More Multiprocessing on seL4: Are efficient SMP Virtual Machines Possible on Verifiable seL4 Kernels today?

Kent McLeod | seL4 Summit 2023 | Minneapolis, USA

KRY10

Last year:

Multiprocessing on seL4 with verified kernels

Kent McLeod | seL4 Summit 2022 | Munich, Germany

KRY10

Usable CPU count by kernel configuration

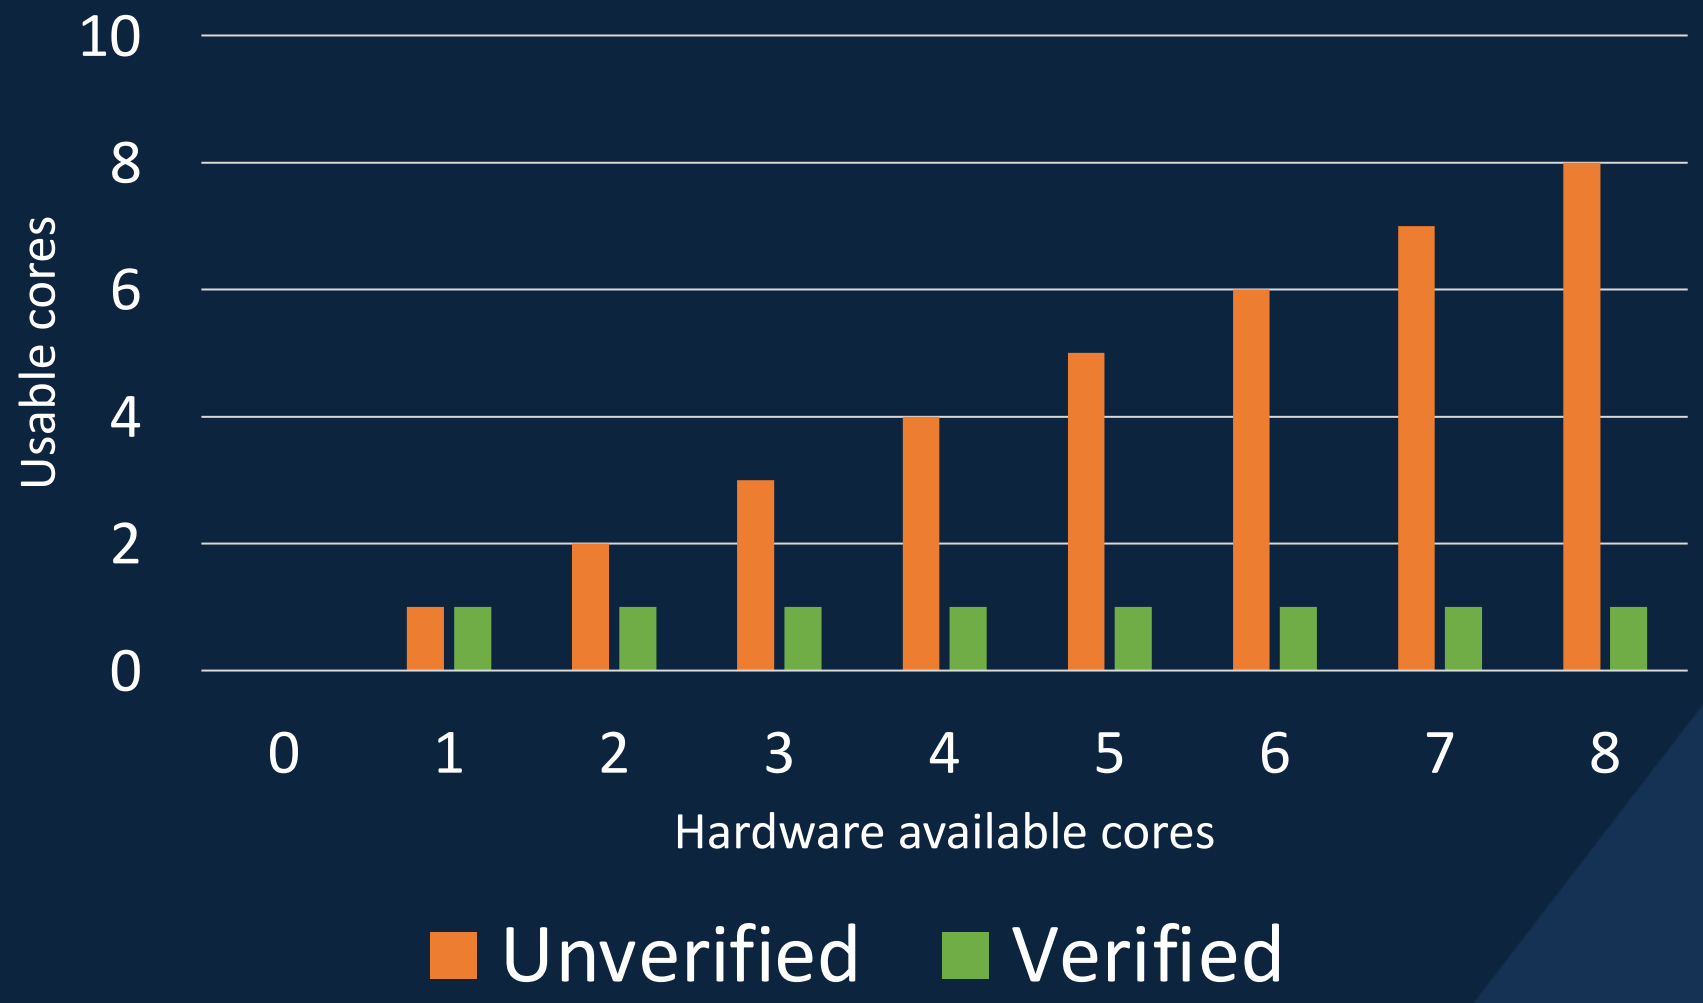

SMP seL4 configurations

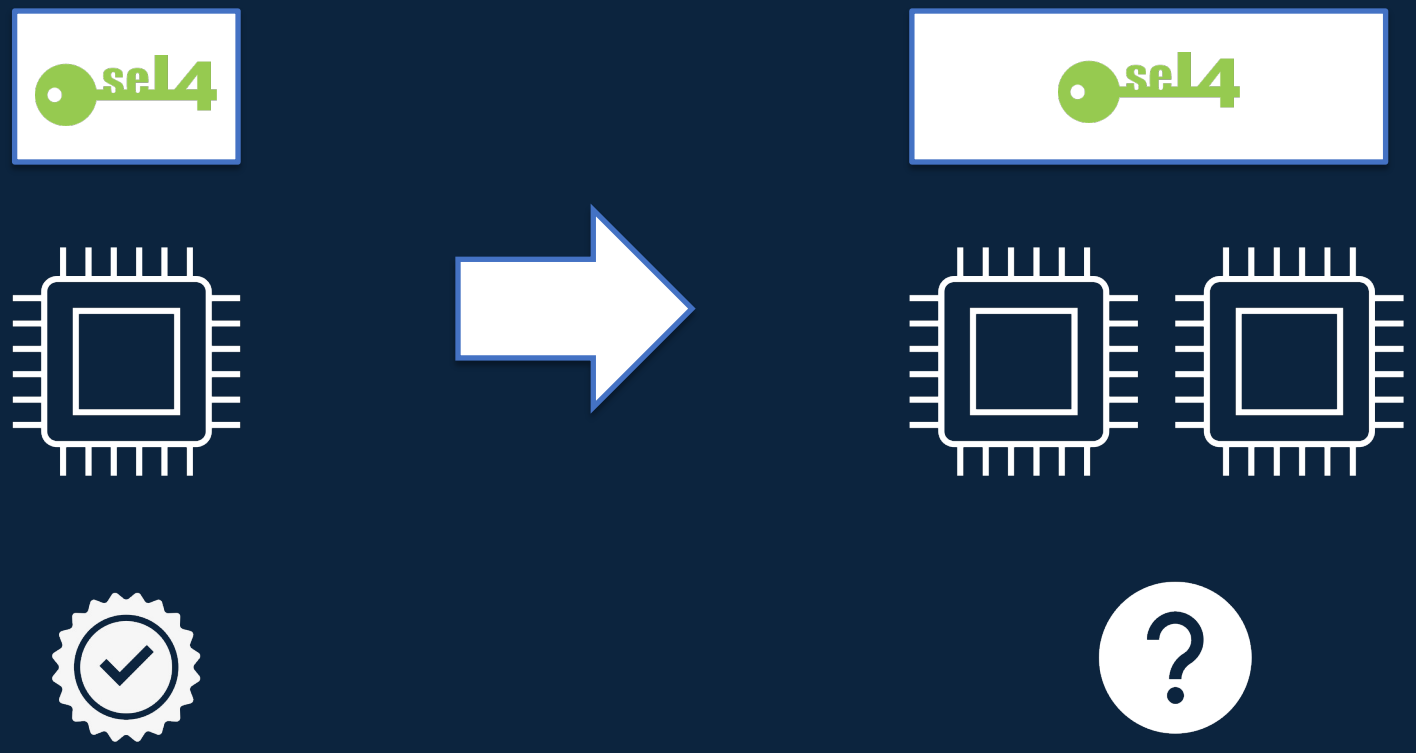

Unicore

SMP

(Re)Introducing: Partitioned multikernel

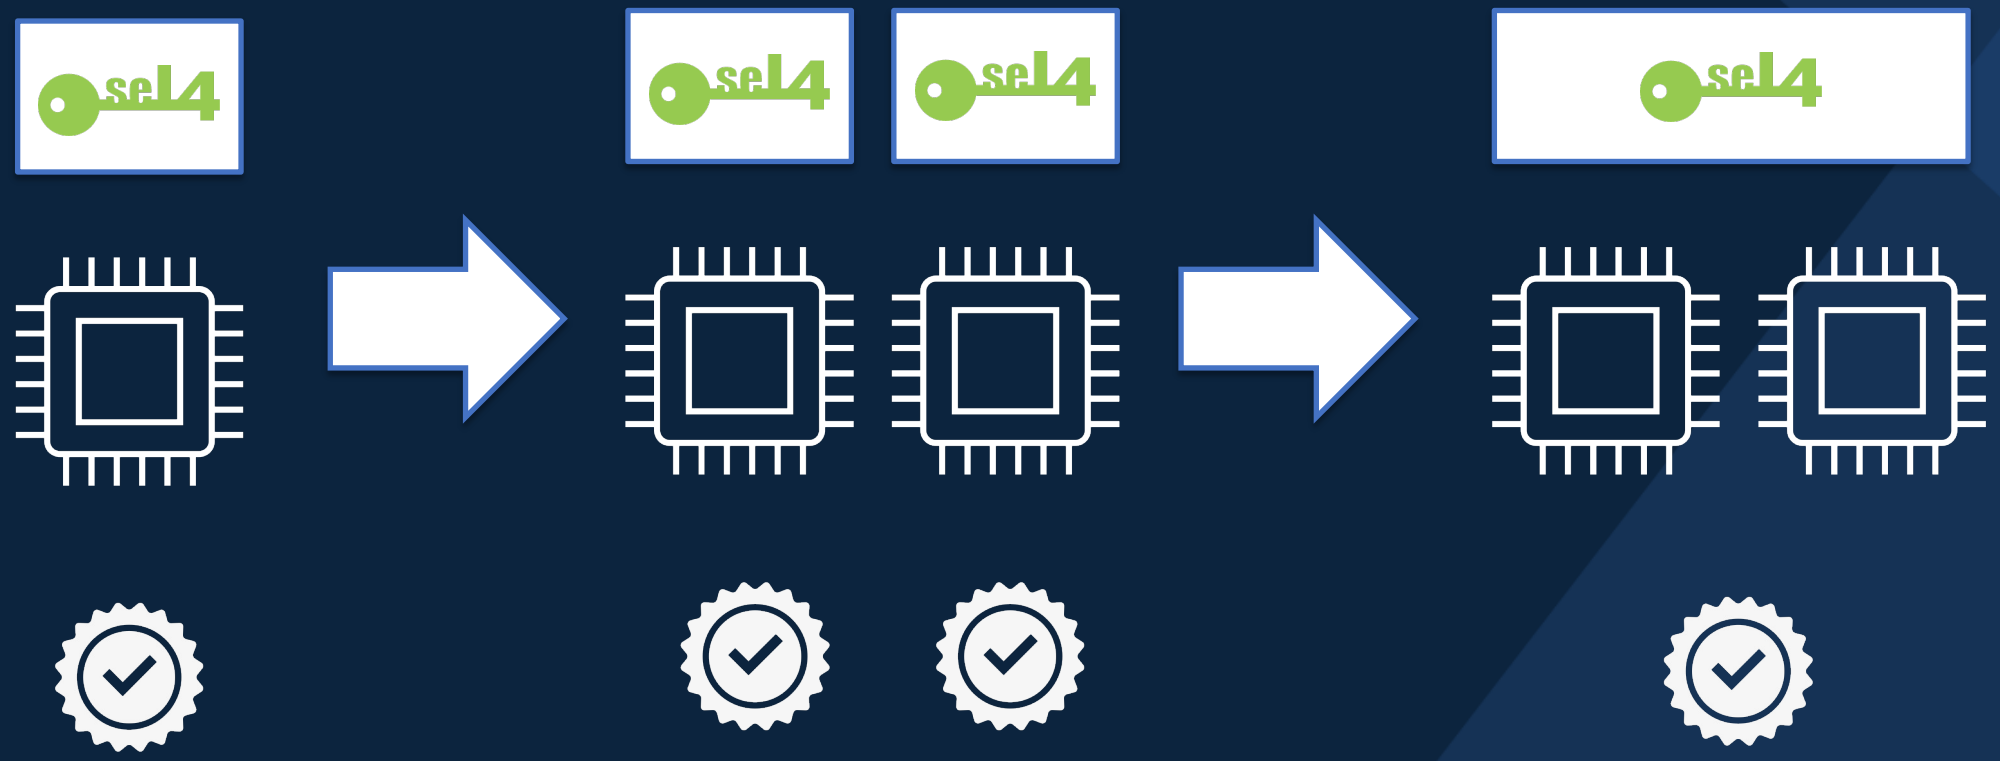

Unicore

Multikernel (AMP)

SMP

Last year:
Follow-up steps

- SMP-like user apps
- Scalable cross-core notifications
- Investigating impact of replicated data on shared caches
- Transparent cross-core seL4RPCCall CAmkES connectors
- Finish off multi-vm multicore example

Usable vCPU count by kernel configuration

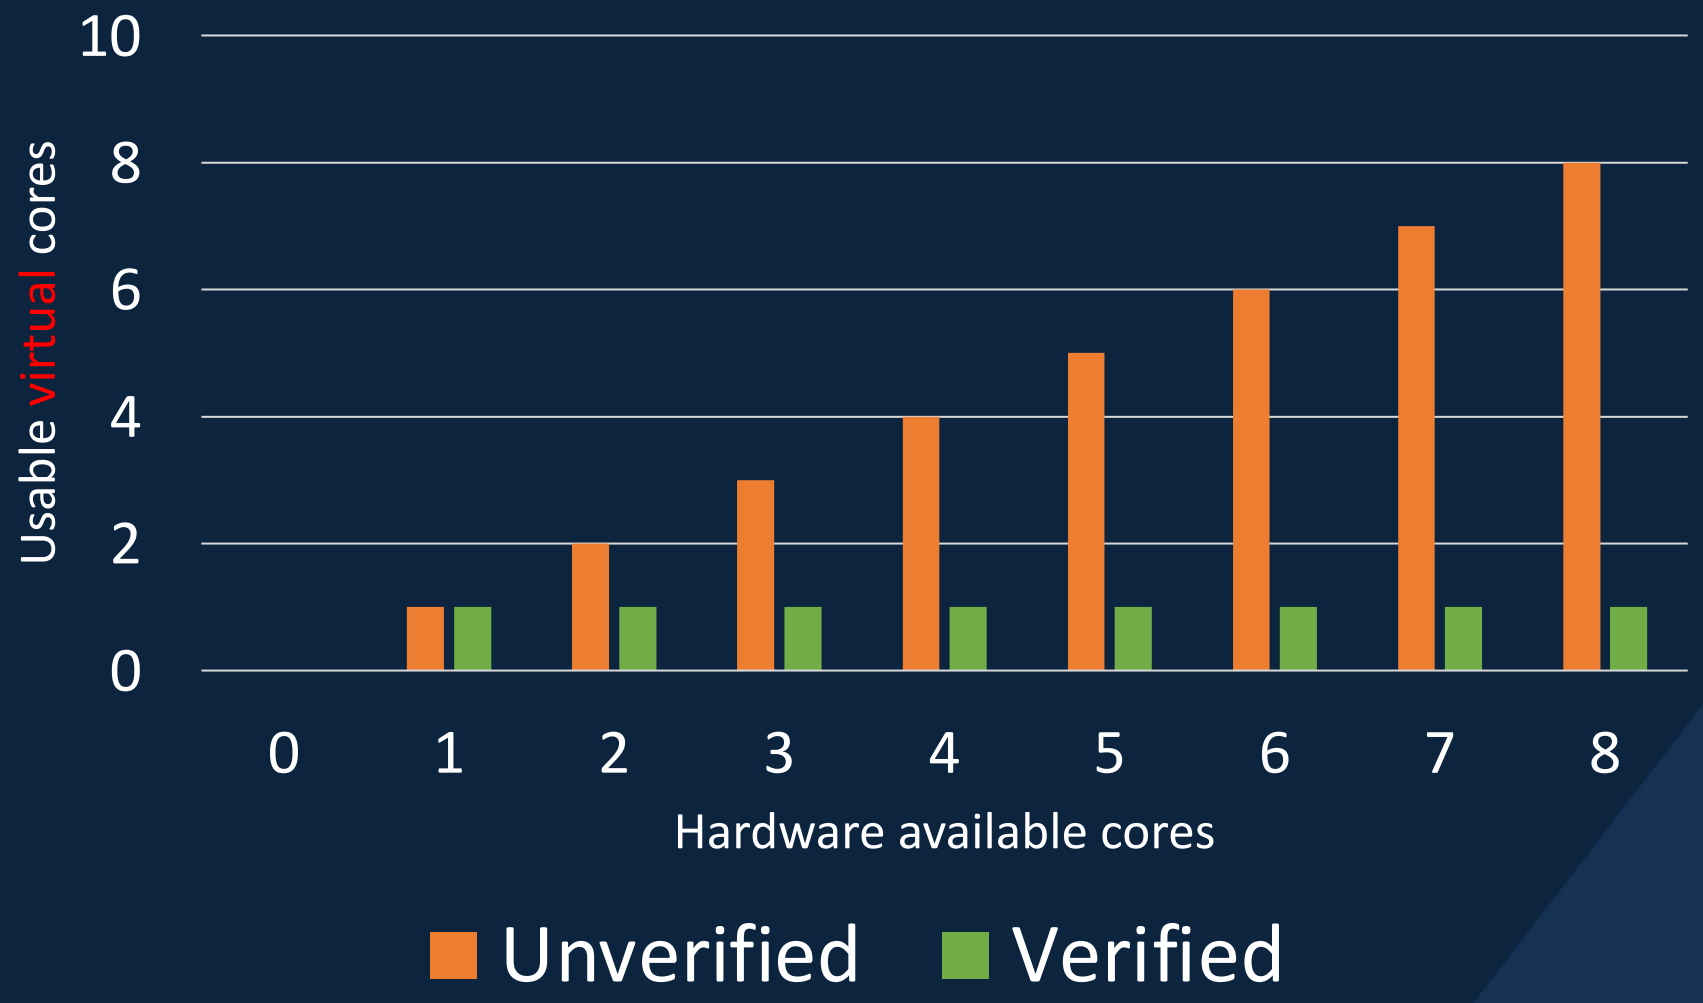

SMP guests on partitioned hosts

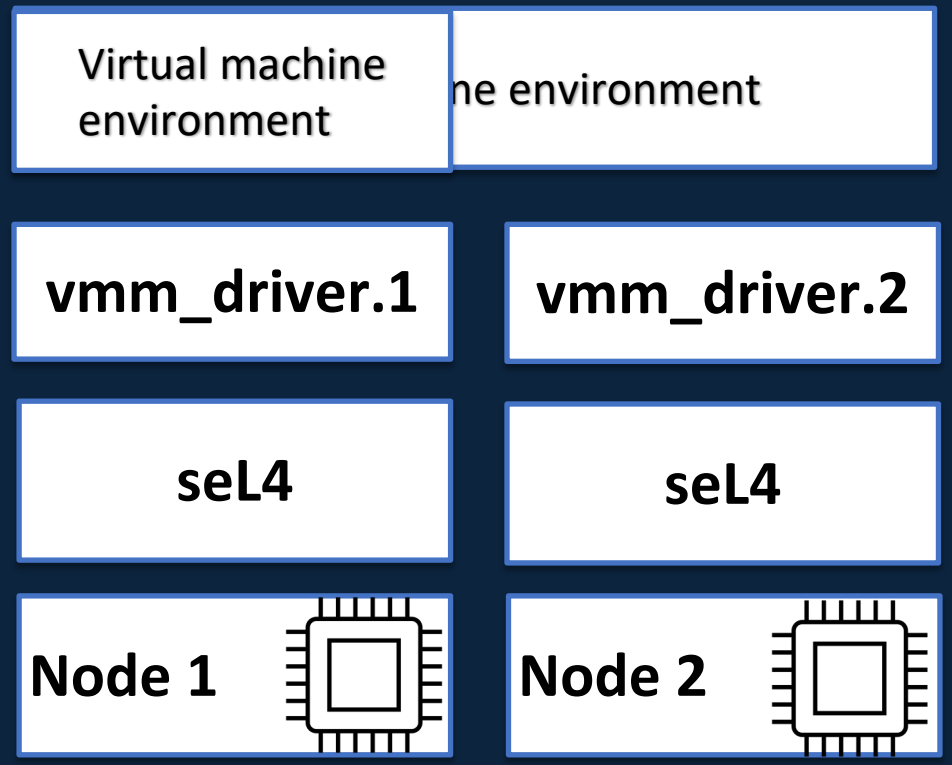

SMP guests on partitioned hosts

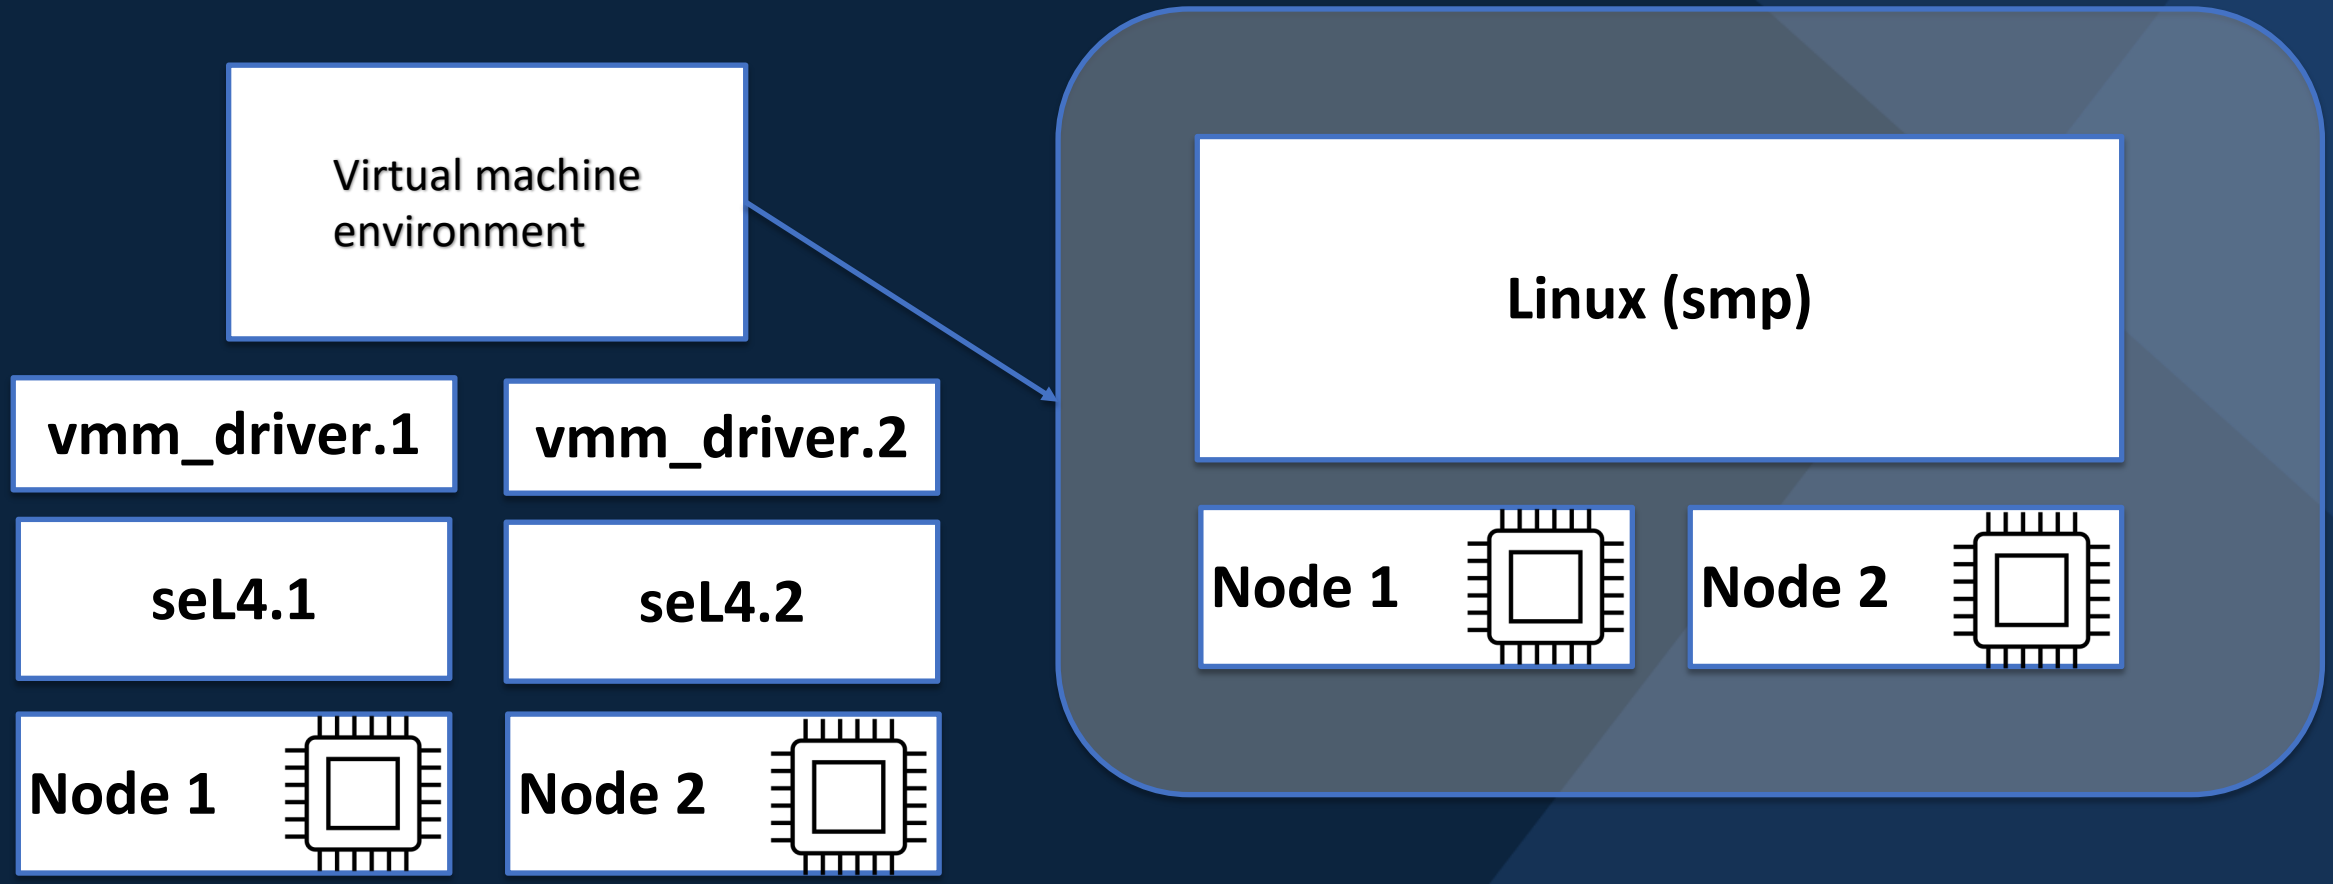

SMP overview

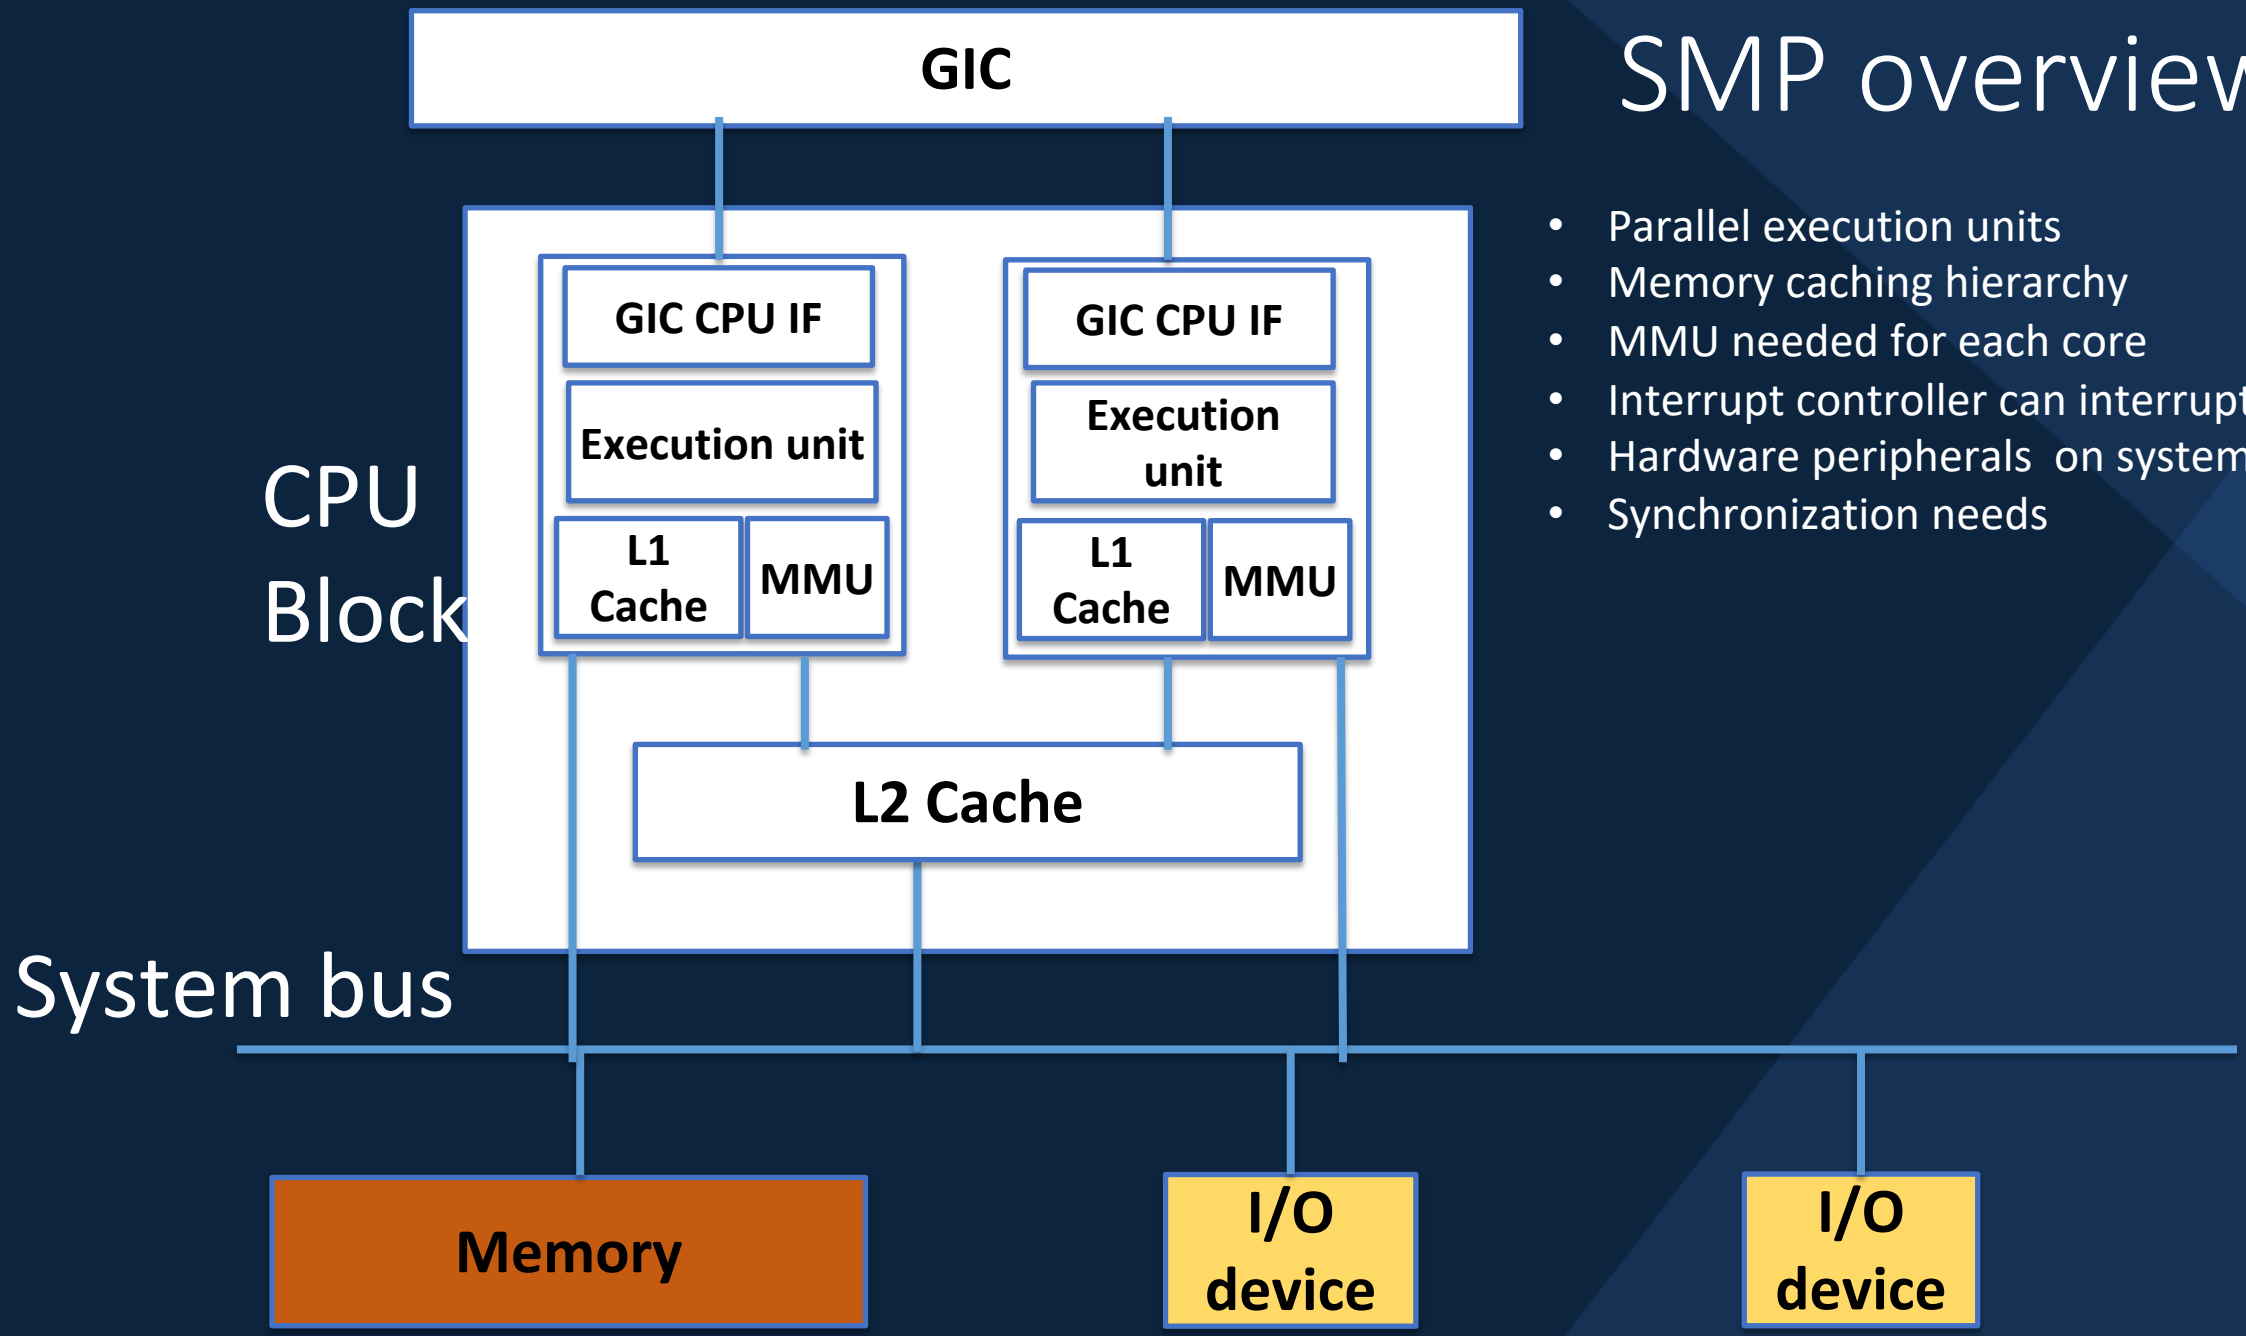

- Parallel execution units
- Memory caching hierarchy
- MMU needed for each core
- Interrupt controller can interrupt per core
- Hardware peripherals on system bus
- Synchronization needs

Arm Hyp extns.

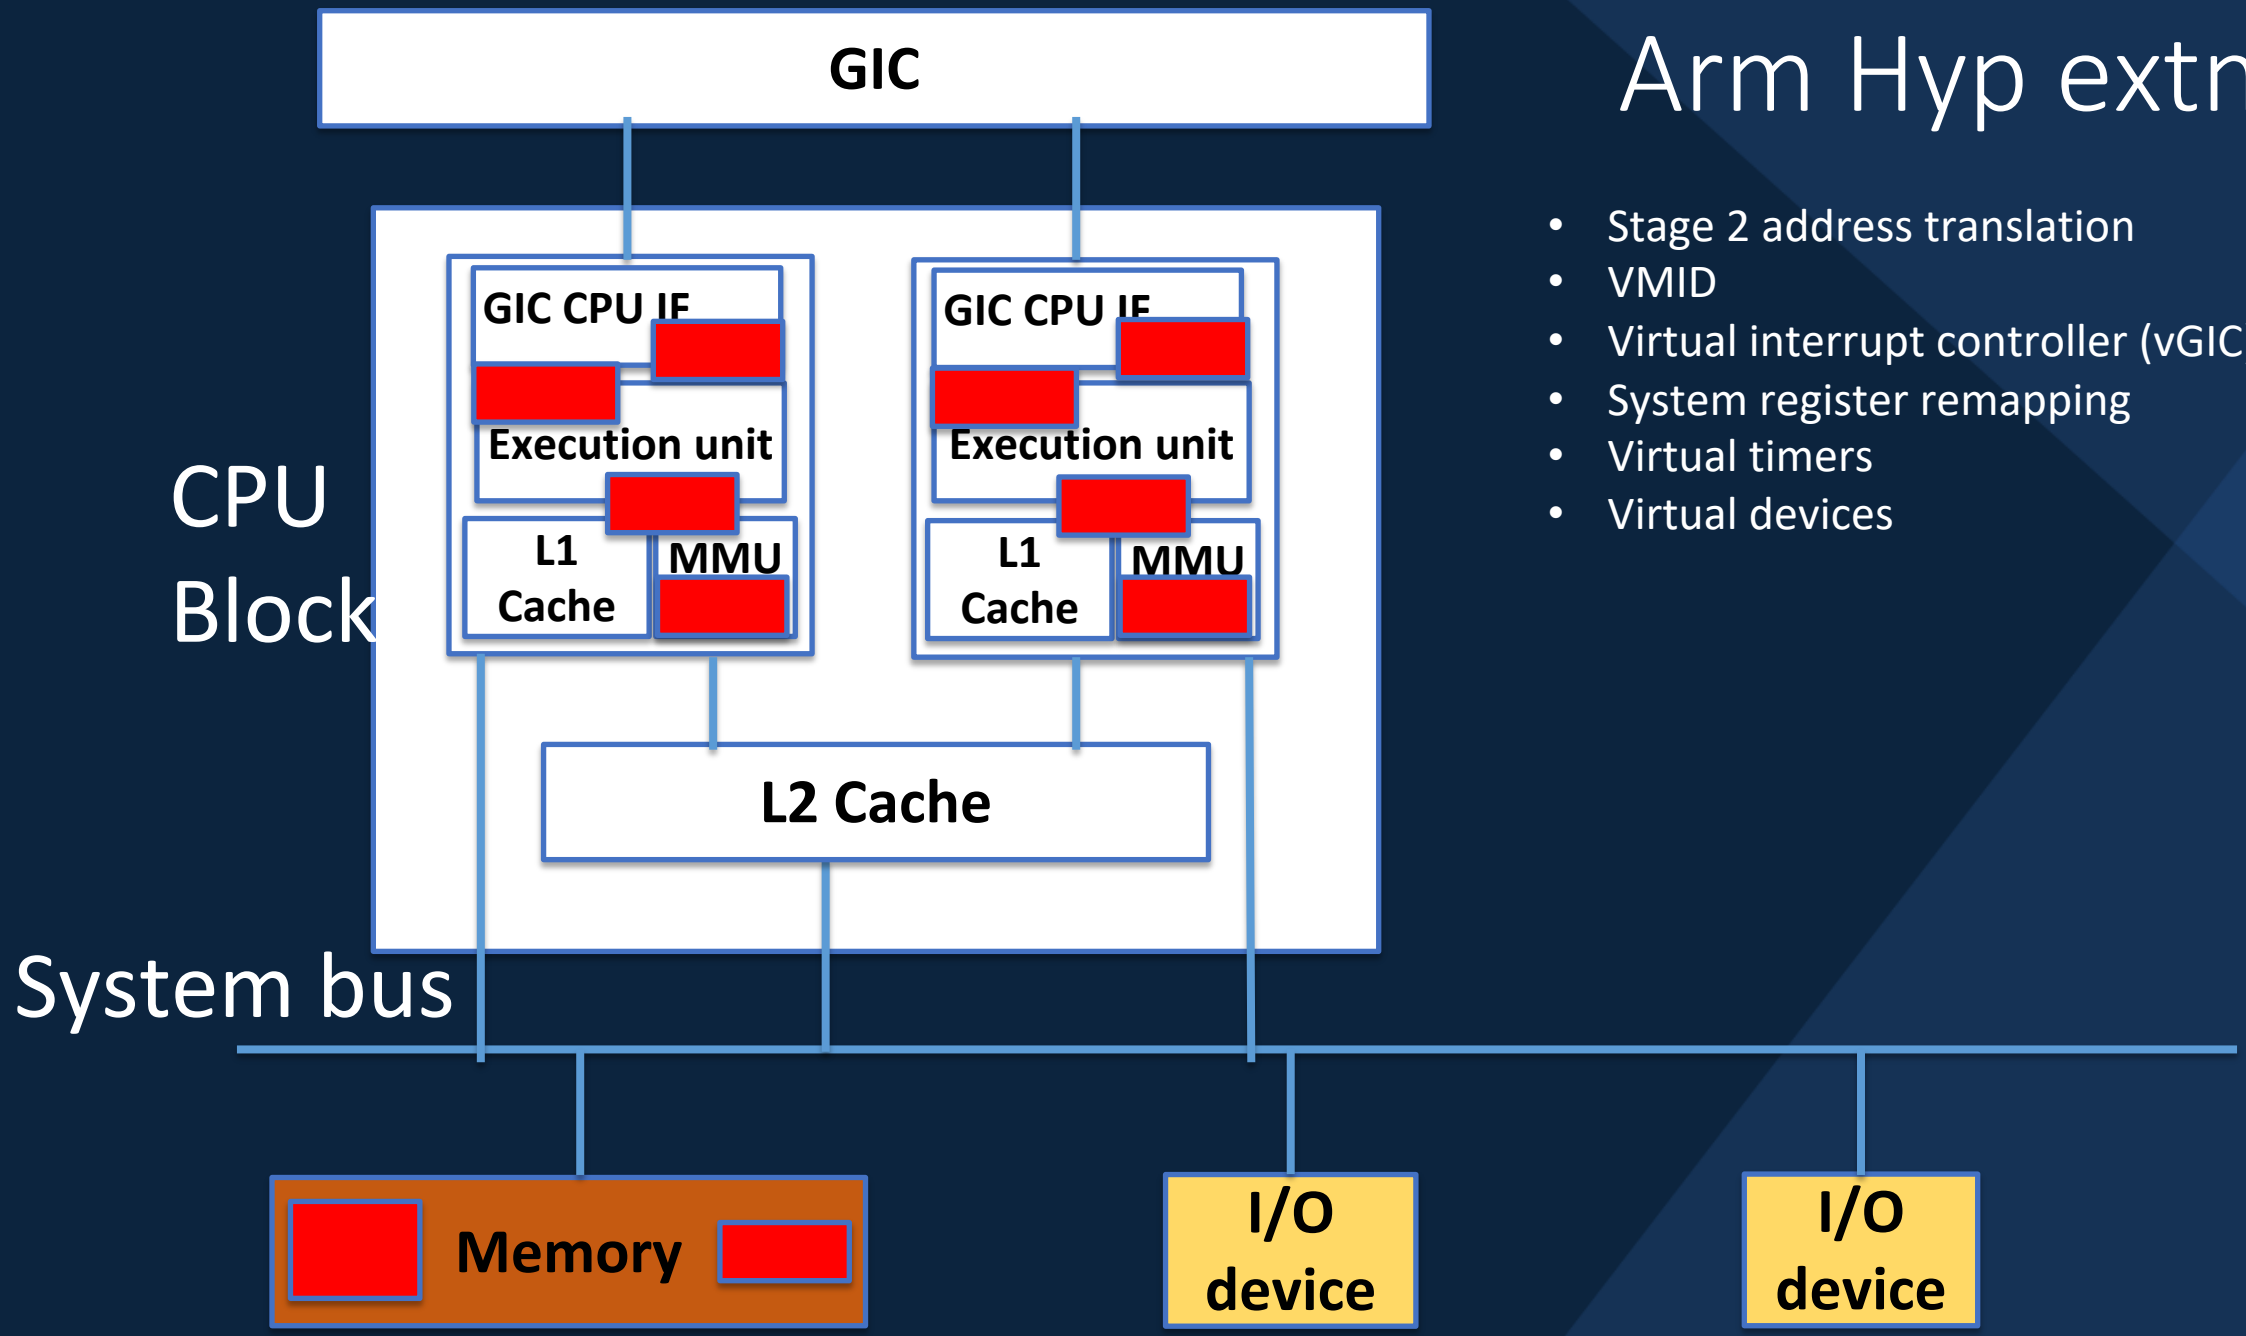

- Stage 2 address translation
- VMID
- Virtual interrupt controller (vGIC)
- System register remapping
- Virtual timers
- Virtual devices

Like a distributed system (with shared mem)

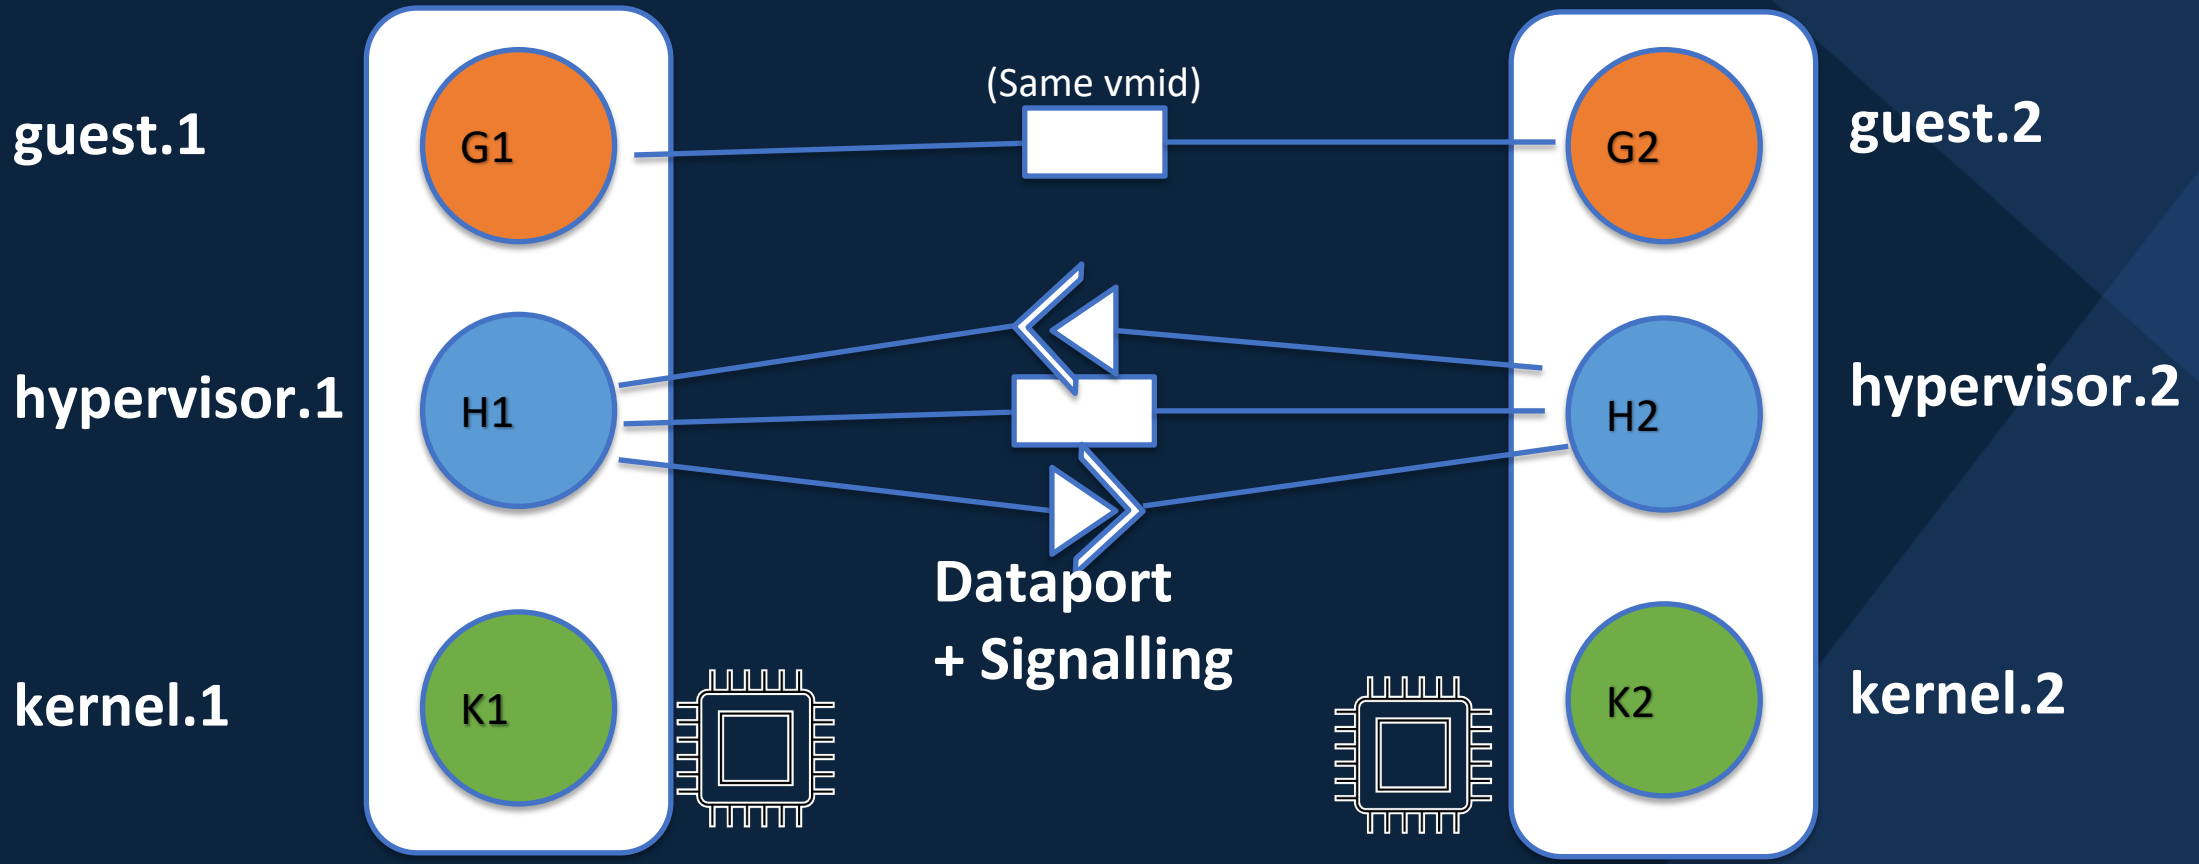

Regular operation

1.

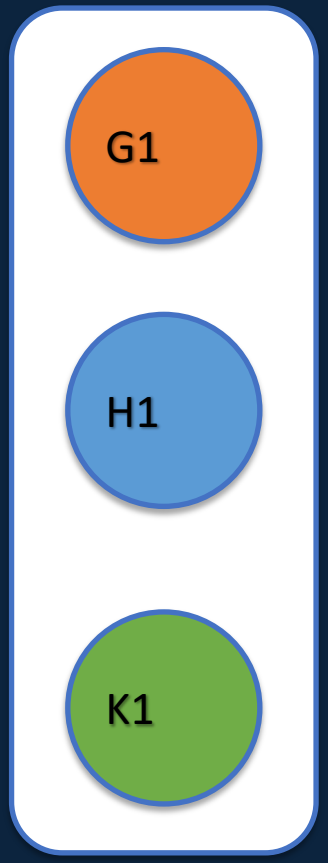

1. Both cores are running in the guest
- Regular ISA operations are largely the same
- Memory operations work the same way
- Process context switching
- TLB shootdown
- Cache invalidation
- Spin lock synchronisation

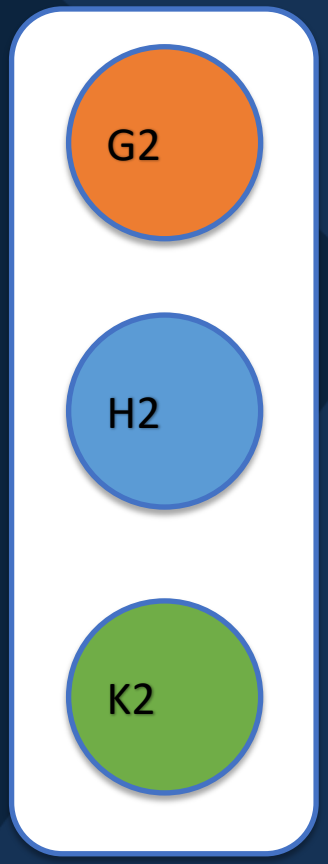

1.

Initialization

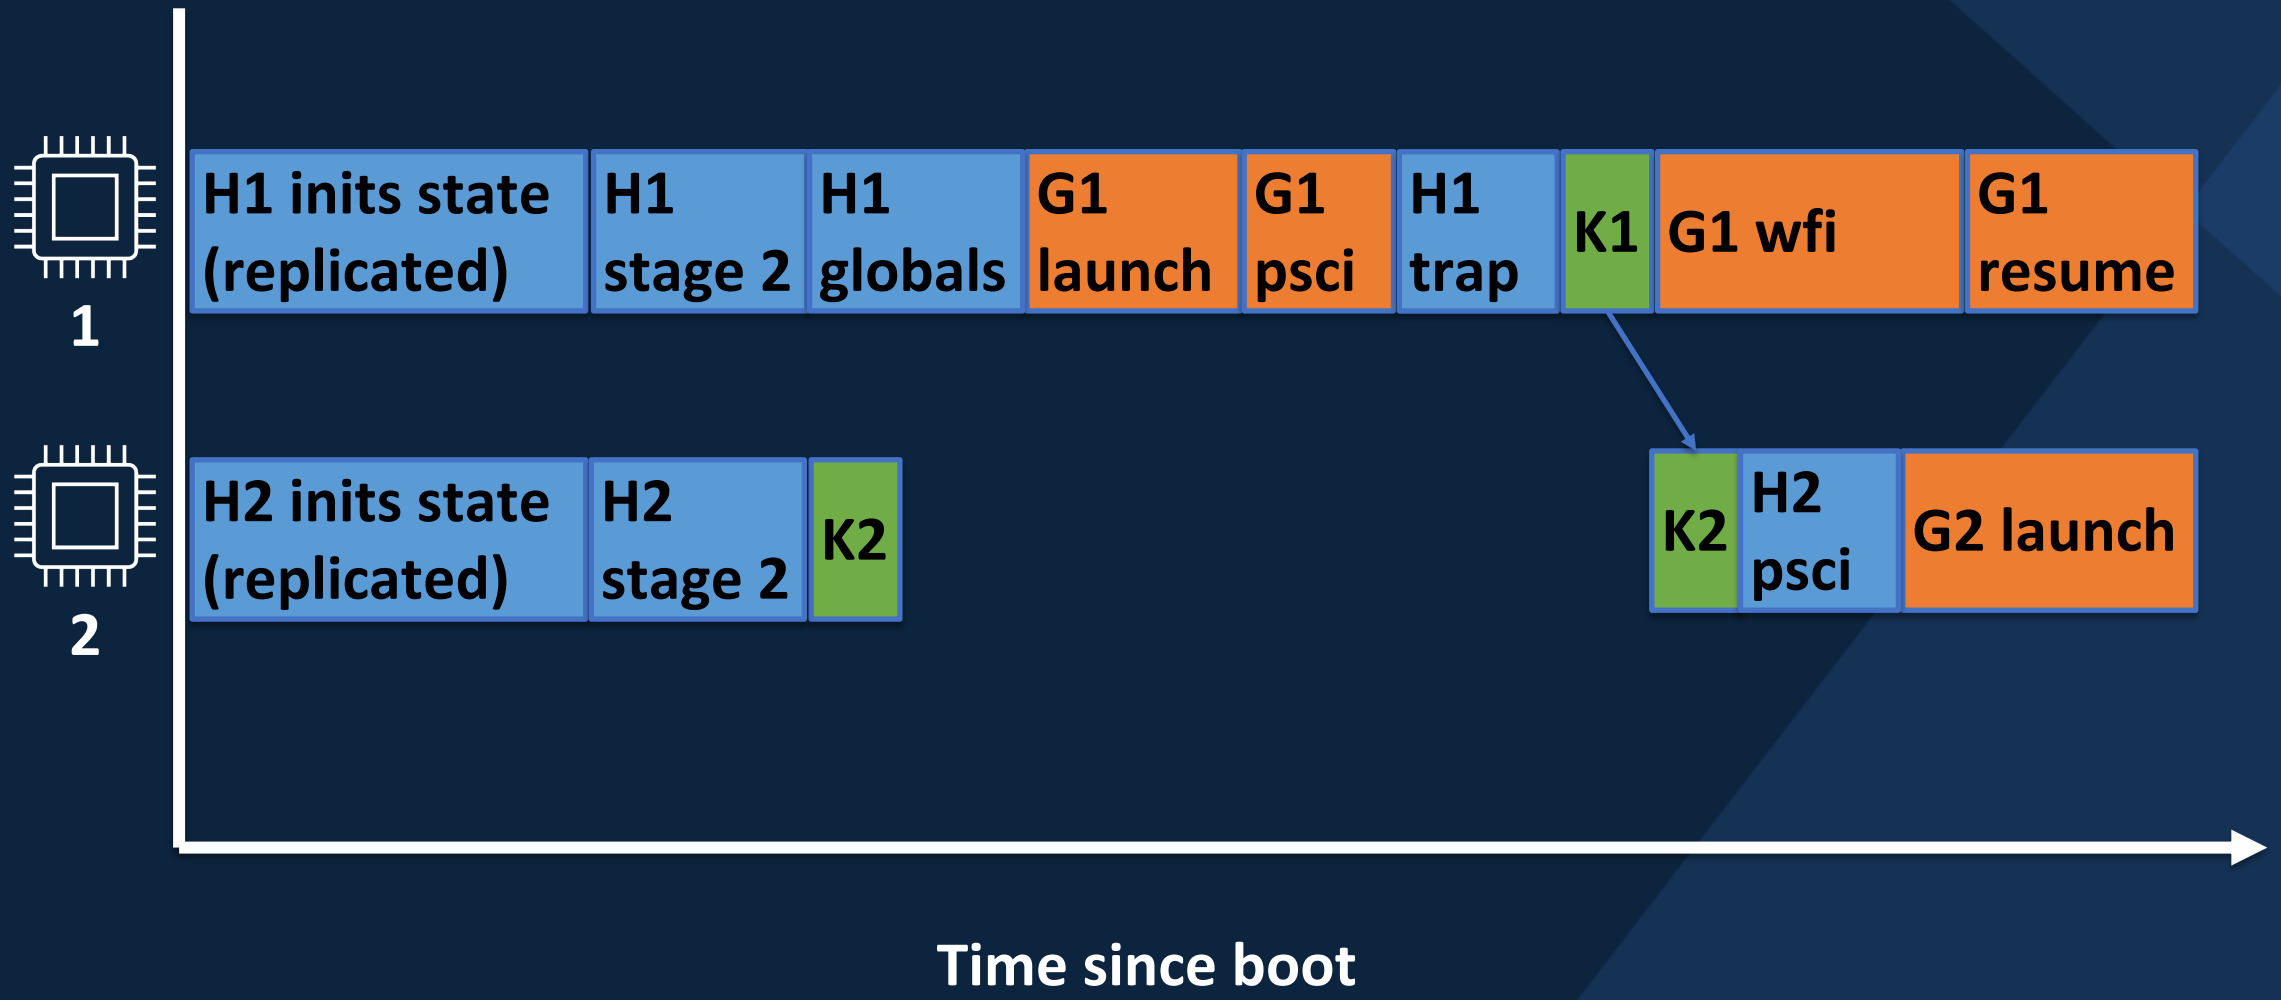

Virtual Interrupt controller (vGIC)

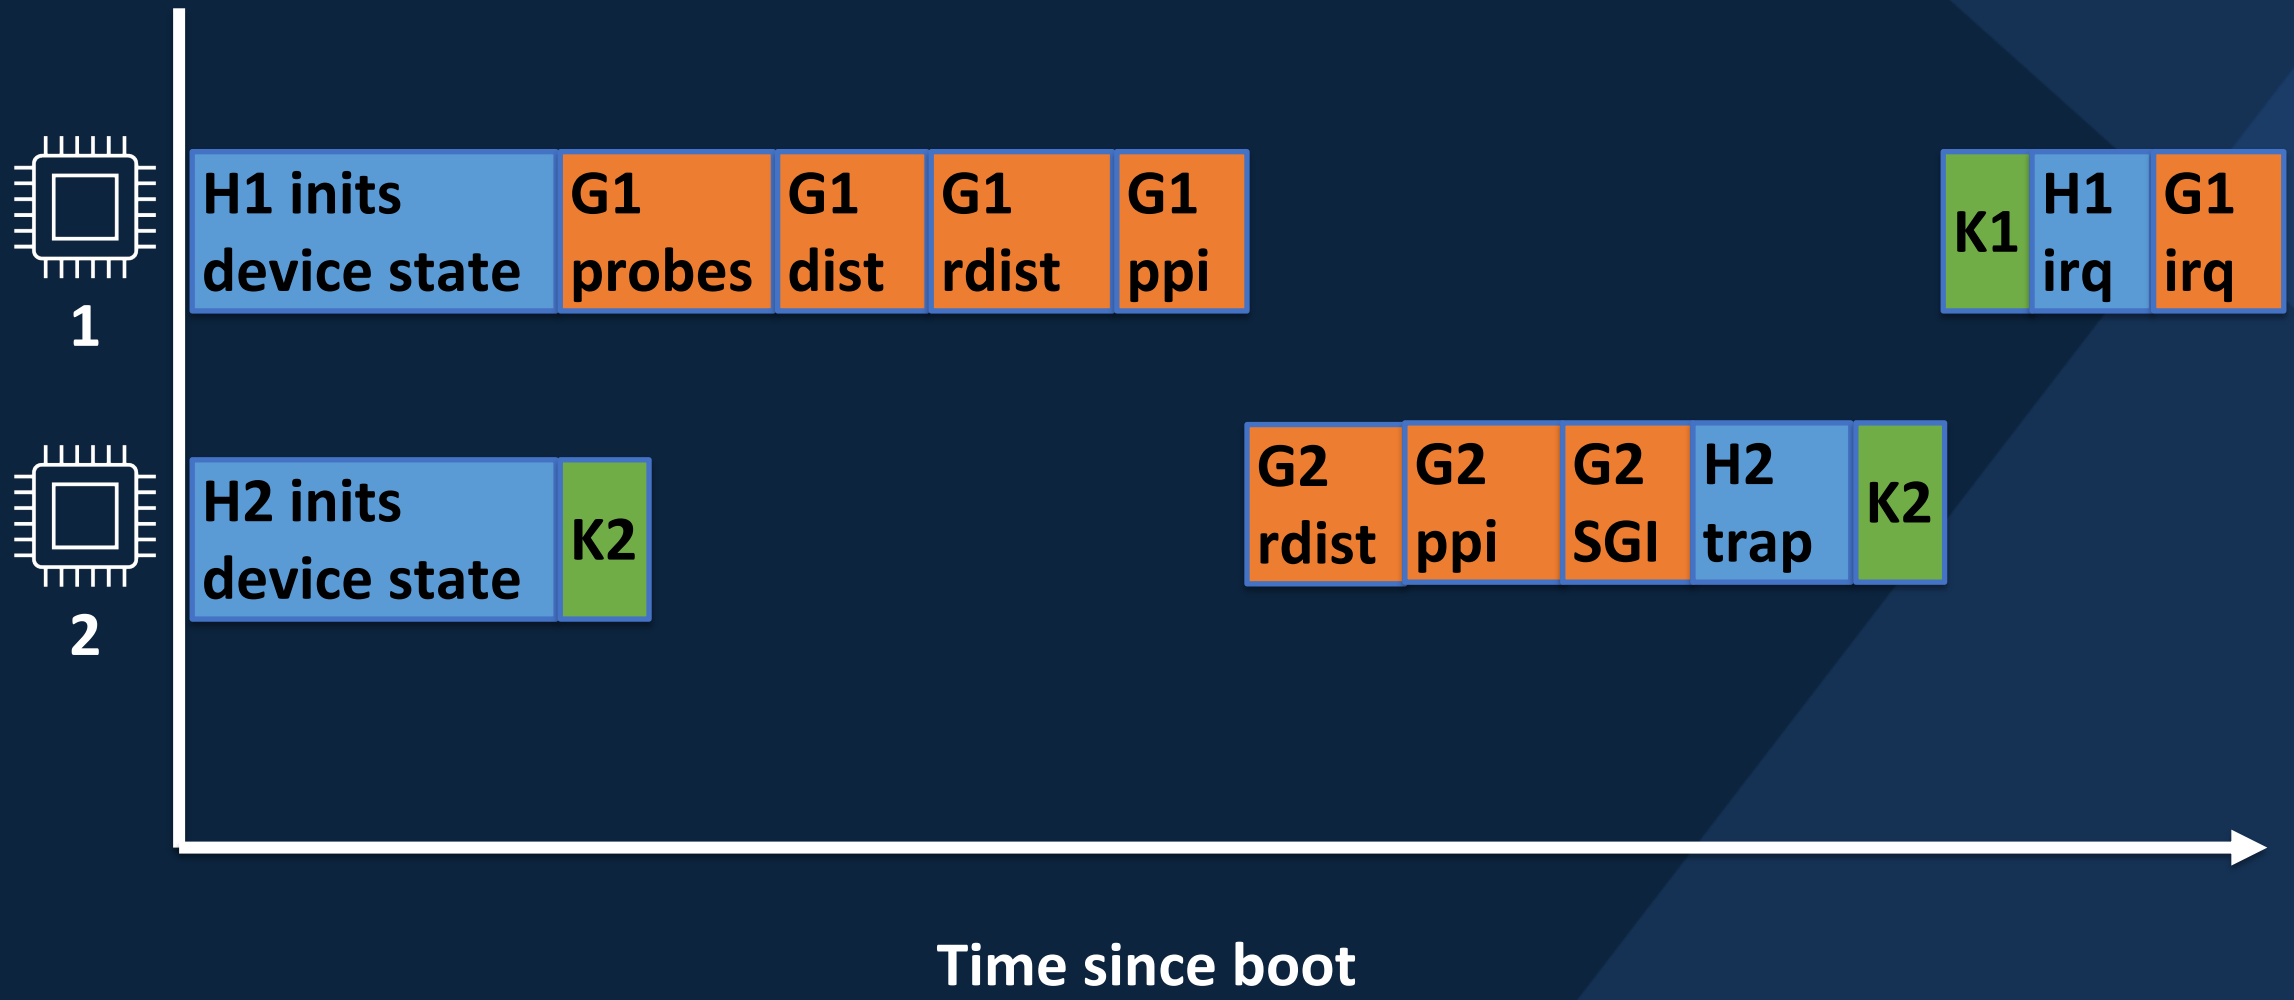

Other Examples

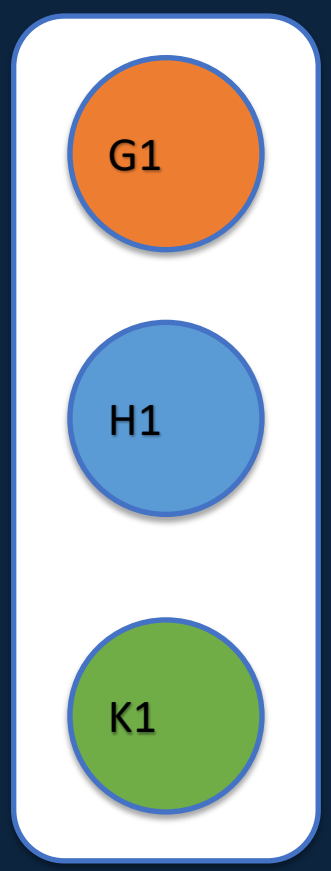

- Virtual devices eg virtio_net
- Cross platform communication
- Privileged operations (SMC calls or protected device access)

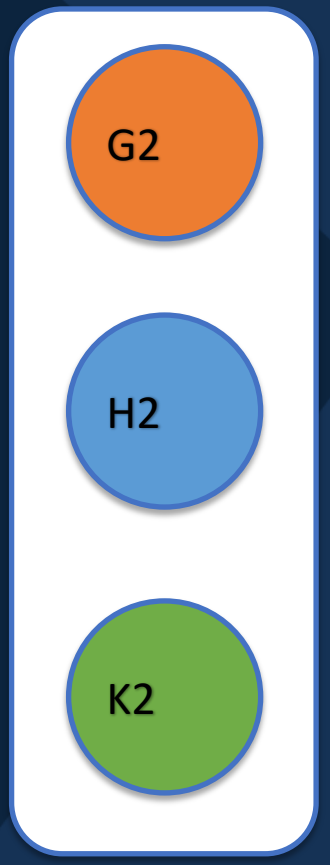

Does it work?

Yes

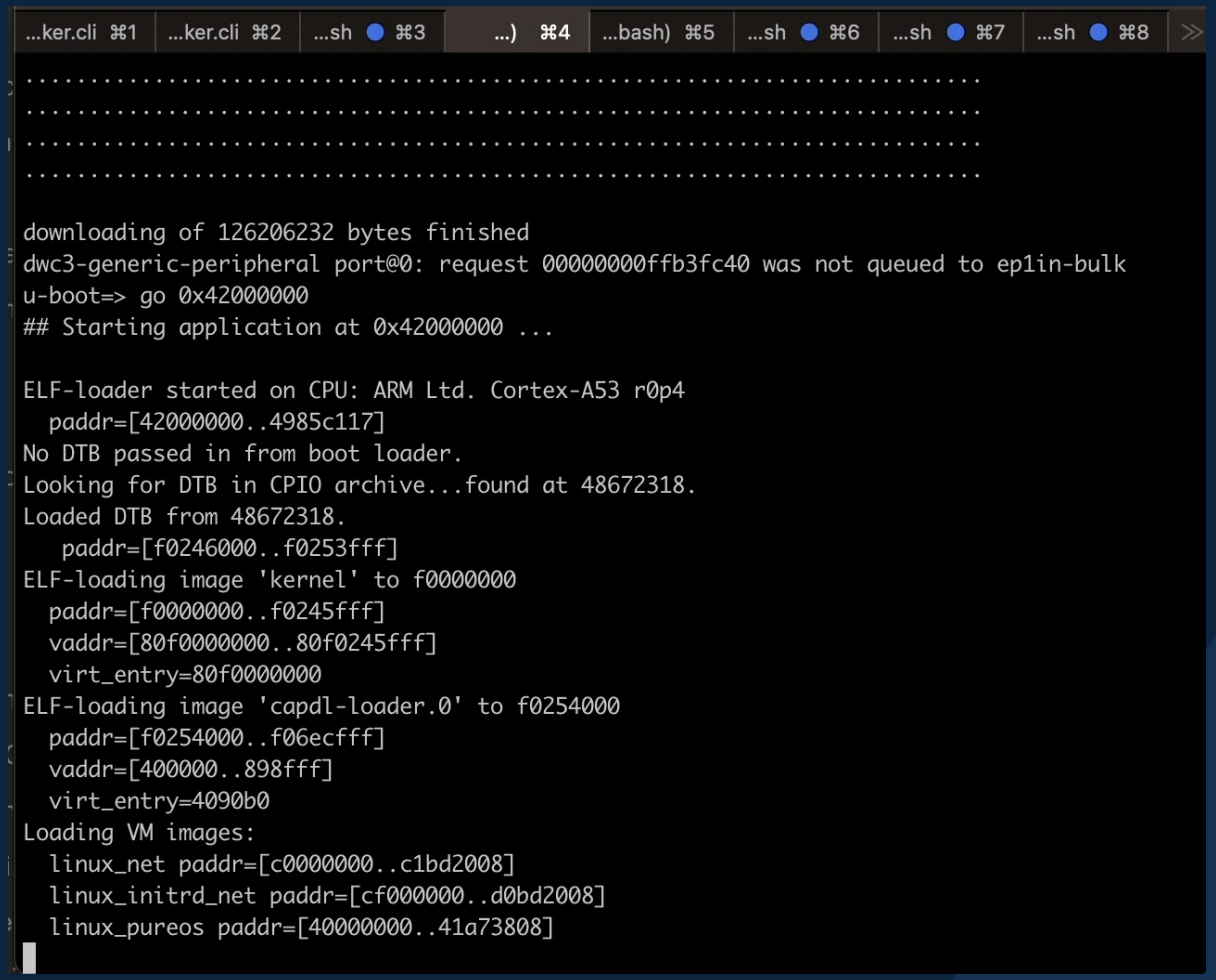

Discussion and extensions

- Dynamic systems
- multiple vms.
- power on and off cores
- Where are scalability issues?
- How trustworthy can VMM layer be?