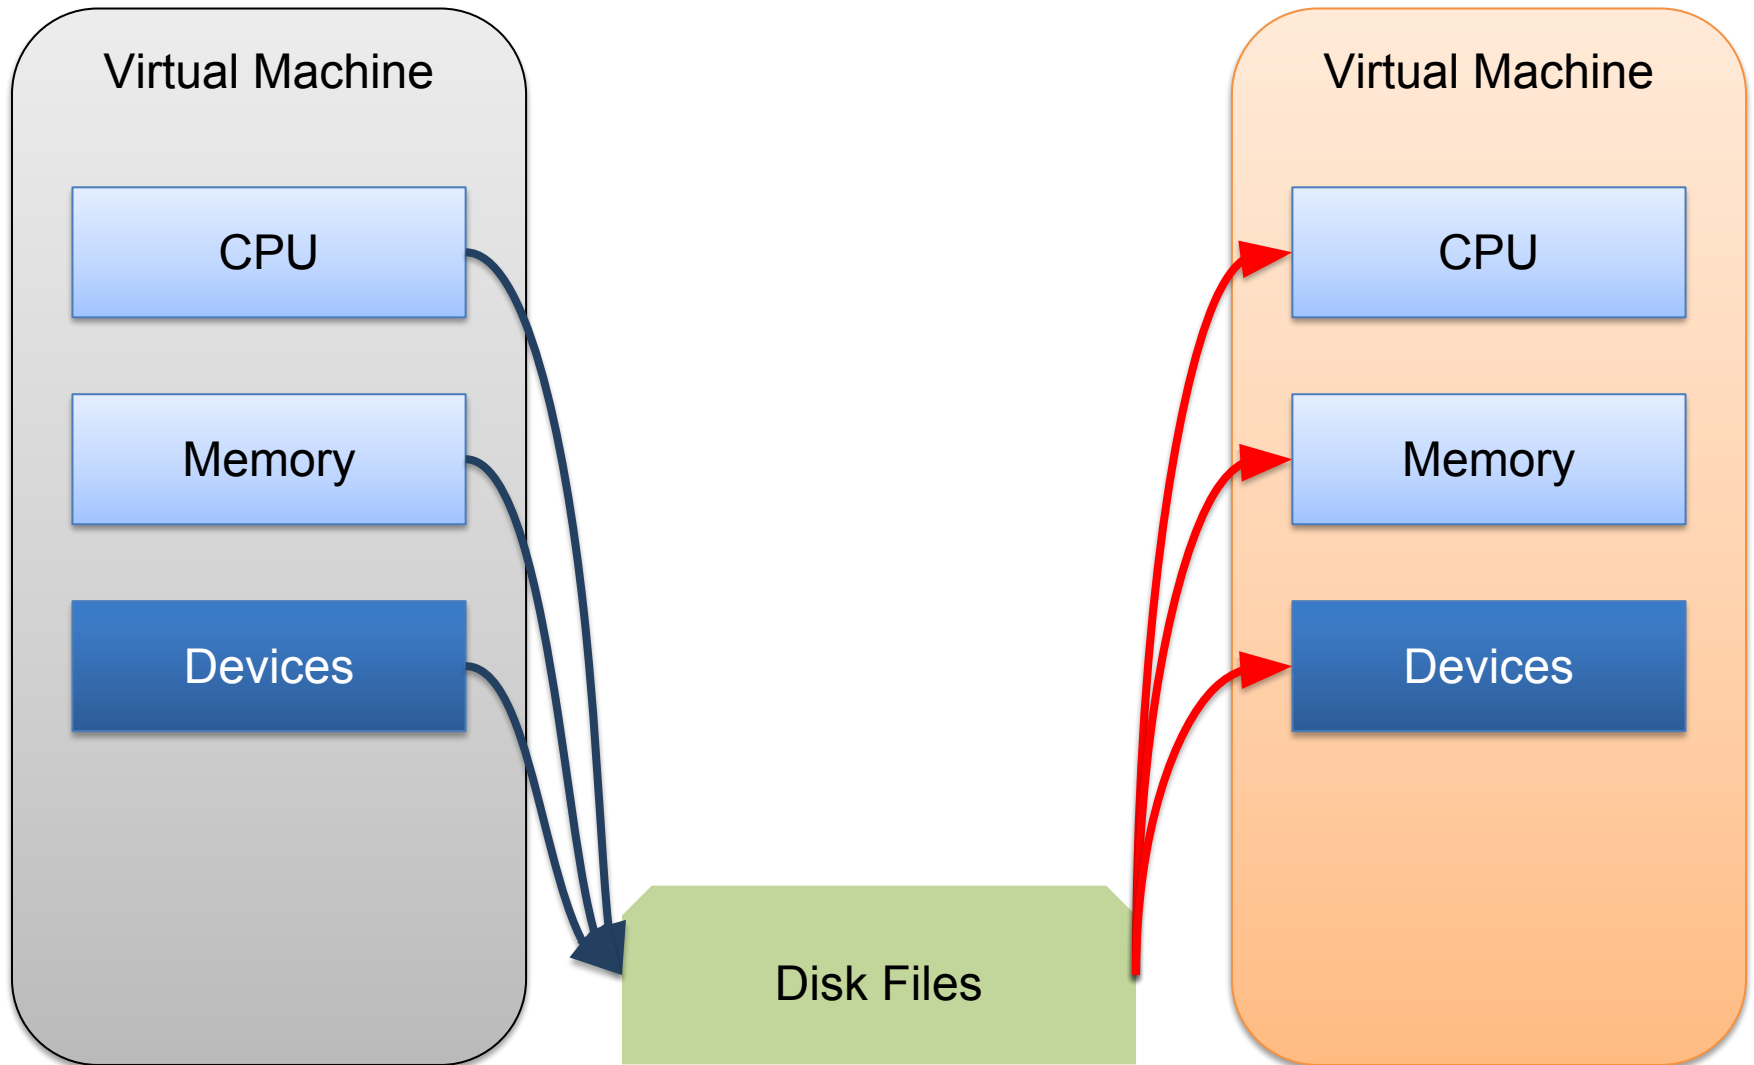


Virtual Machine Snapshot Improvements Timekeeping

Presented by
Darius Mihai
dariusmihaim@gmail.com

bhyve – VM Snapshot




Timekeeping issues - Linux

```
a@test-server:~$ while sleep 1; do date; done  
Thu Jan 17 20:35:29 UTC 2019  
Thu Jan 17 20:35:30 UTC 2019  
Thu Jan 17 20:35:31 UTC 2019  
Thu Jan 17 20:35:32 UTC 2019
```

Snapshot taken

VM restored

Timekeeping issues - Windows



Timers

- Used by OS for periodic tasks

```
shell# sleep $N
```



<https://en.wikipedia.org/wiki/Timer>

LAPIC Virtualization - Snapshot

- Save Current Count Register
- Reprogram callouts at reset time

HPET Virtualization - Snapshot

- Save the value of the counter
- Add offset to the current value of the counter

Clocksource

- Used by the OS to measure time
- Monotonic counter

Time Stamp Counter

- Per-CPU register
- Incremented at CPU core frequency
- Incremented at a static rate

Time Stamp Counter Virtualization

- Shared with host
- Intel/AMD CPUs have a TSC Offset register
- $TSC_guest = TSC_system + TSC_offset$

Timekeeping Virtualization - Results

- VM does not freeze or complain about time

Special Thanks

- Matthew Grooms
- iXsystems
- Marcelo Araujo

Online Resources

[https://github.com/FreeBSD-UPB/freebsd/](https://github.com/FreeBSD-UPB/freebsd/-/projects/bhyve_snapshot)
- projects/bhyve_snapshot branch

<https://reviews.freebsd.org/D19495>