



Linaro
connect
San Francisco 2017

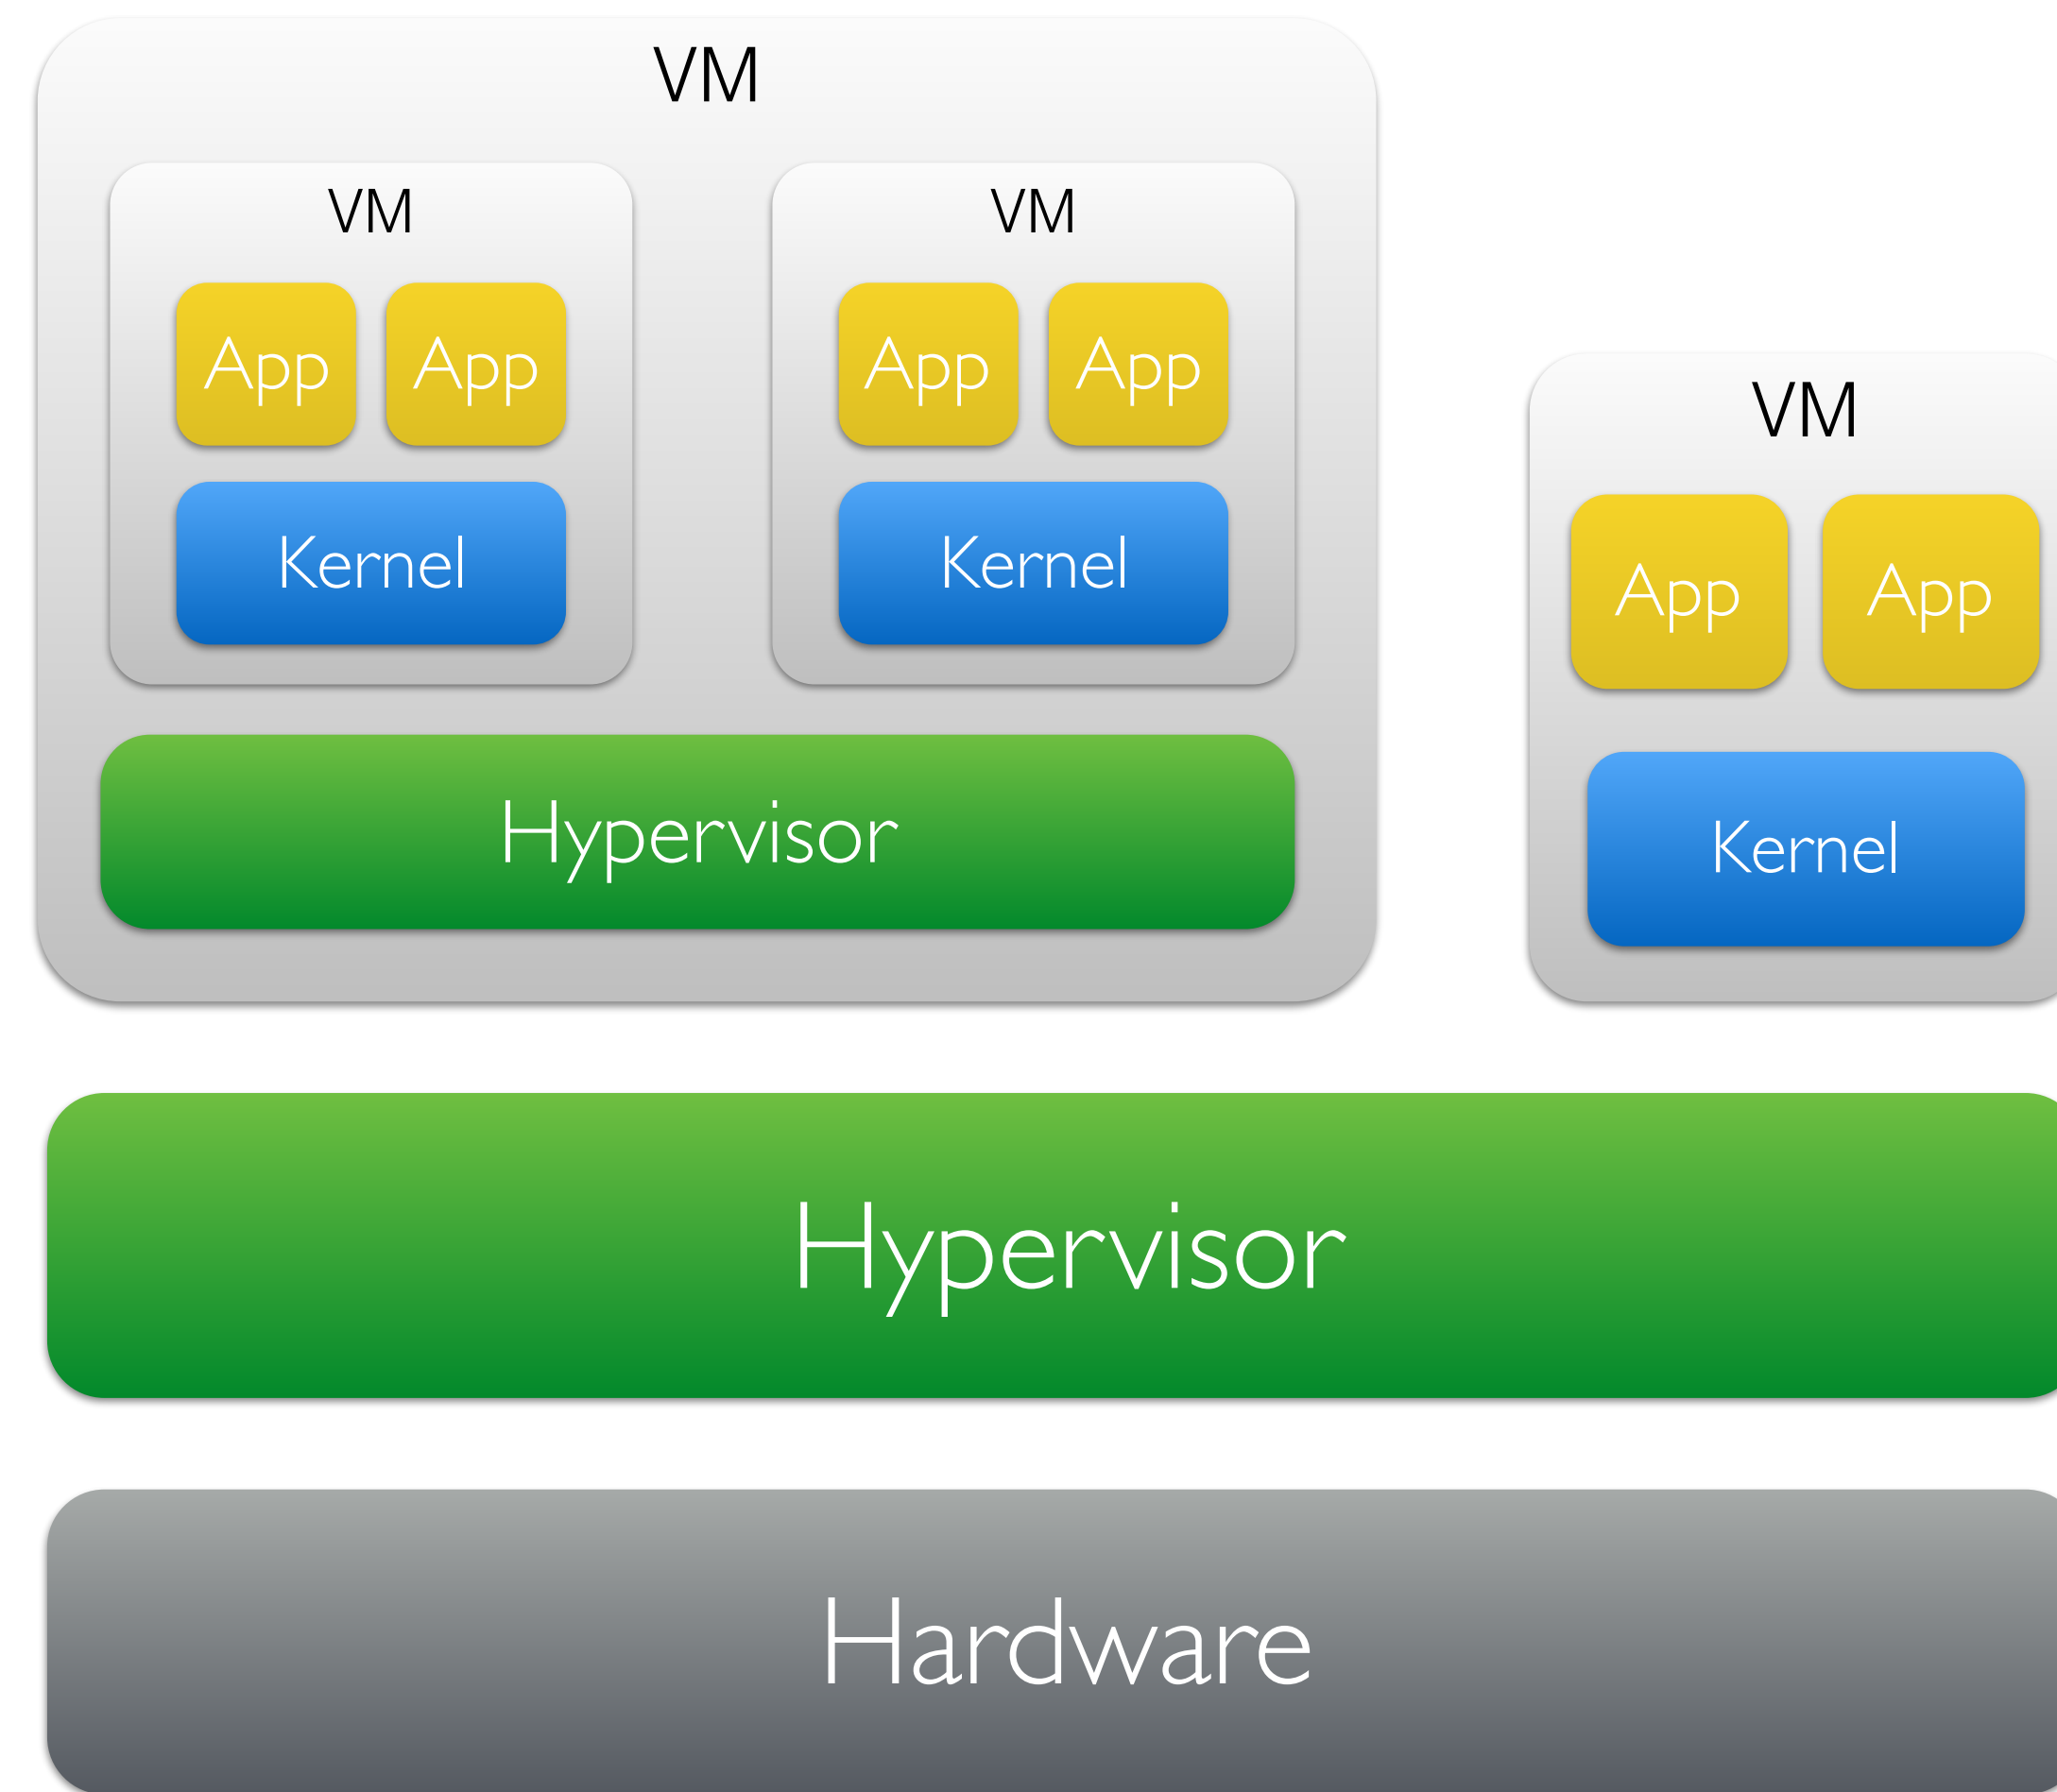
SFO17-410: NEVE: Nested Virtualization Extensions for ARM

Jintack Lim, Christoffer Dall, Shih-Wei Li, Jason Nieh, and Marc Zyngier



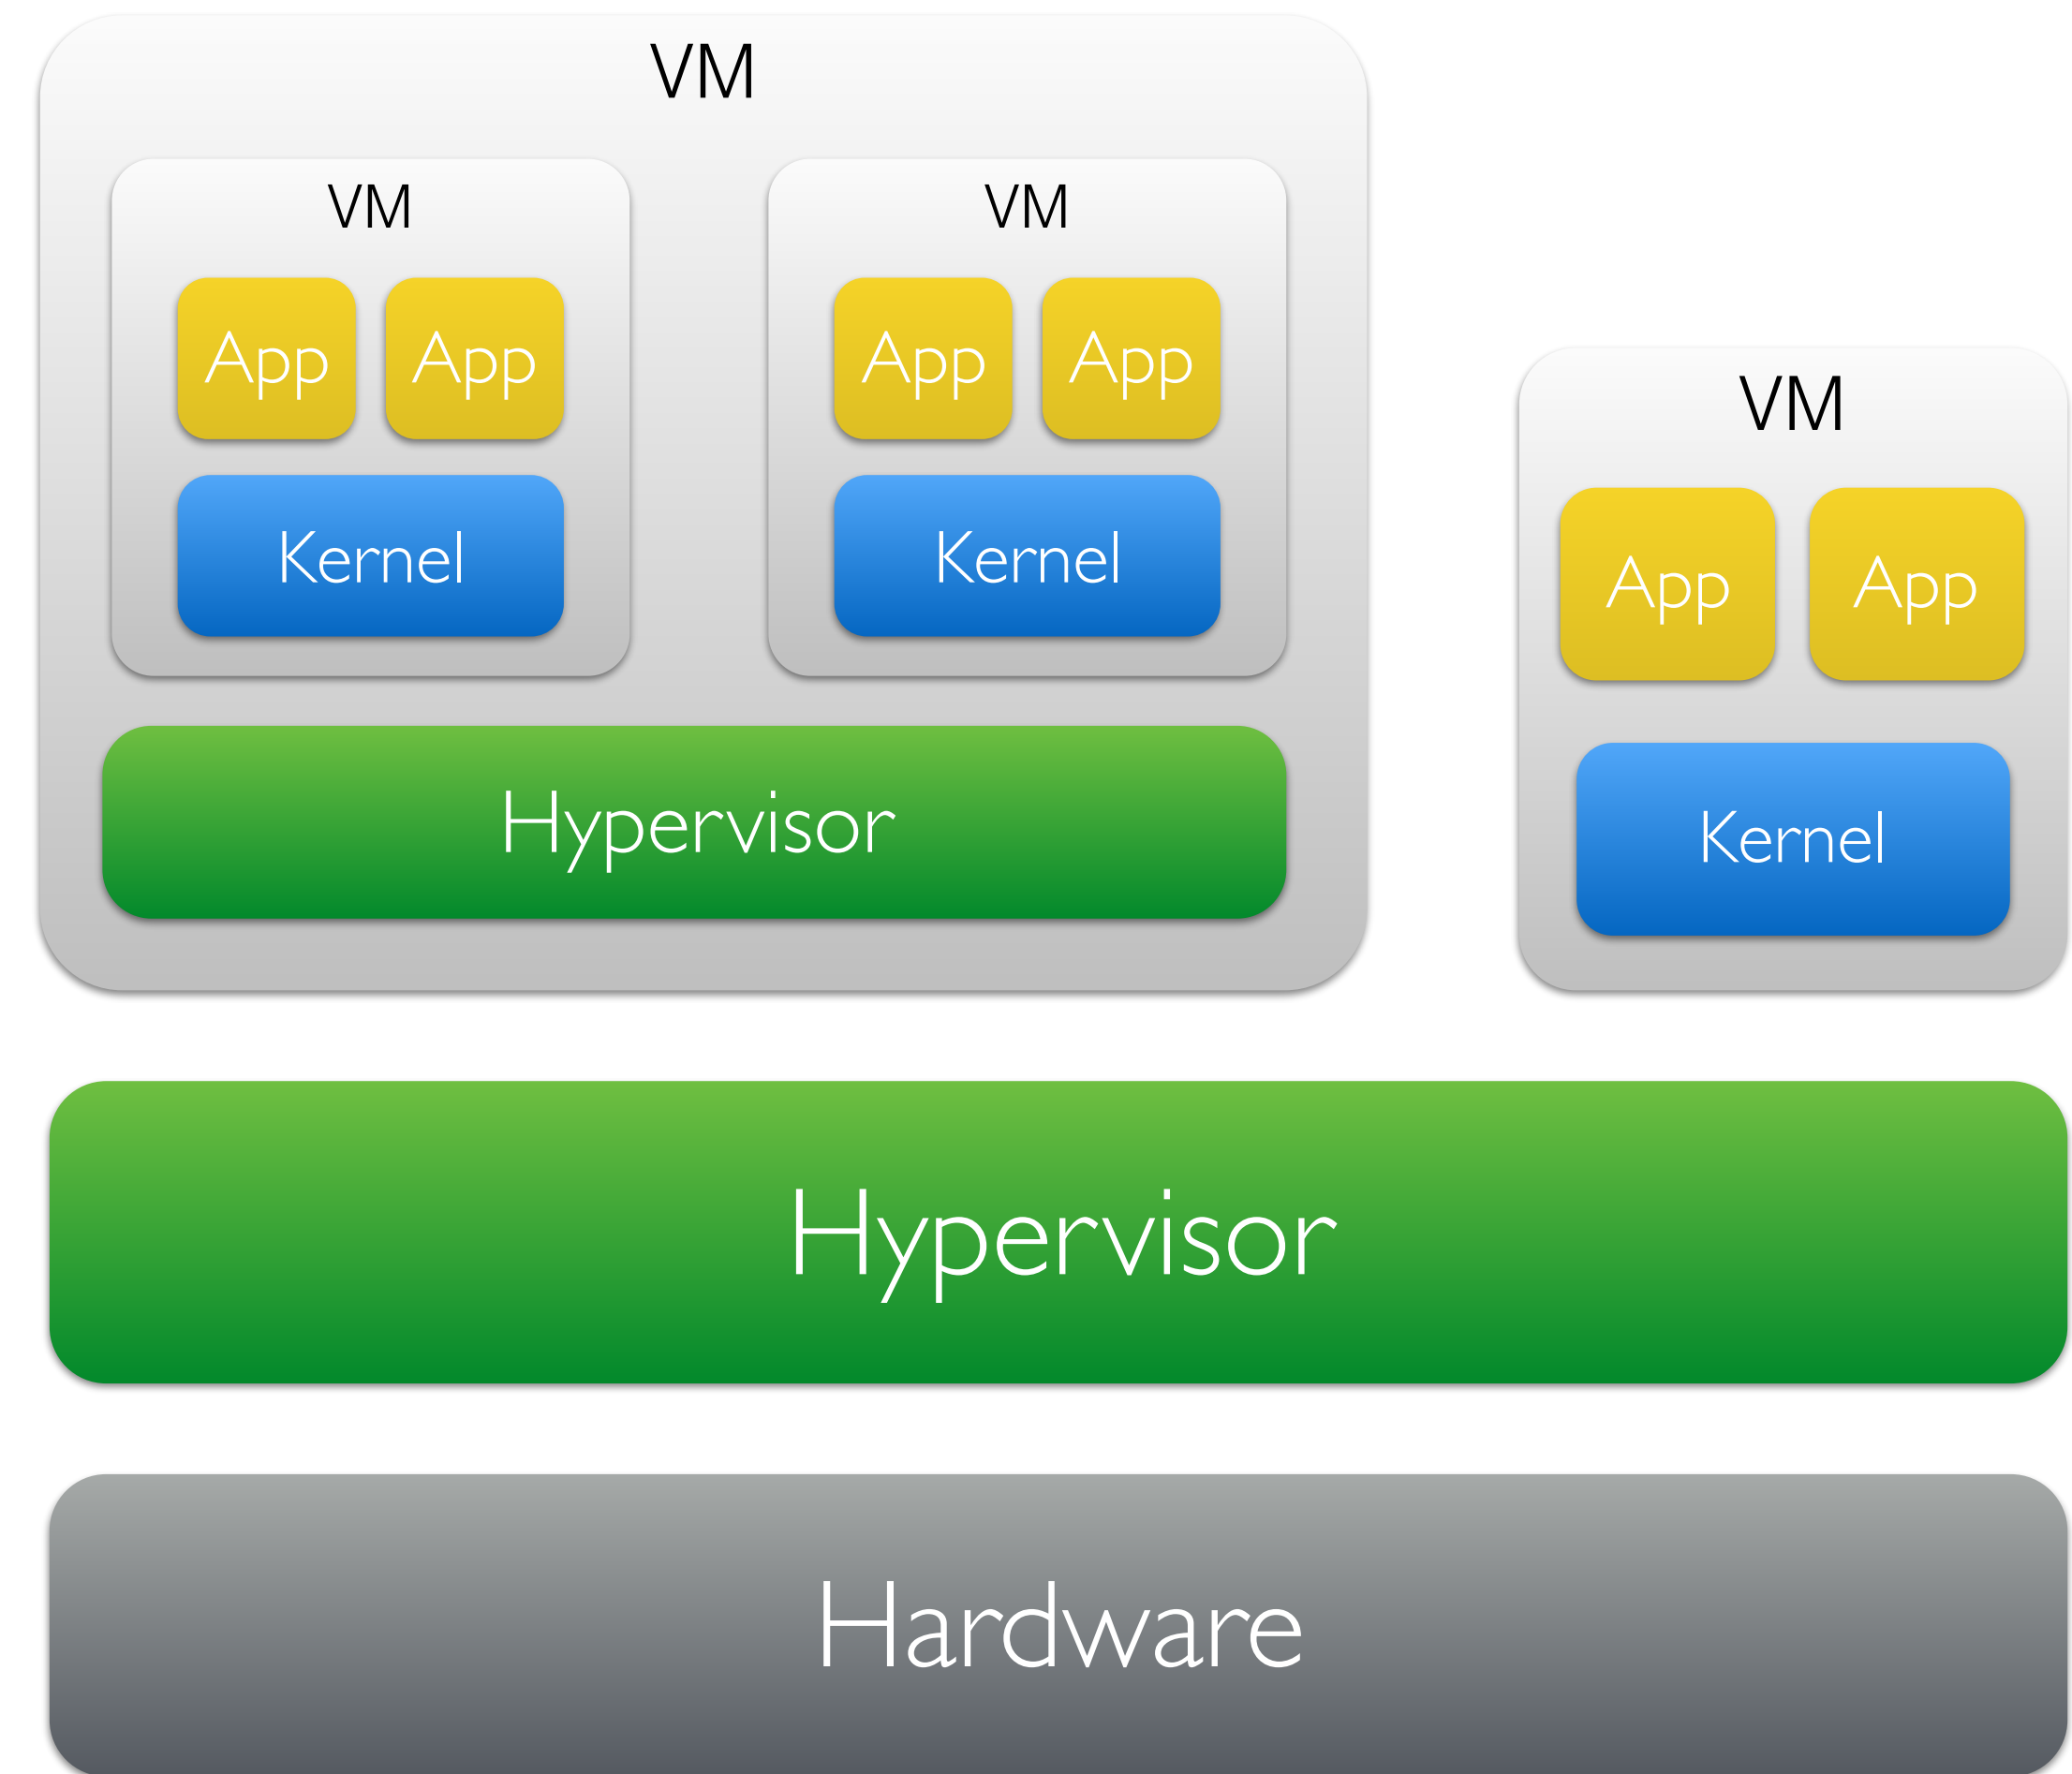
connect.linaro.org

Nested Virtualization

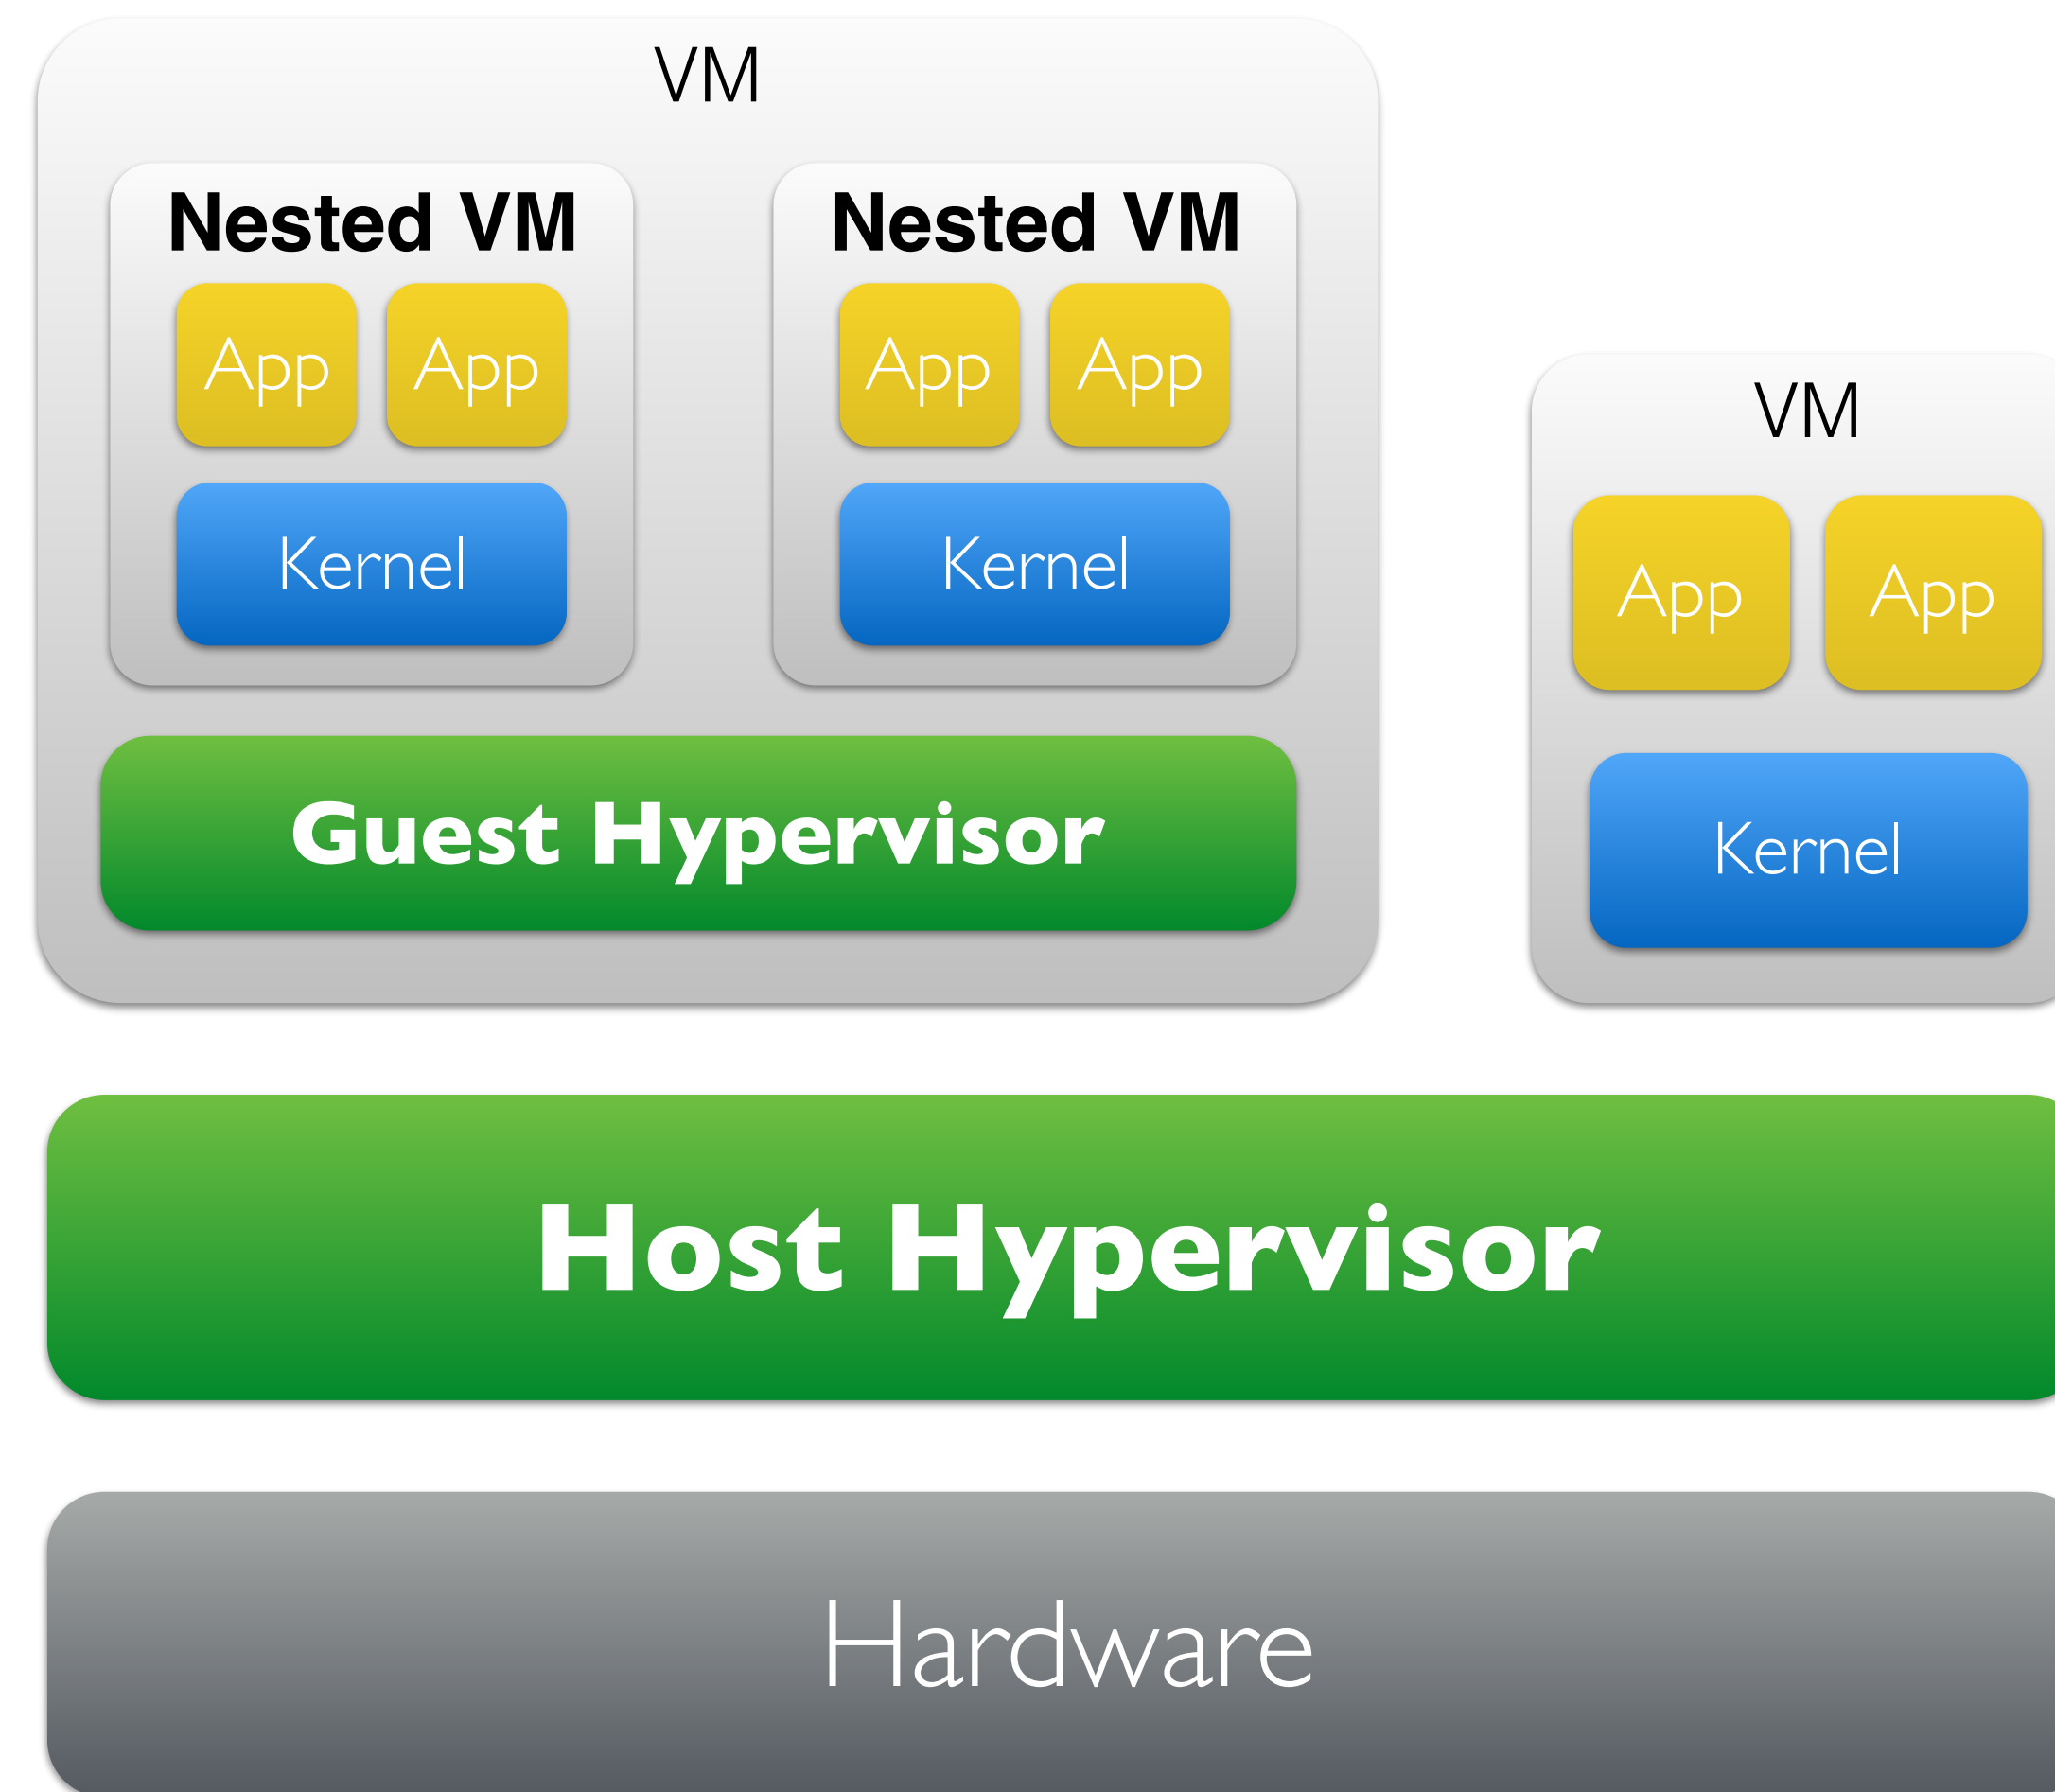


Nested Virtualization

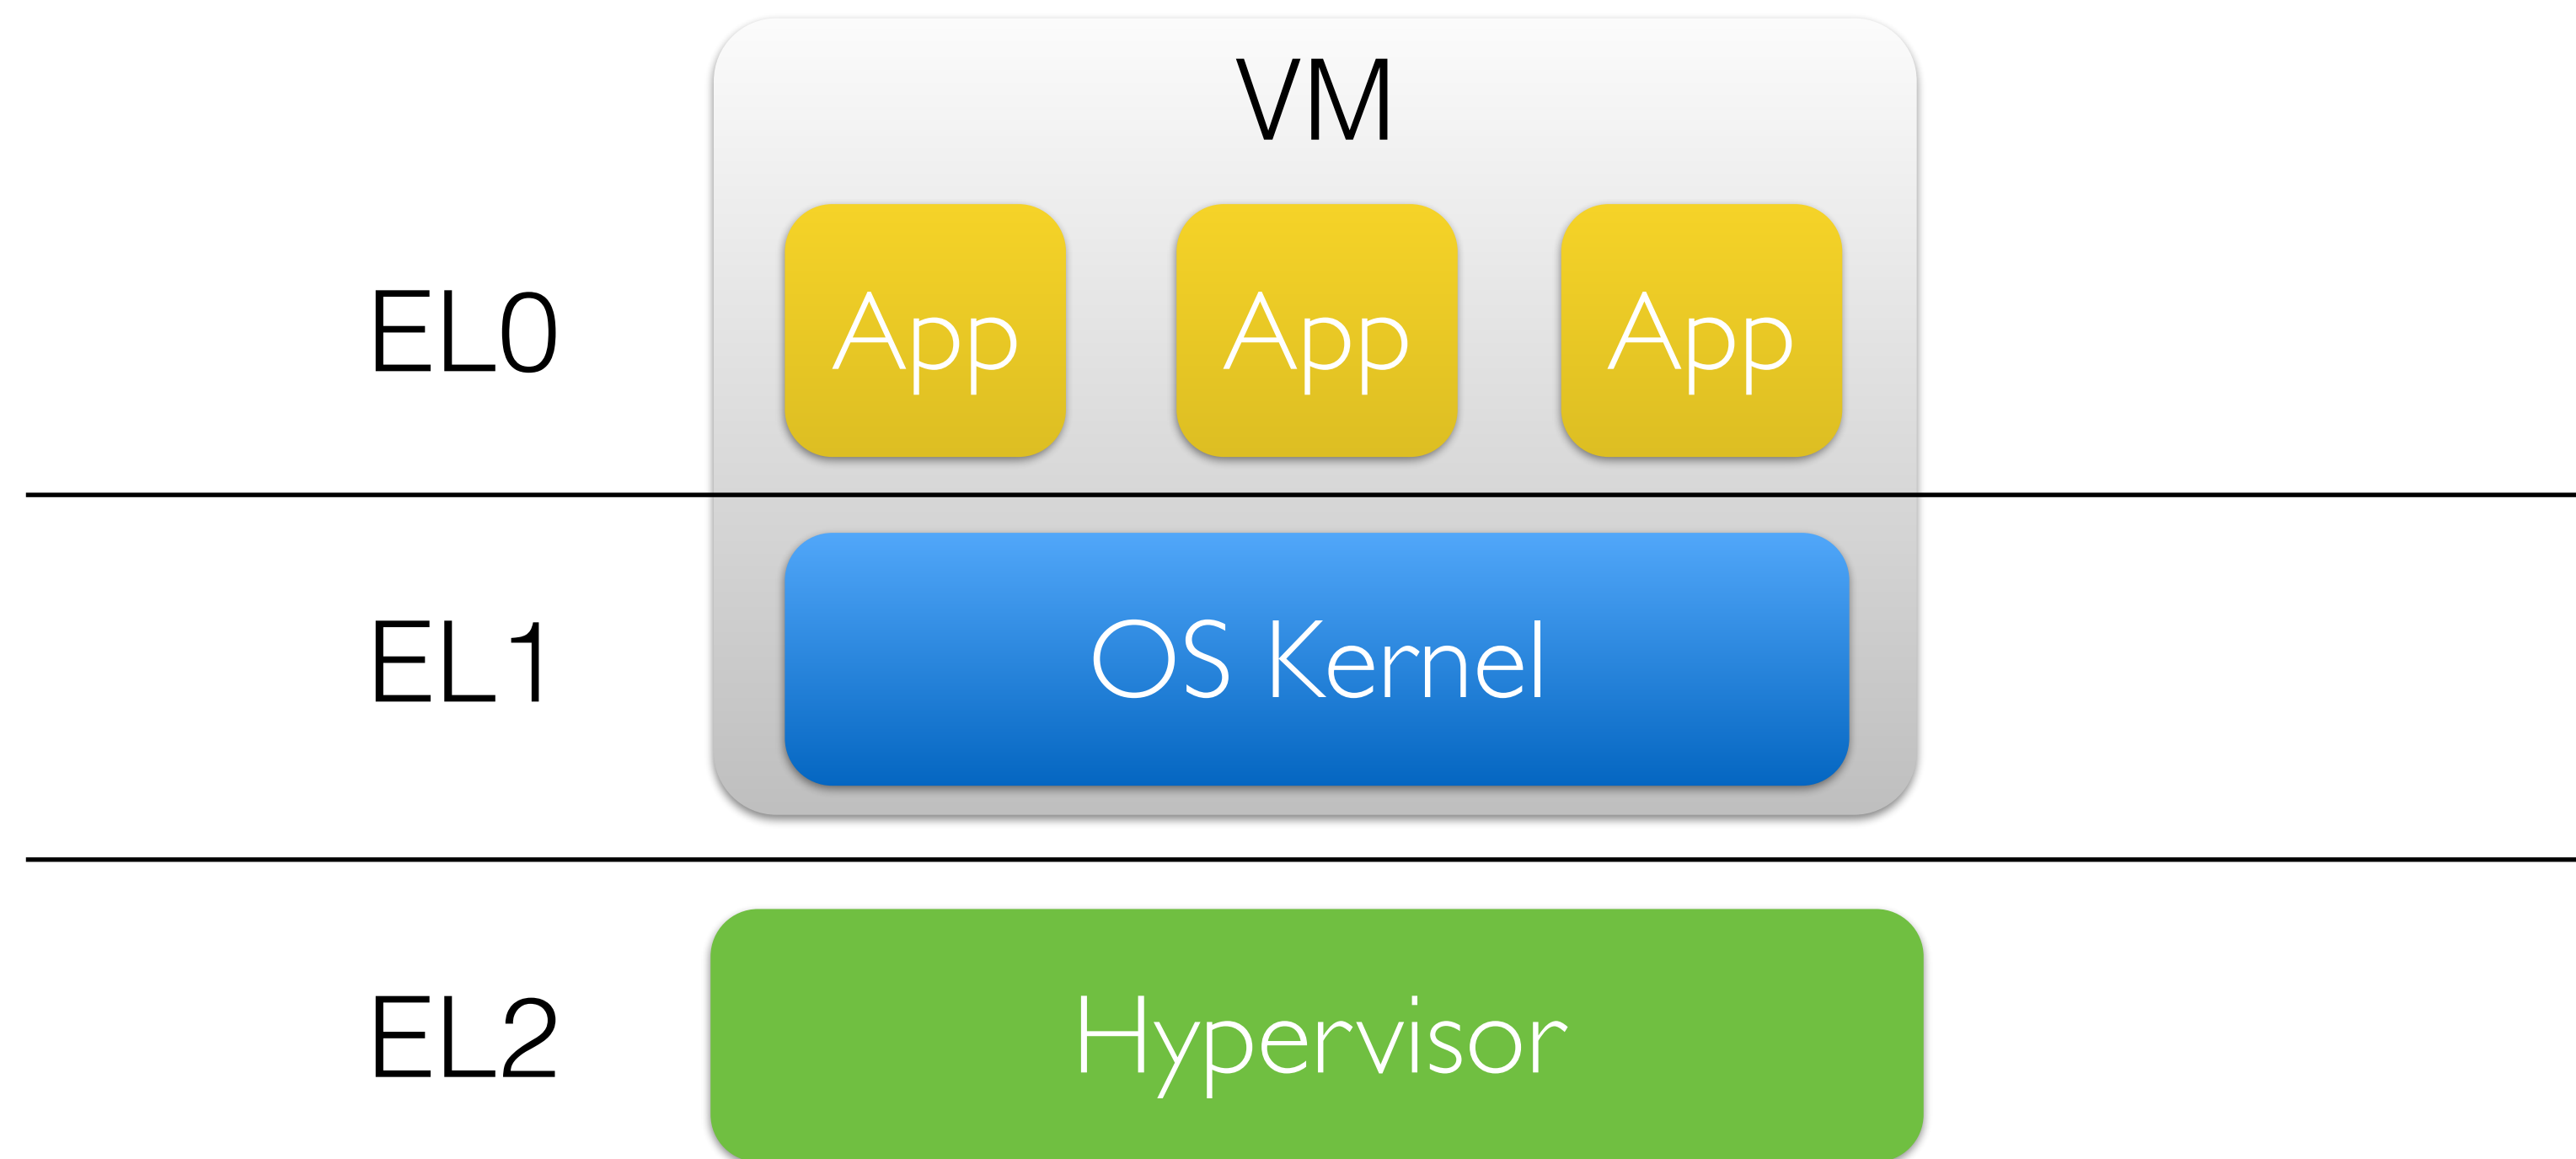
- IaaS hosting private clouds
- Develop and test your hypervisor



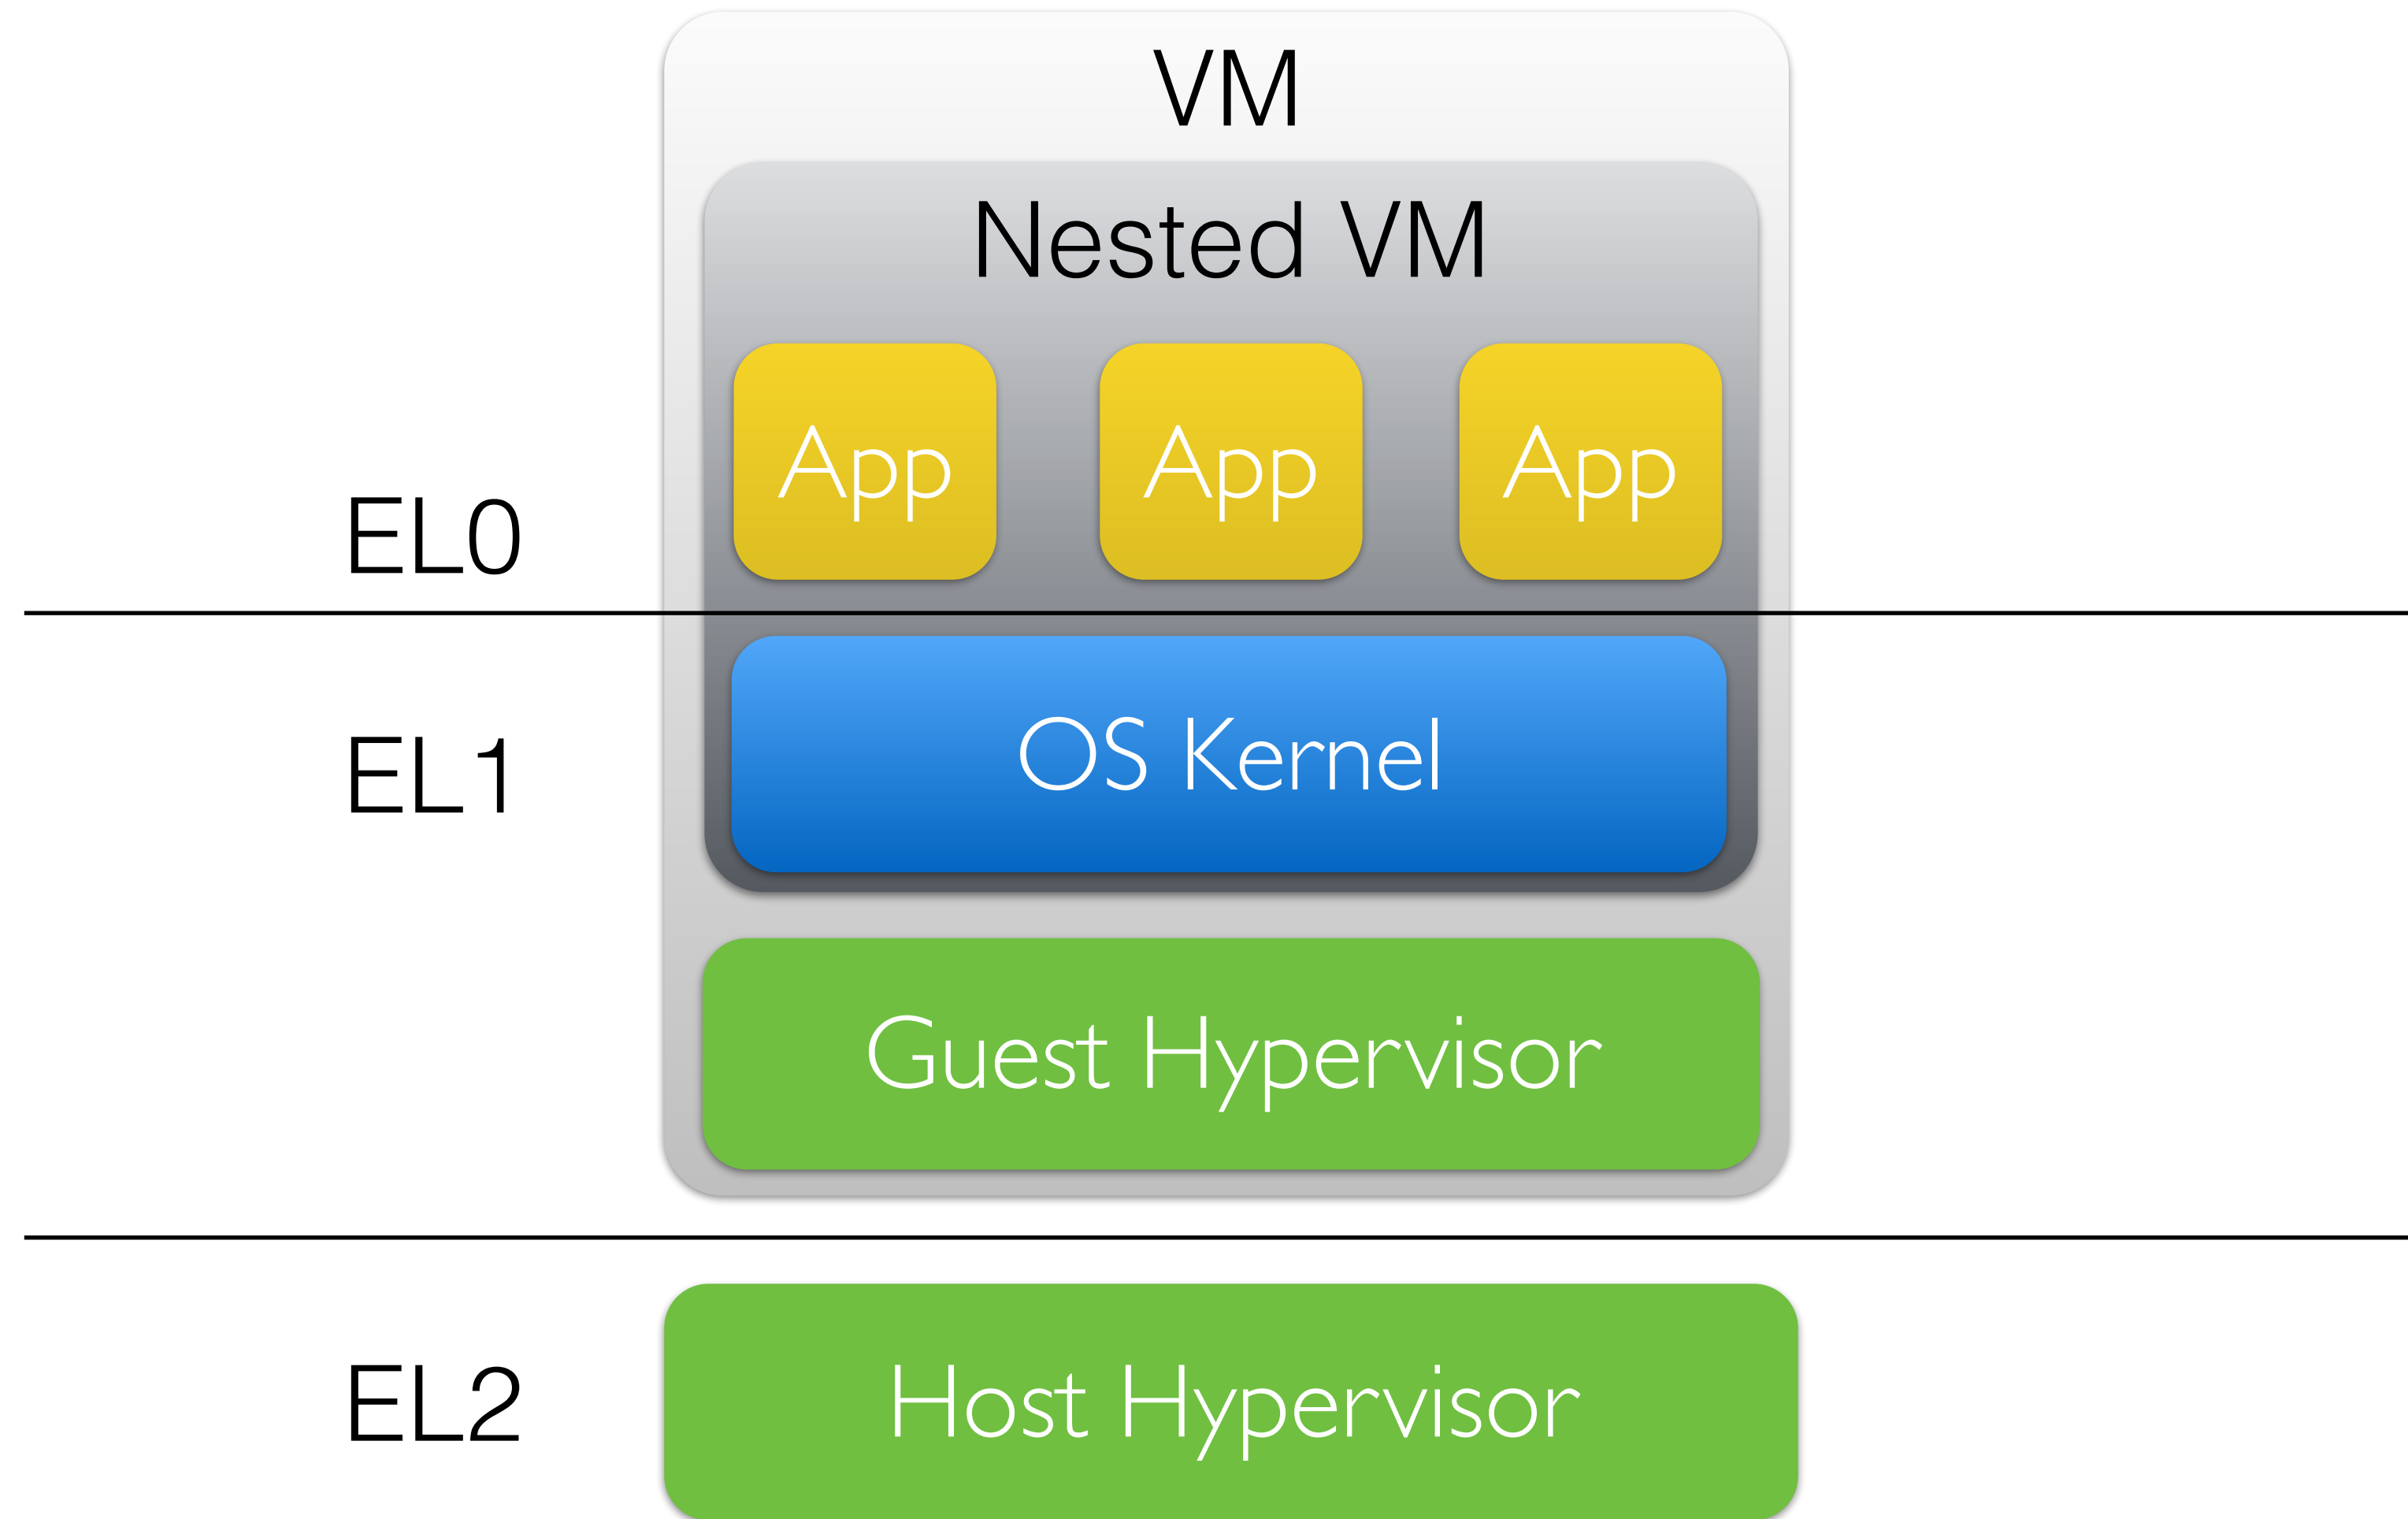
Terminology



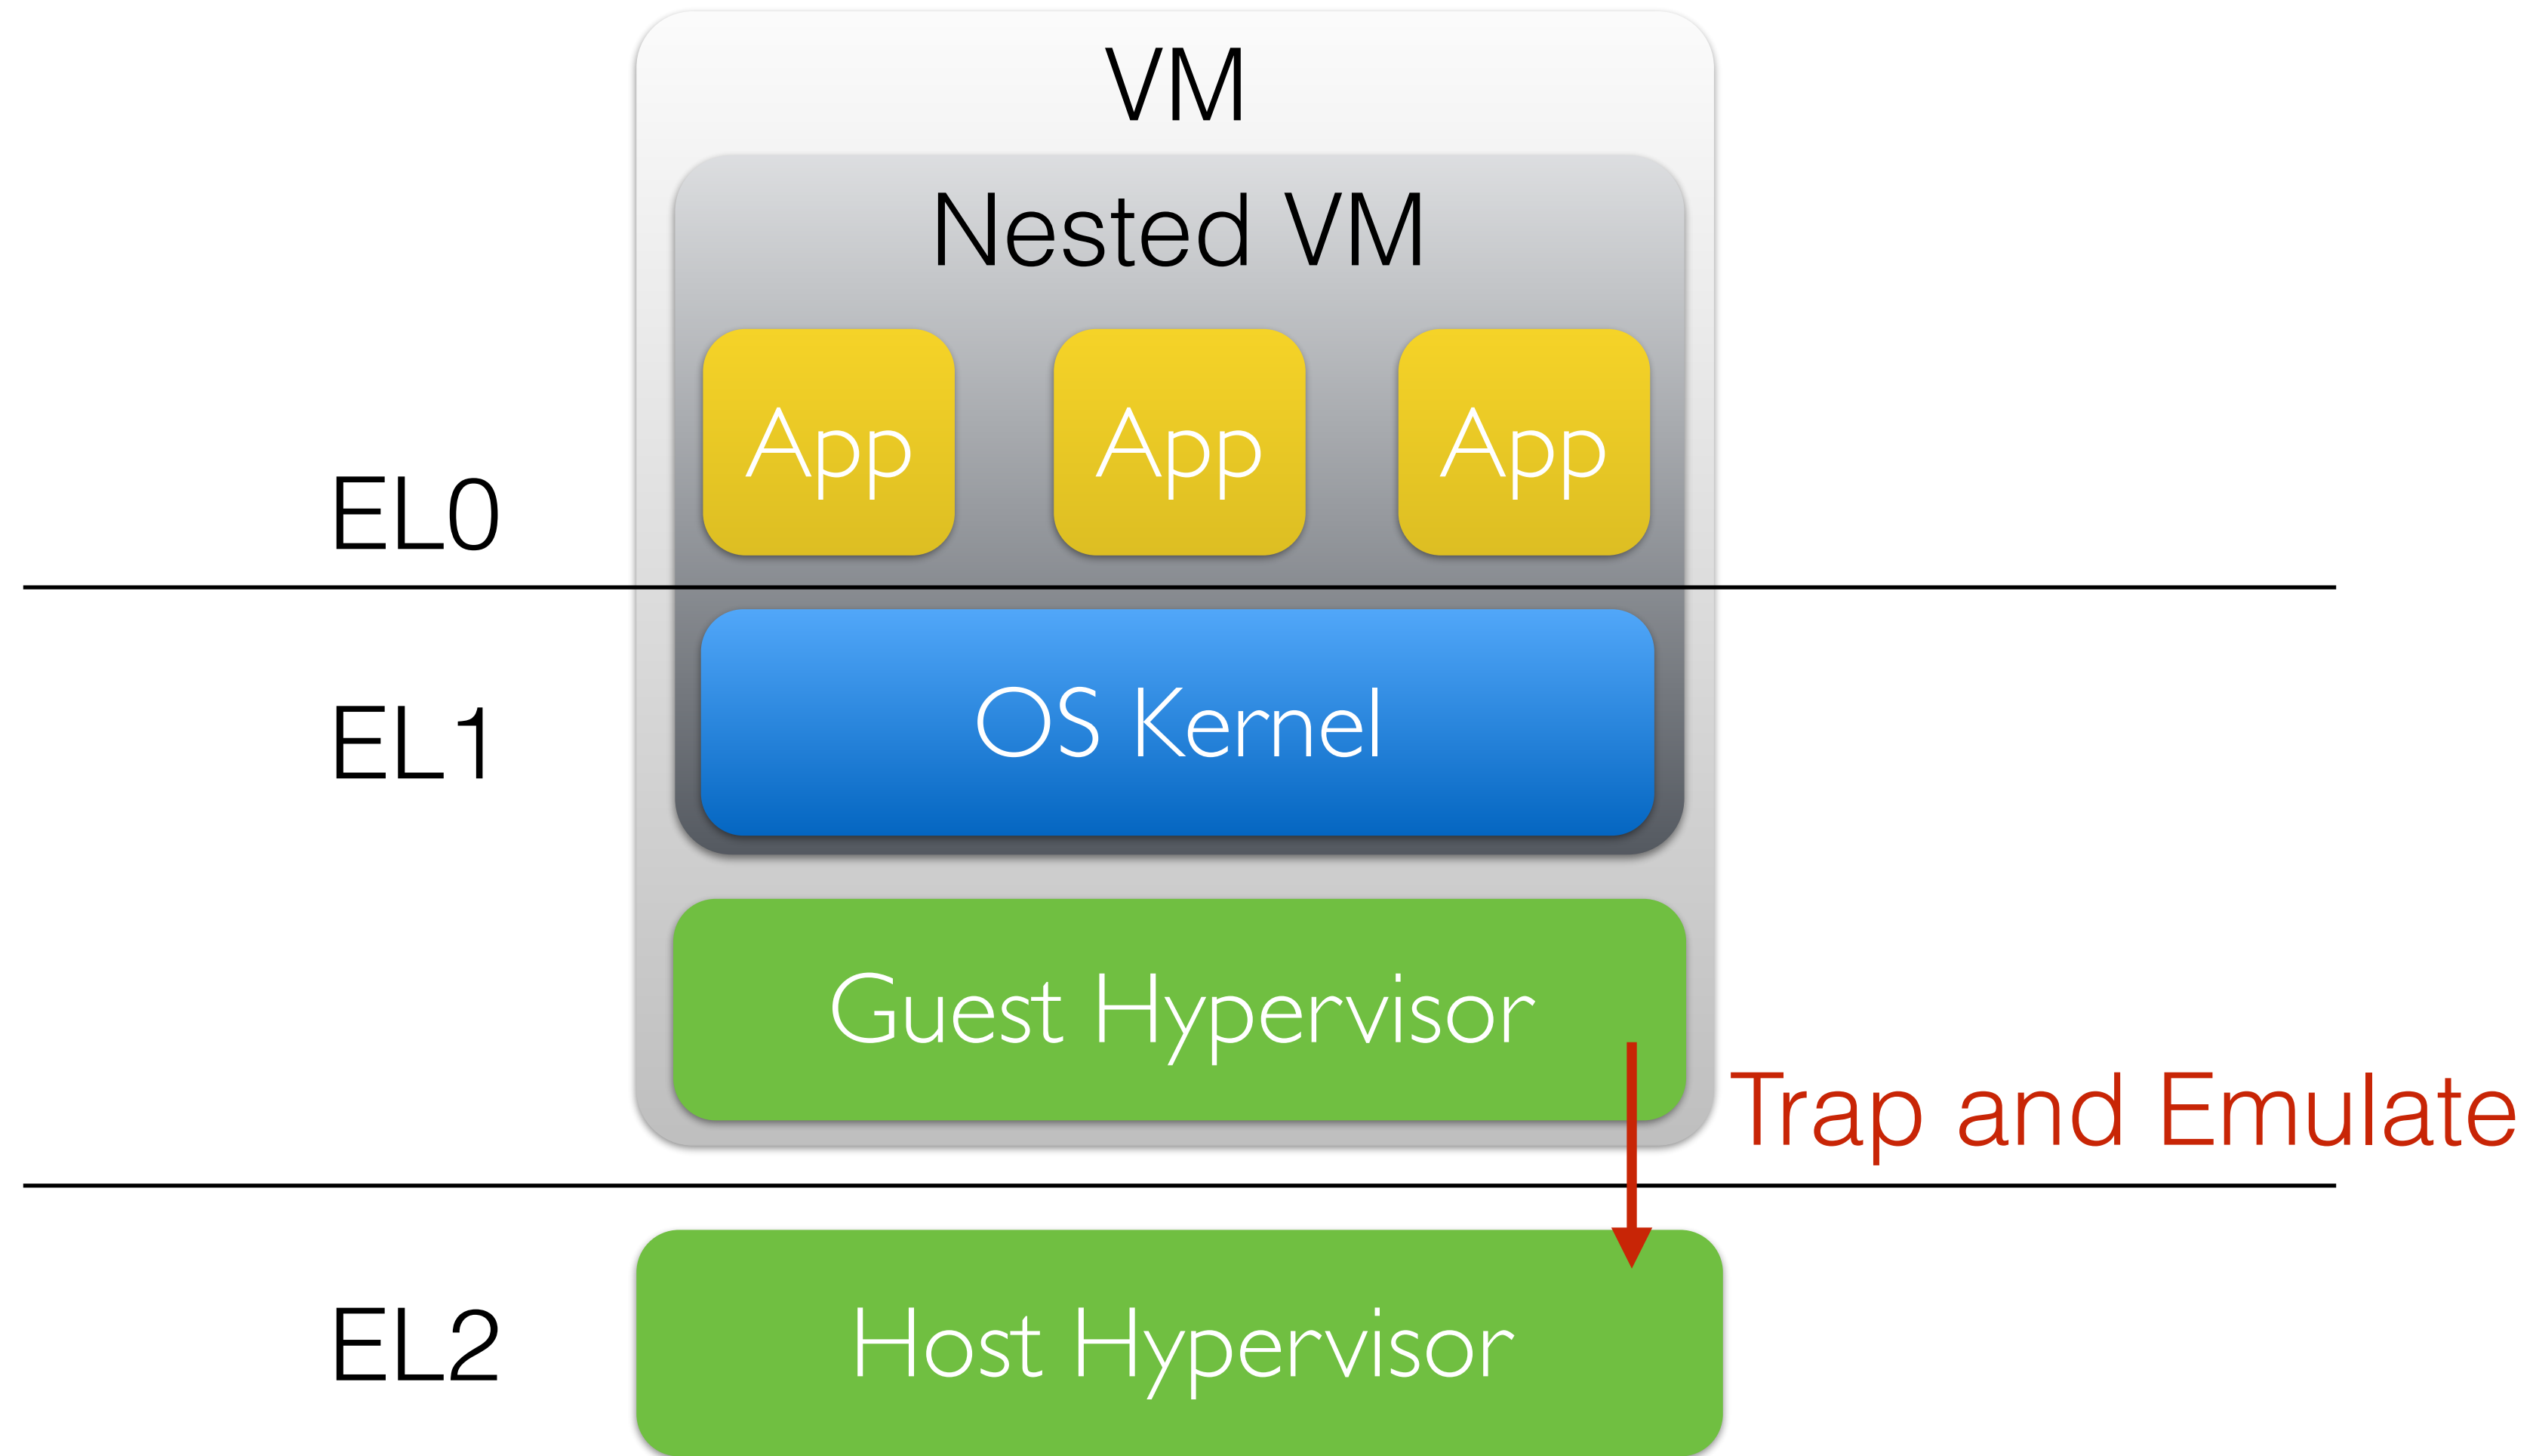
ARM Virtualization Extensions



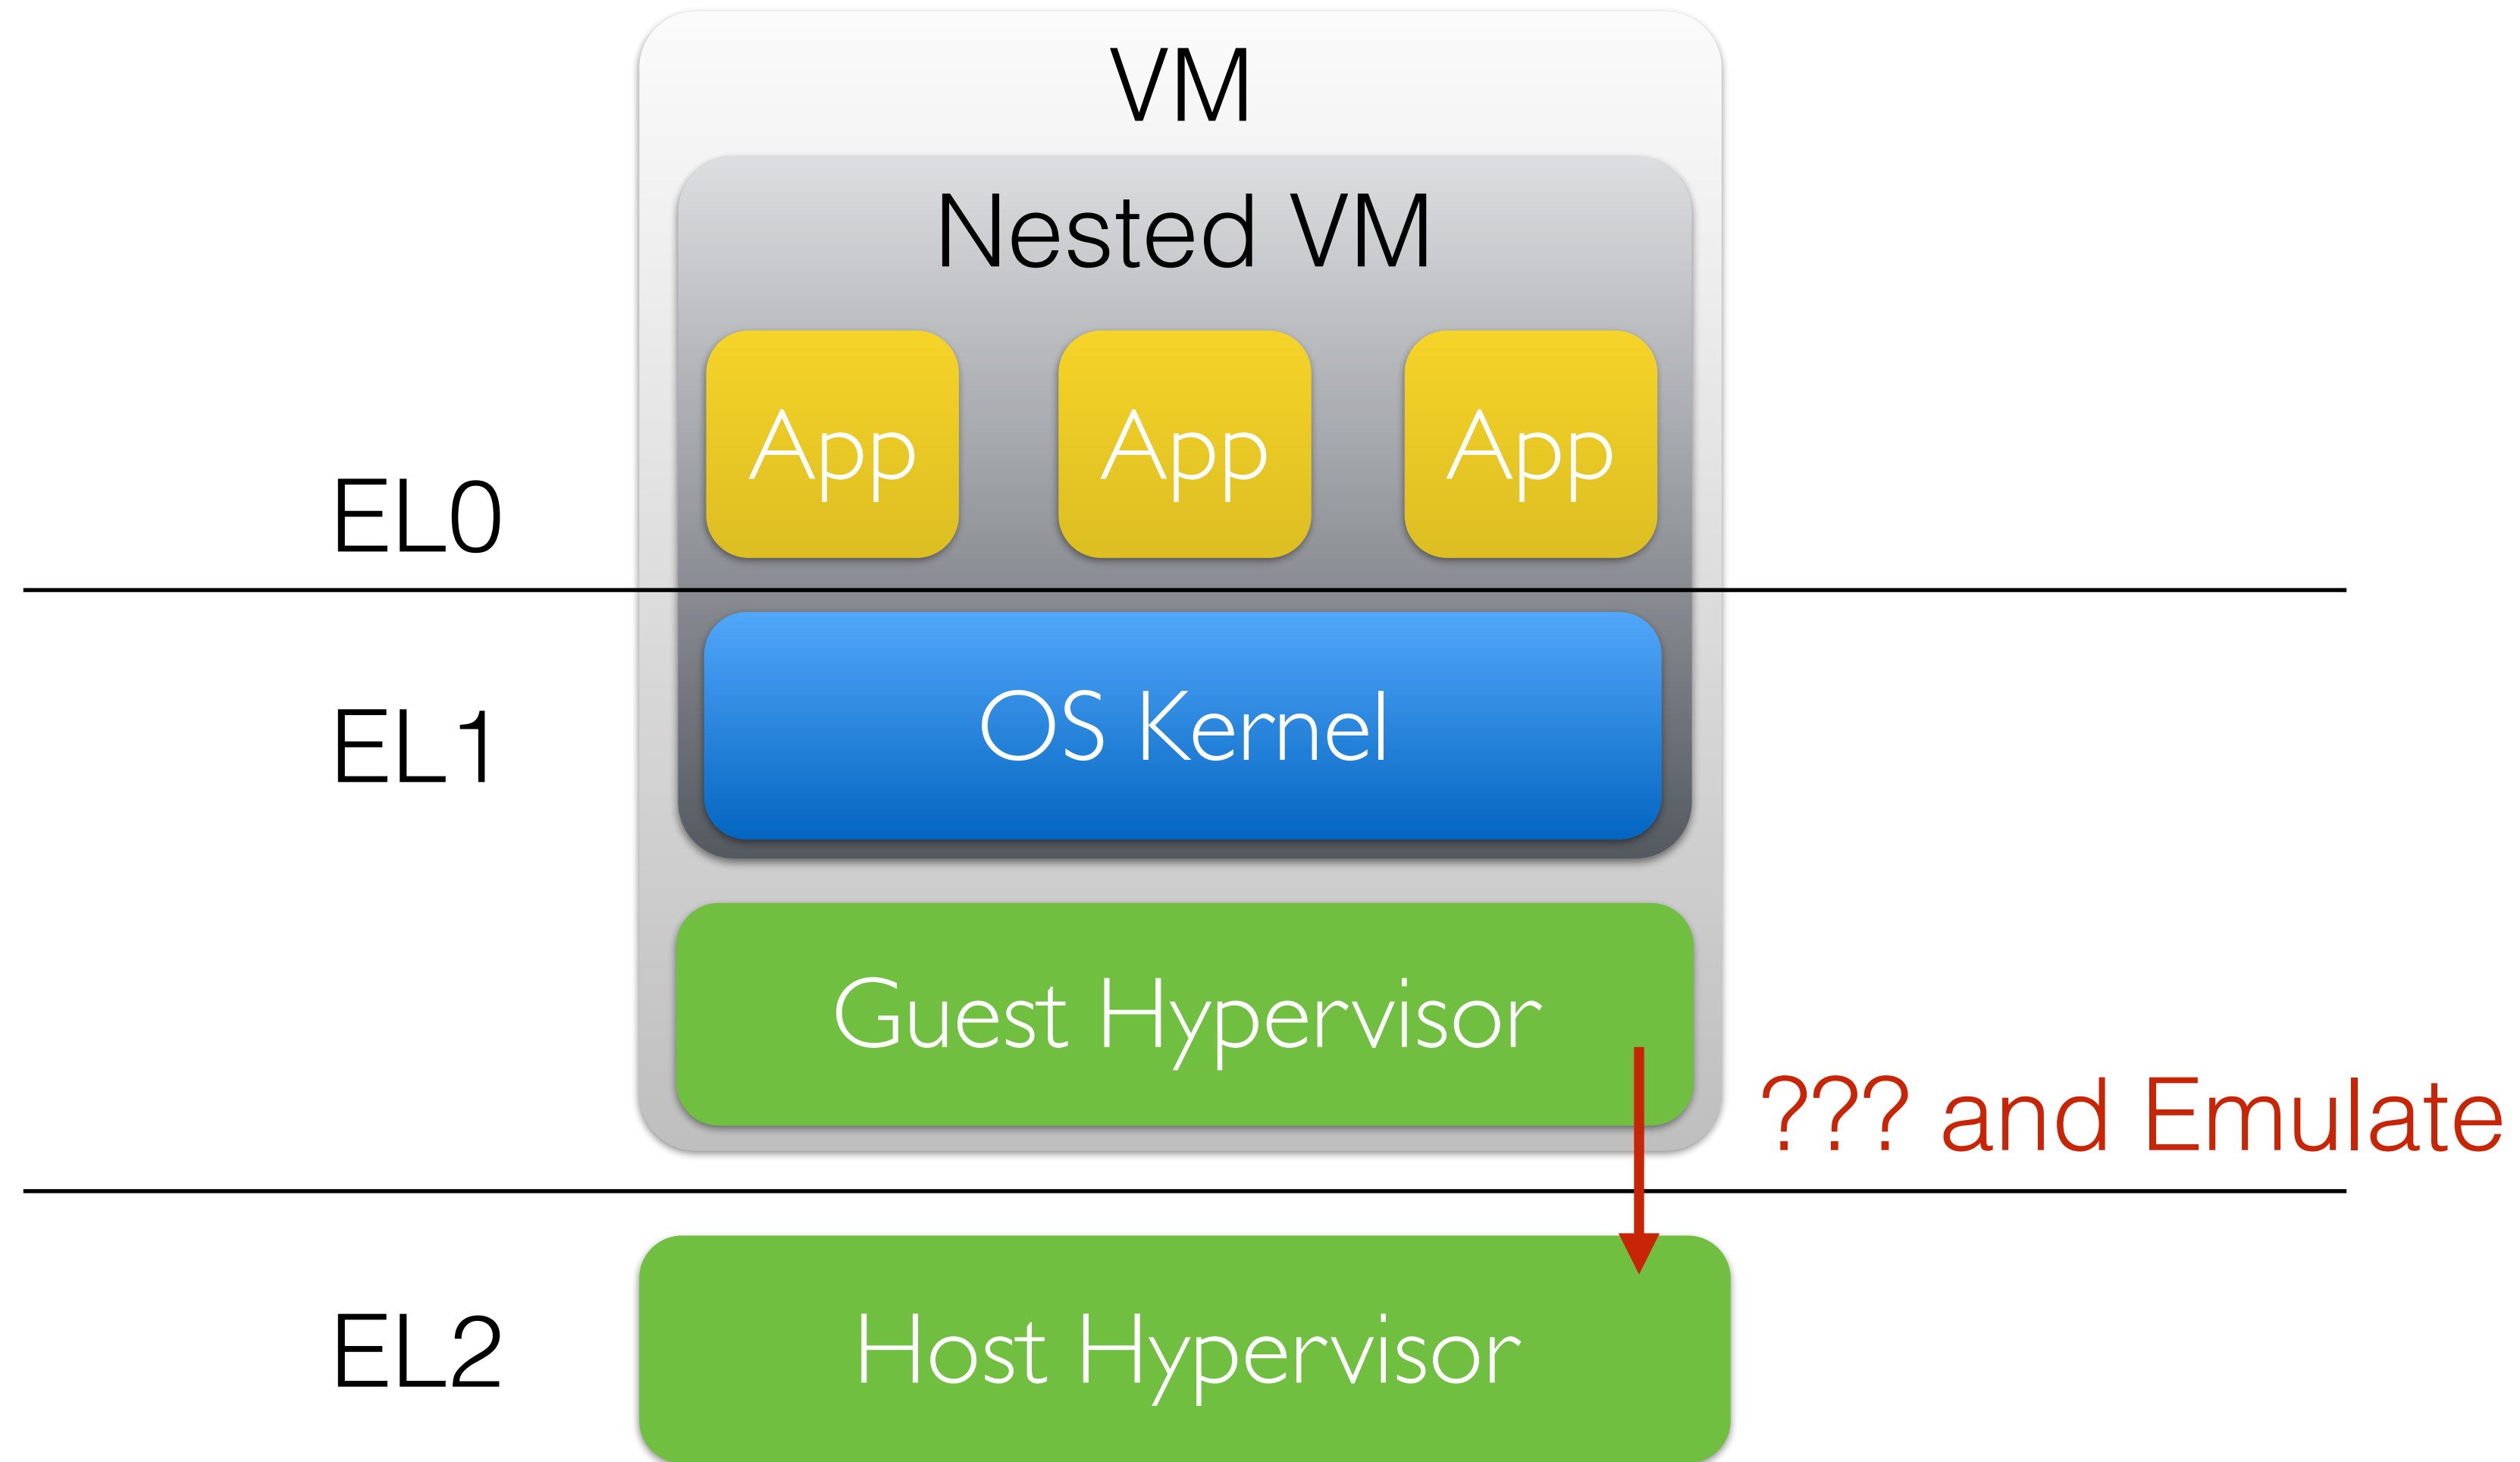
ARM Nested Virtualization



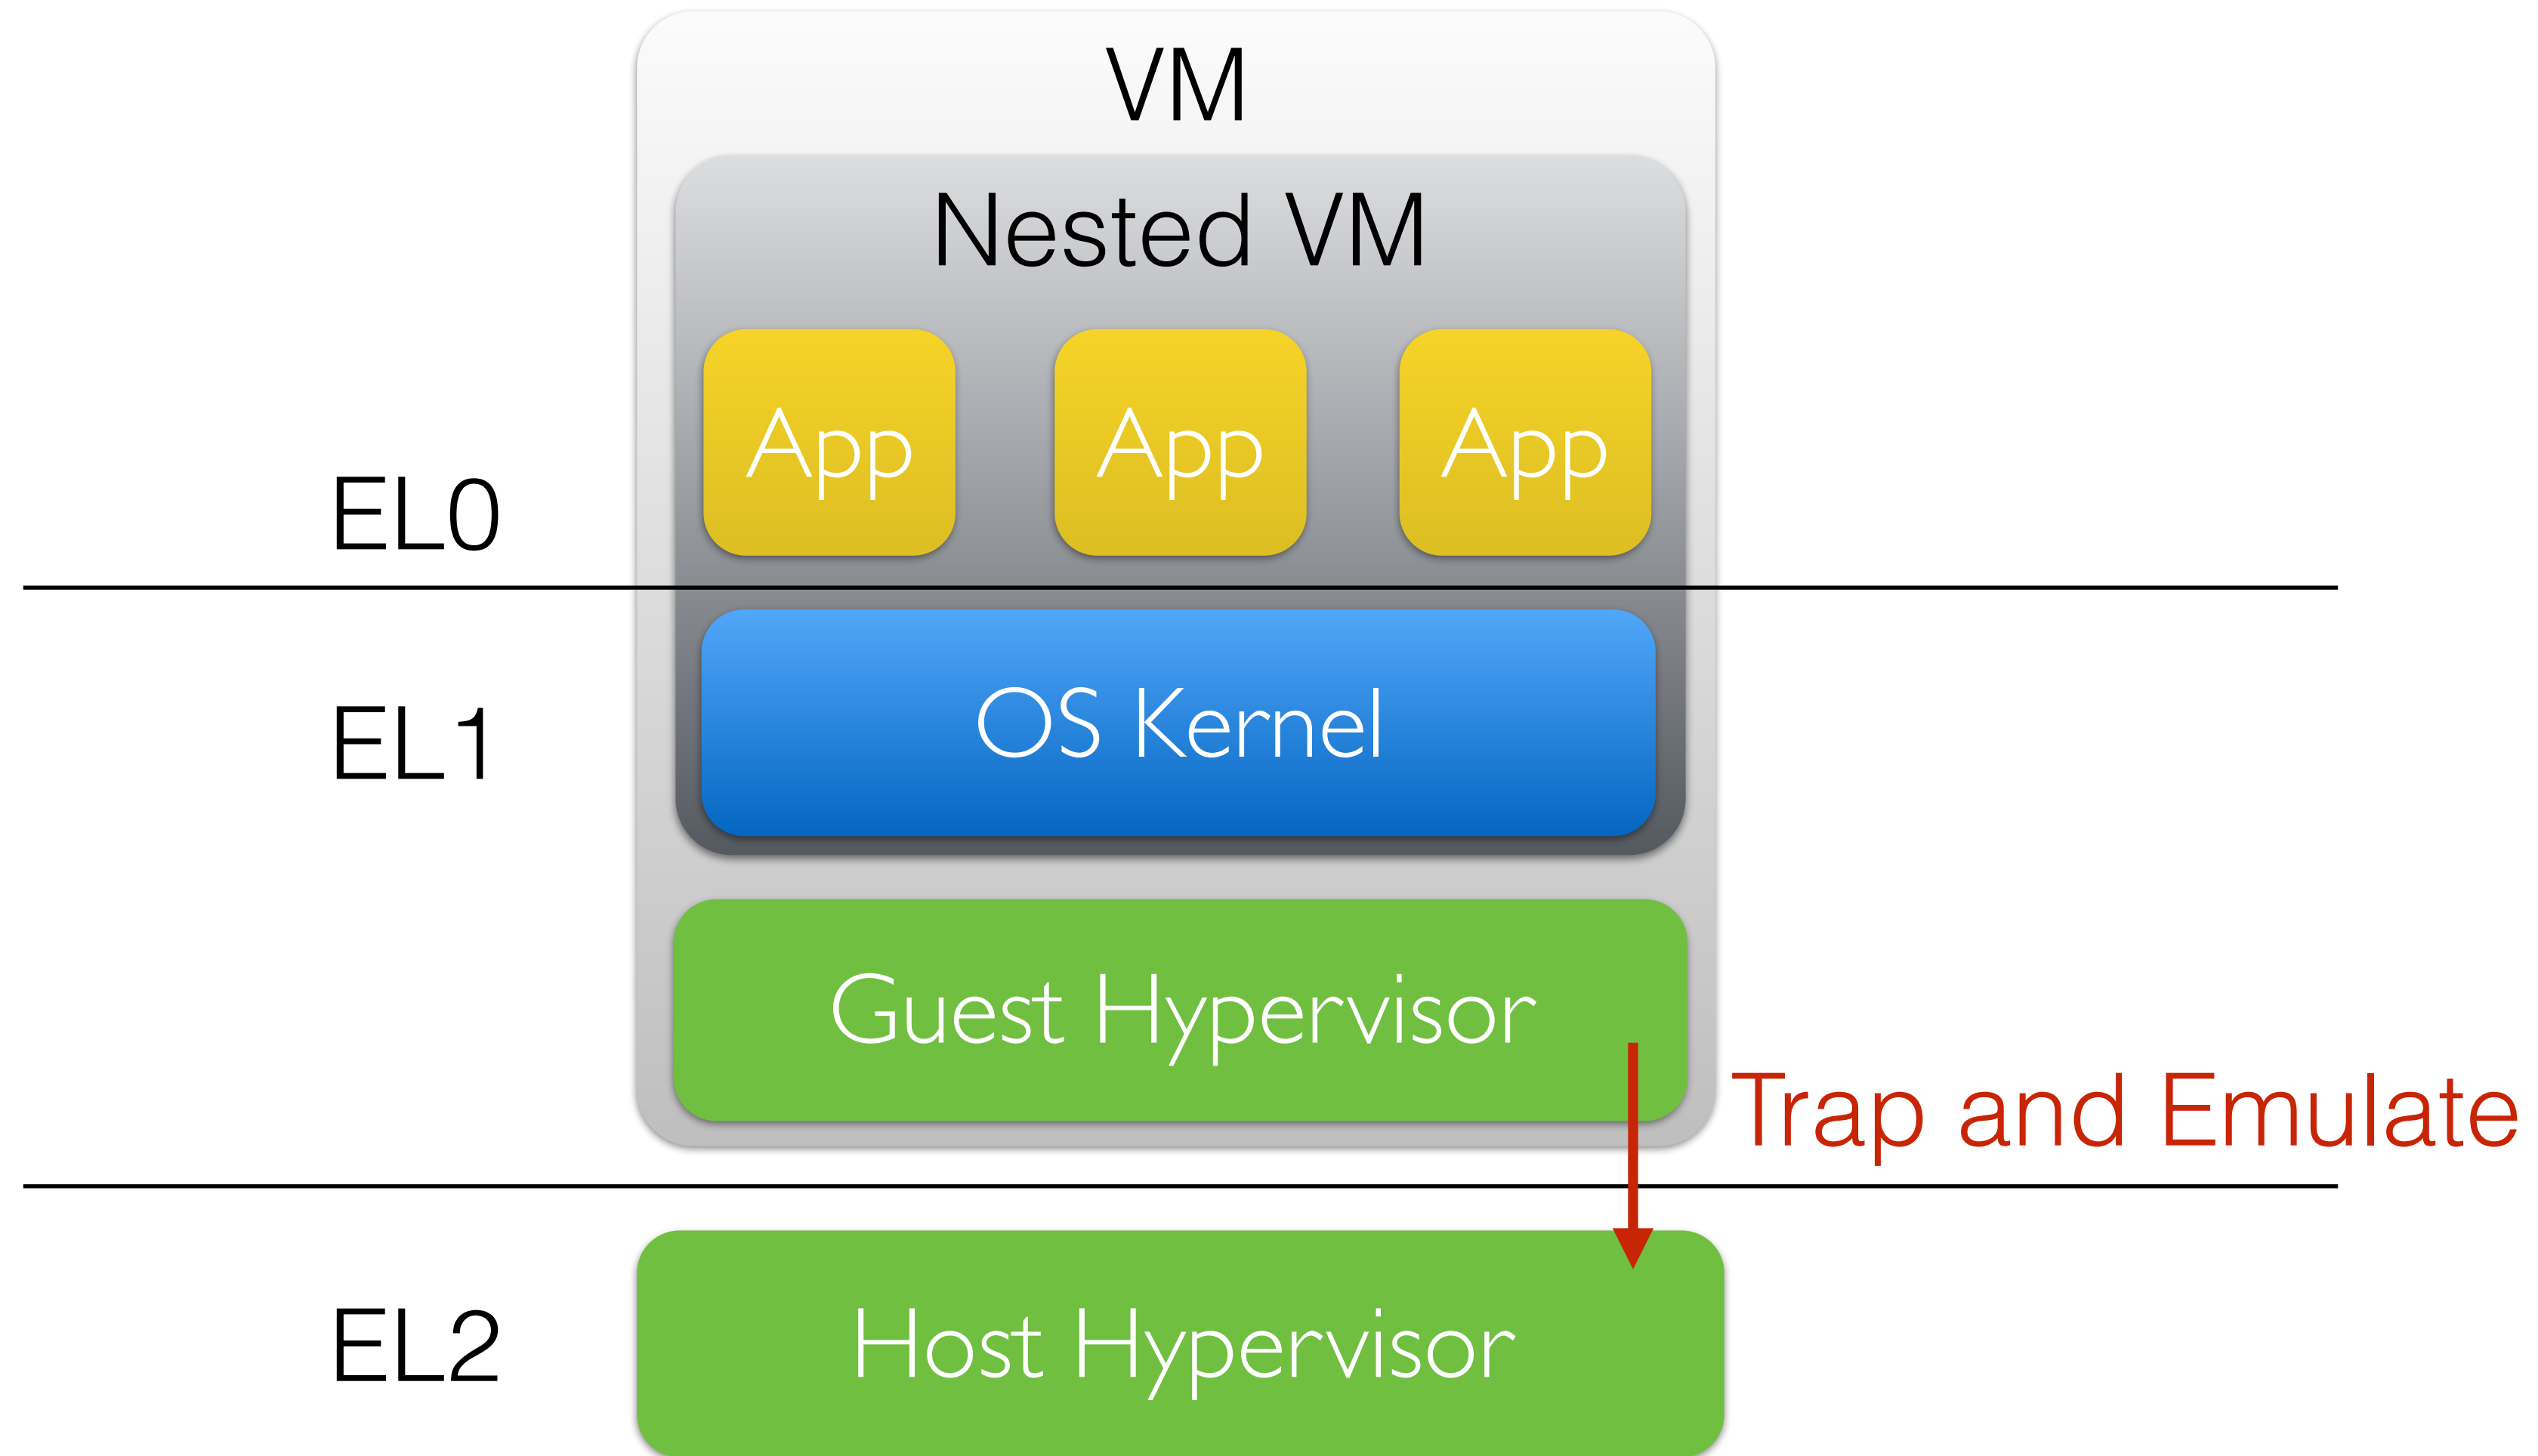
ARM Nested Virtualization



ARMv8.0



ARMv8.3

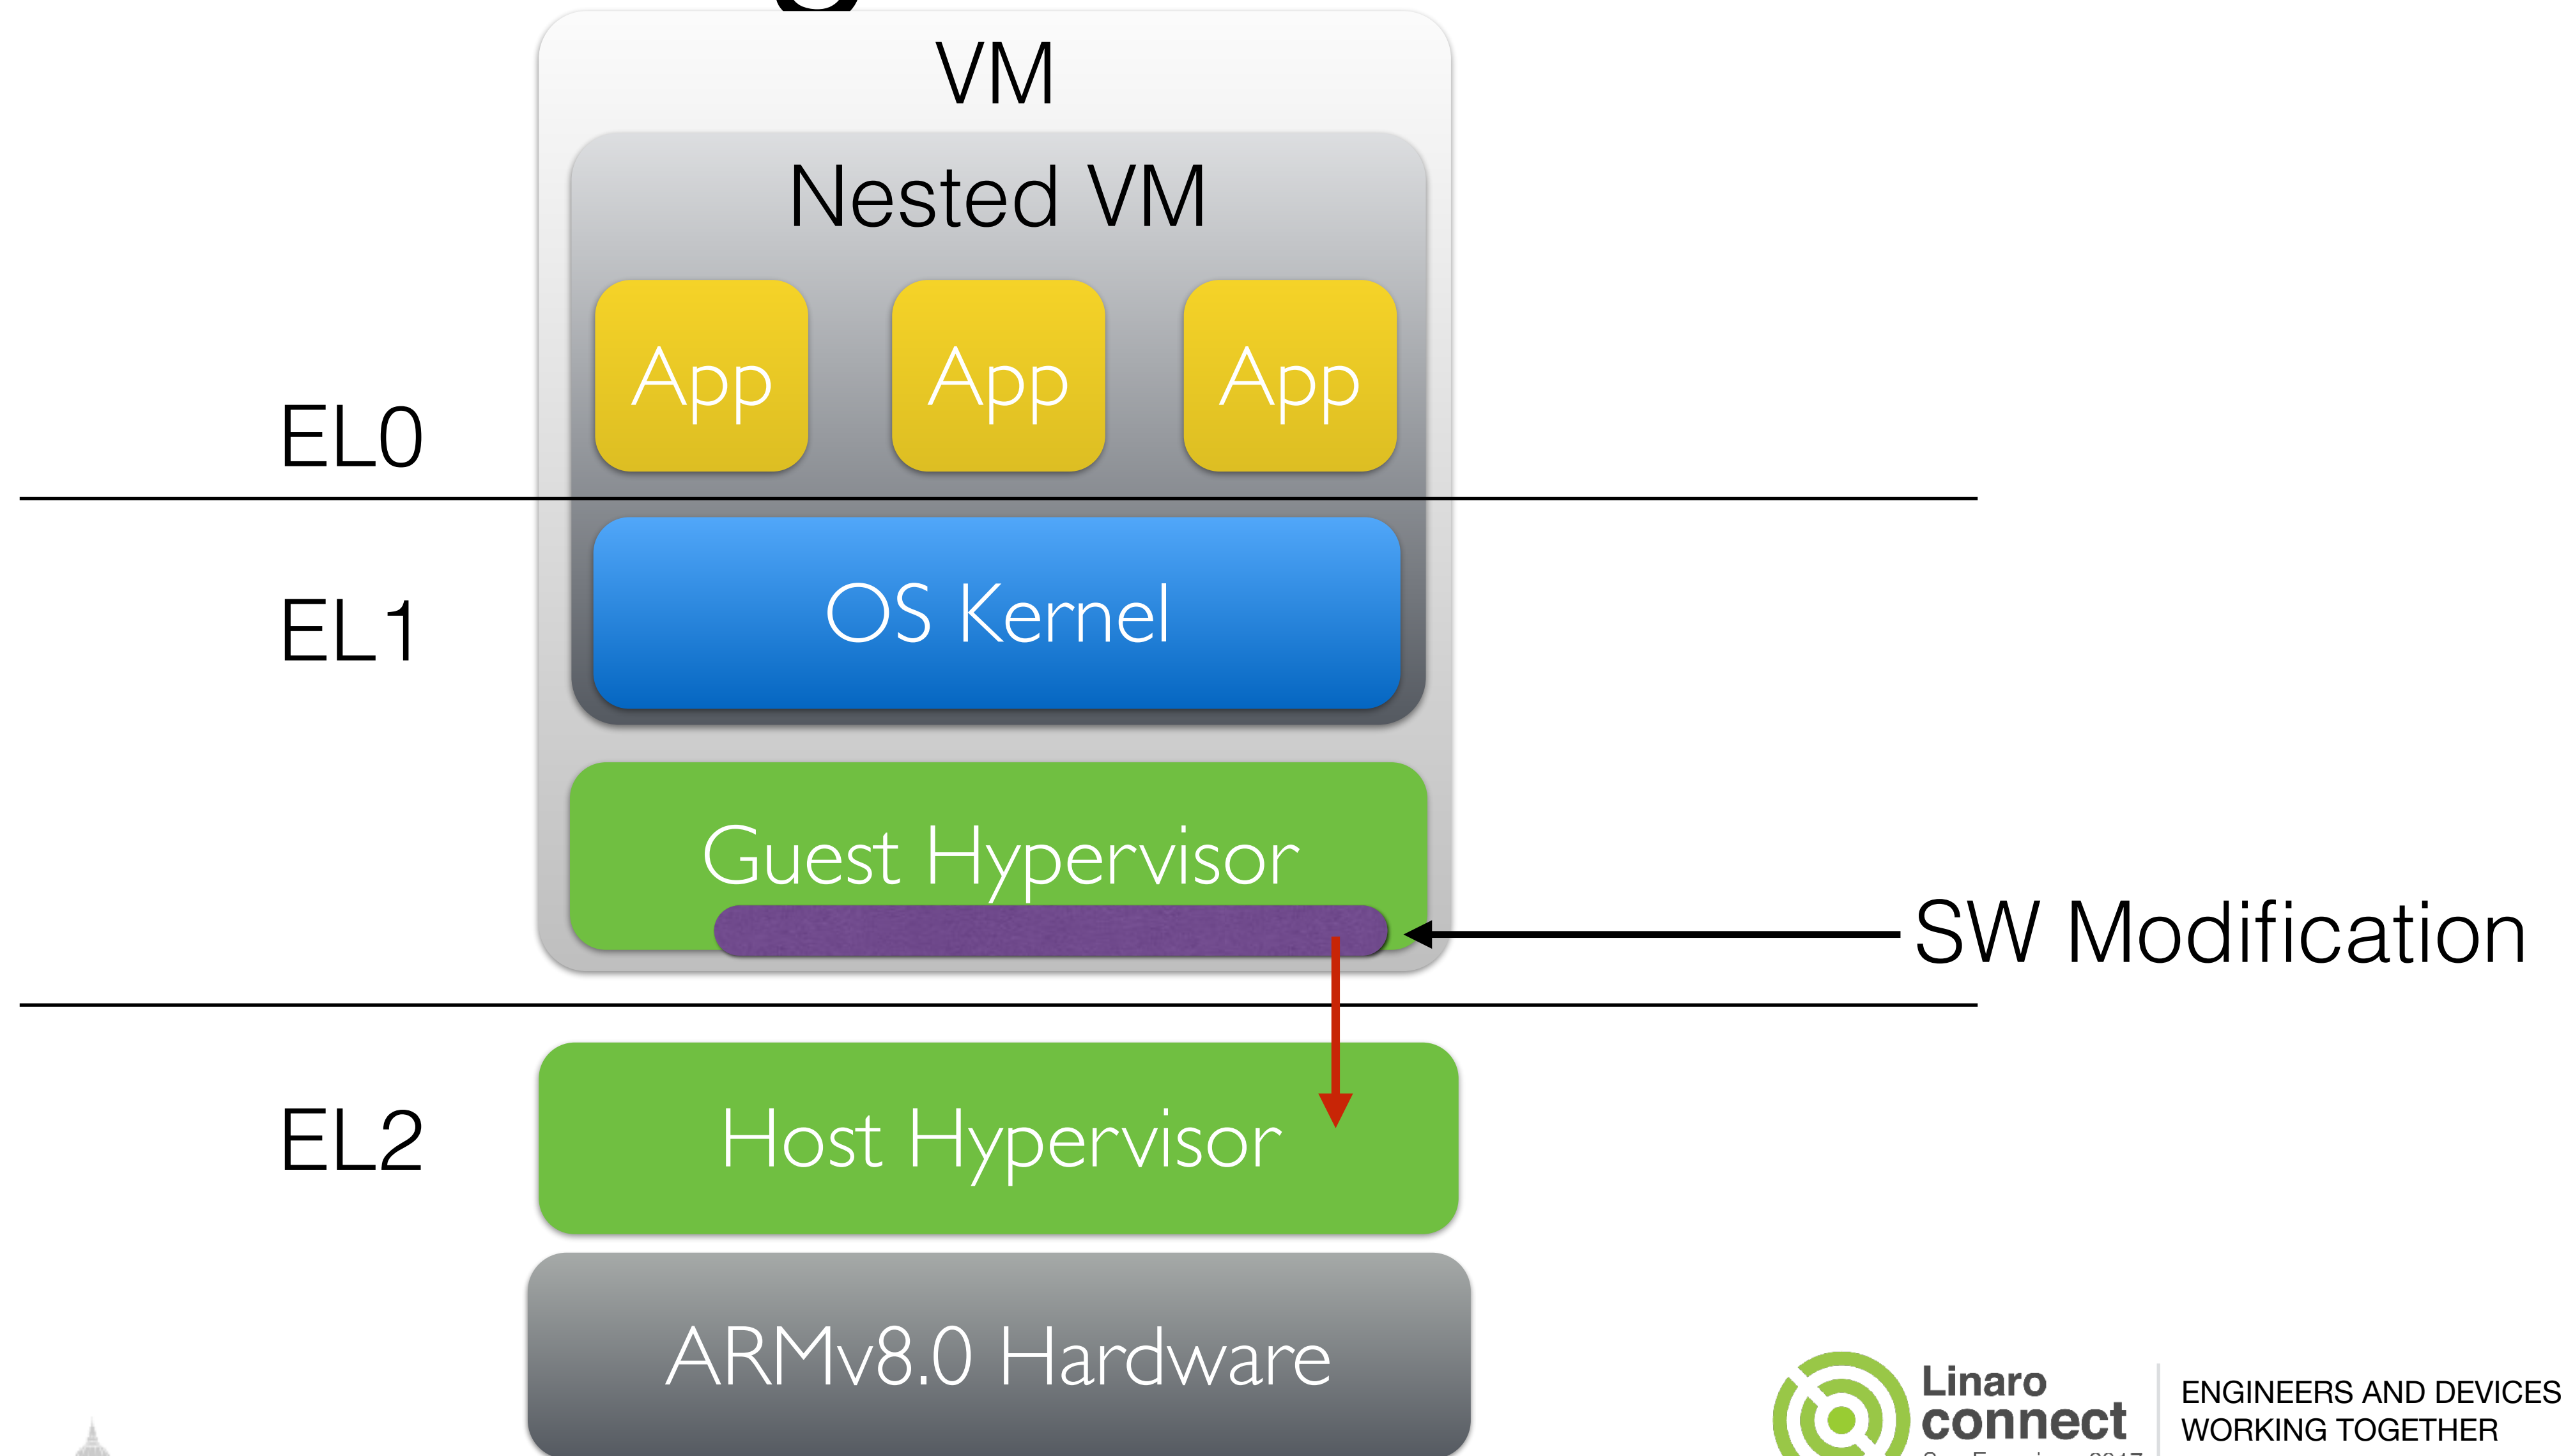


Performance Evaluation

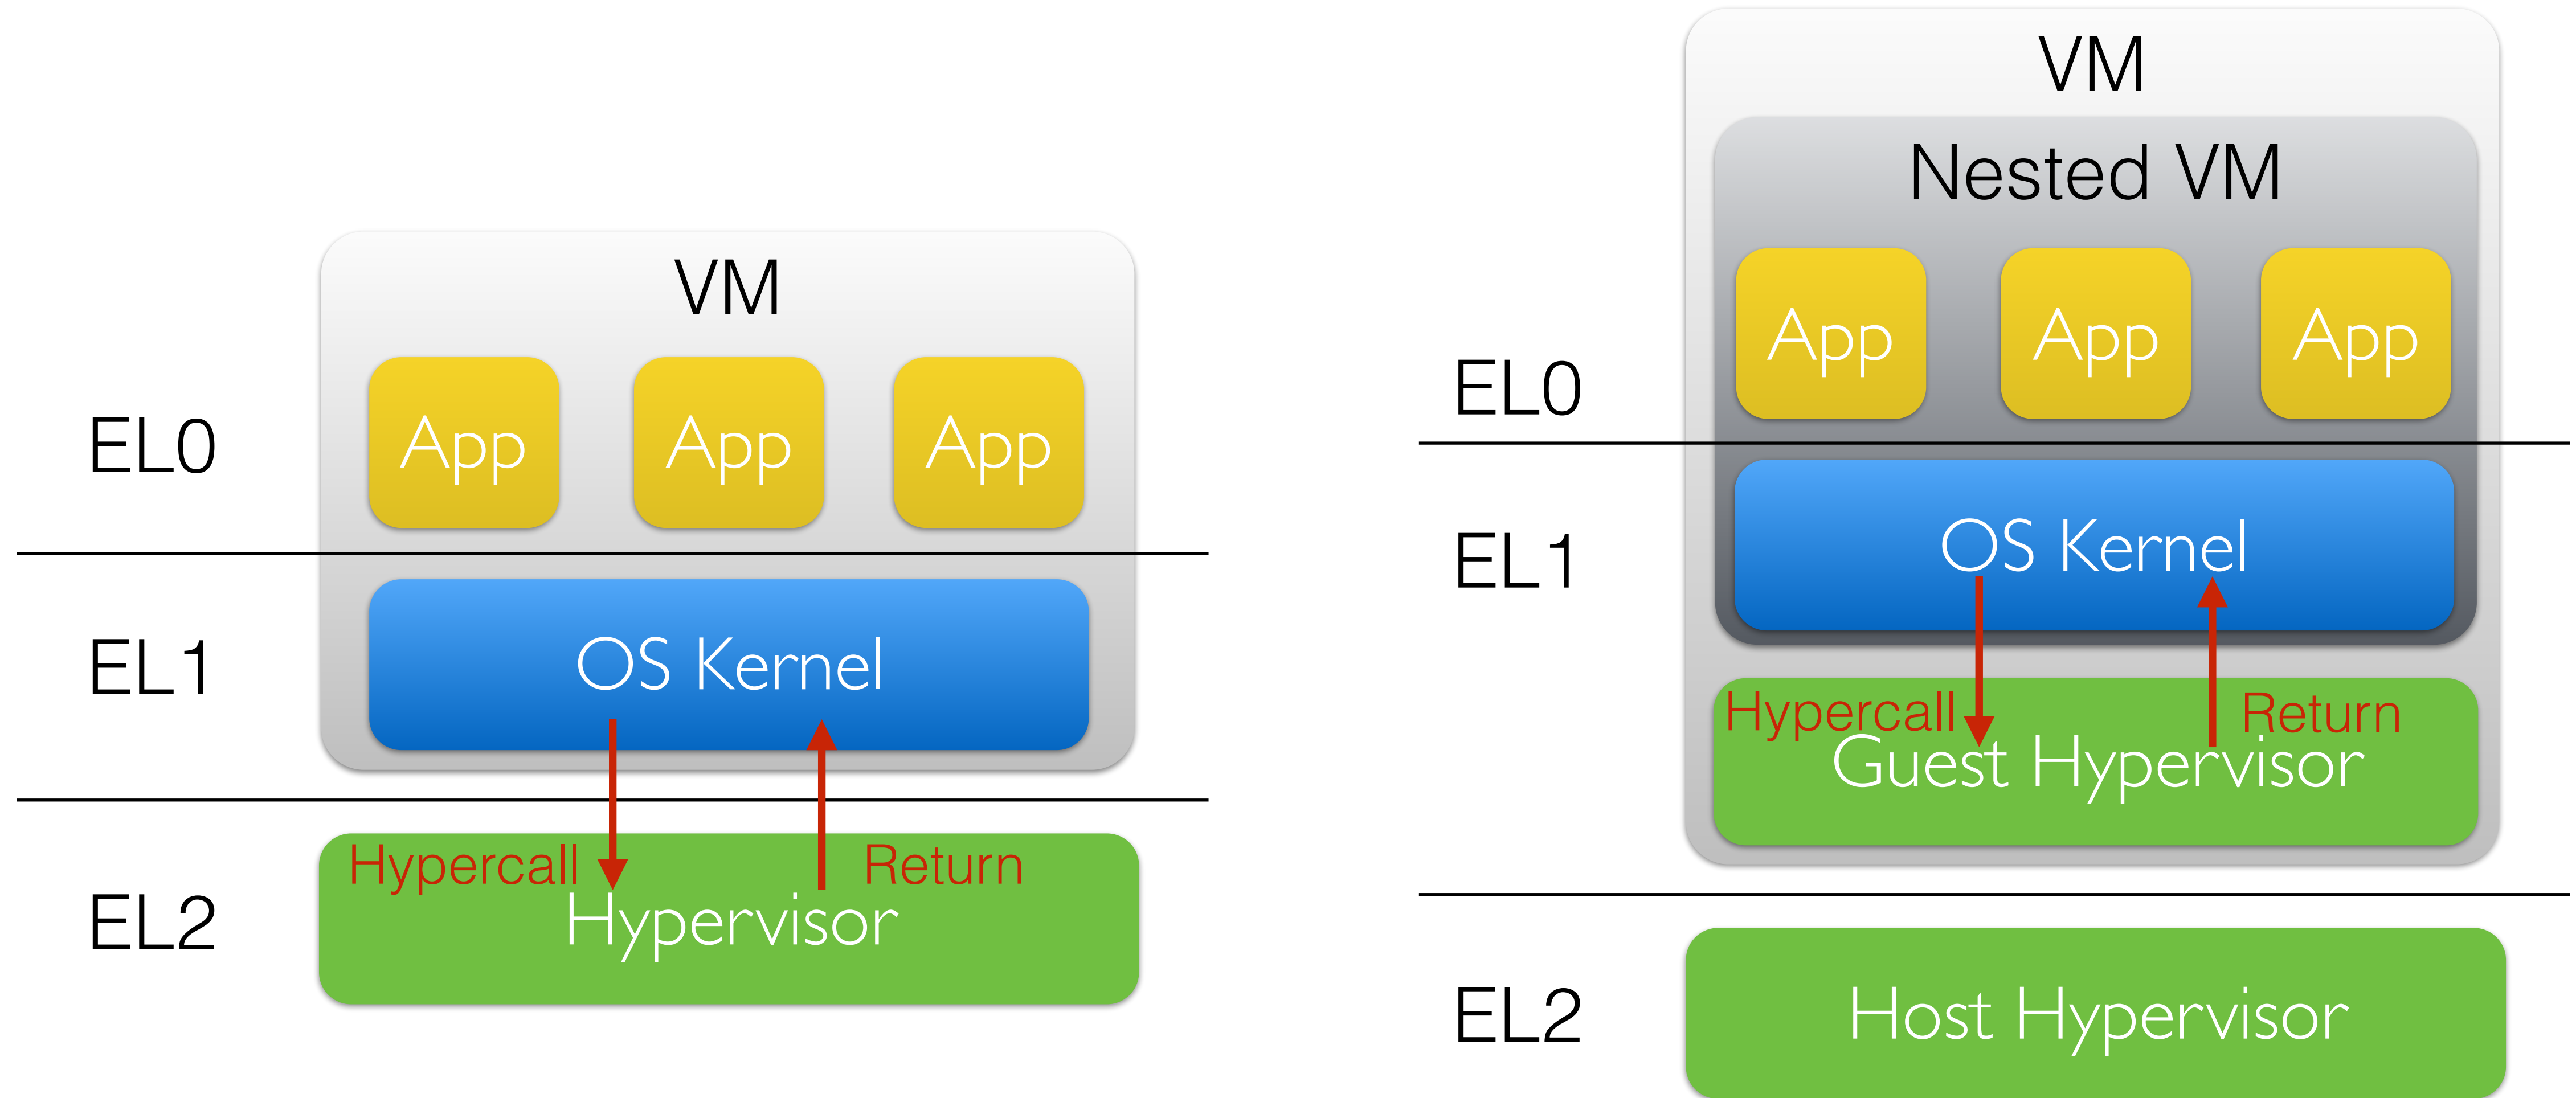
- Problem: No ARMv8.3 hardware available.
- Solution: Use ARMv8.0 hardware with the software modification



Emulating v8.3 on v8.0



Hypercall MicroBenchmark




Hypercall MicroBenchmark

	ARMv8.3	
	VM	Nested VM
Cycle counts	2,729	422,720
Ratio to VM	1	155x



Nested VM Hypercall

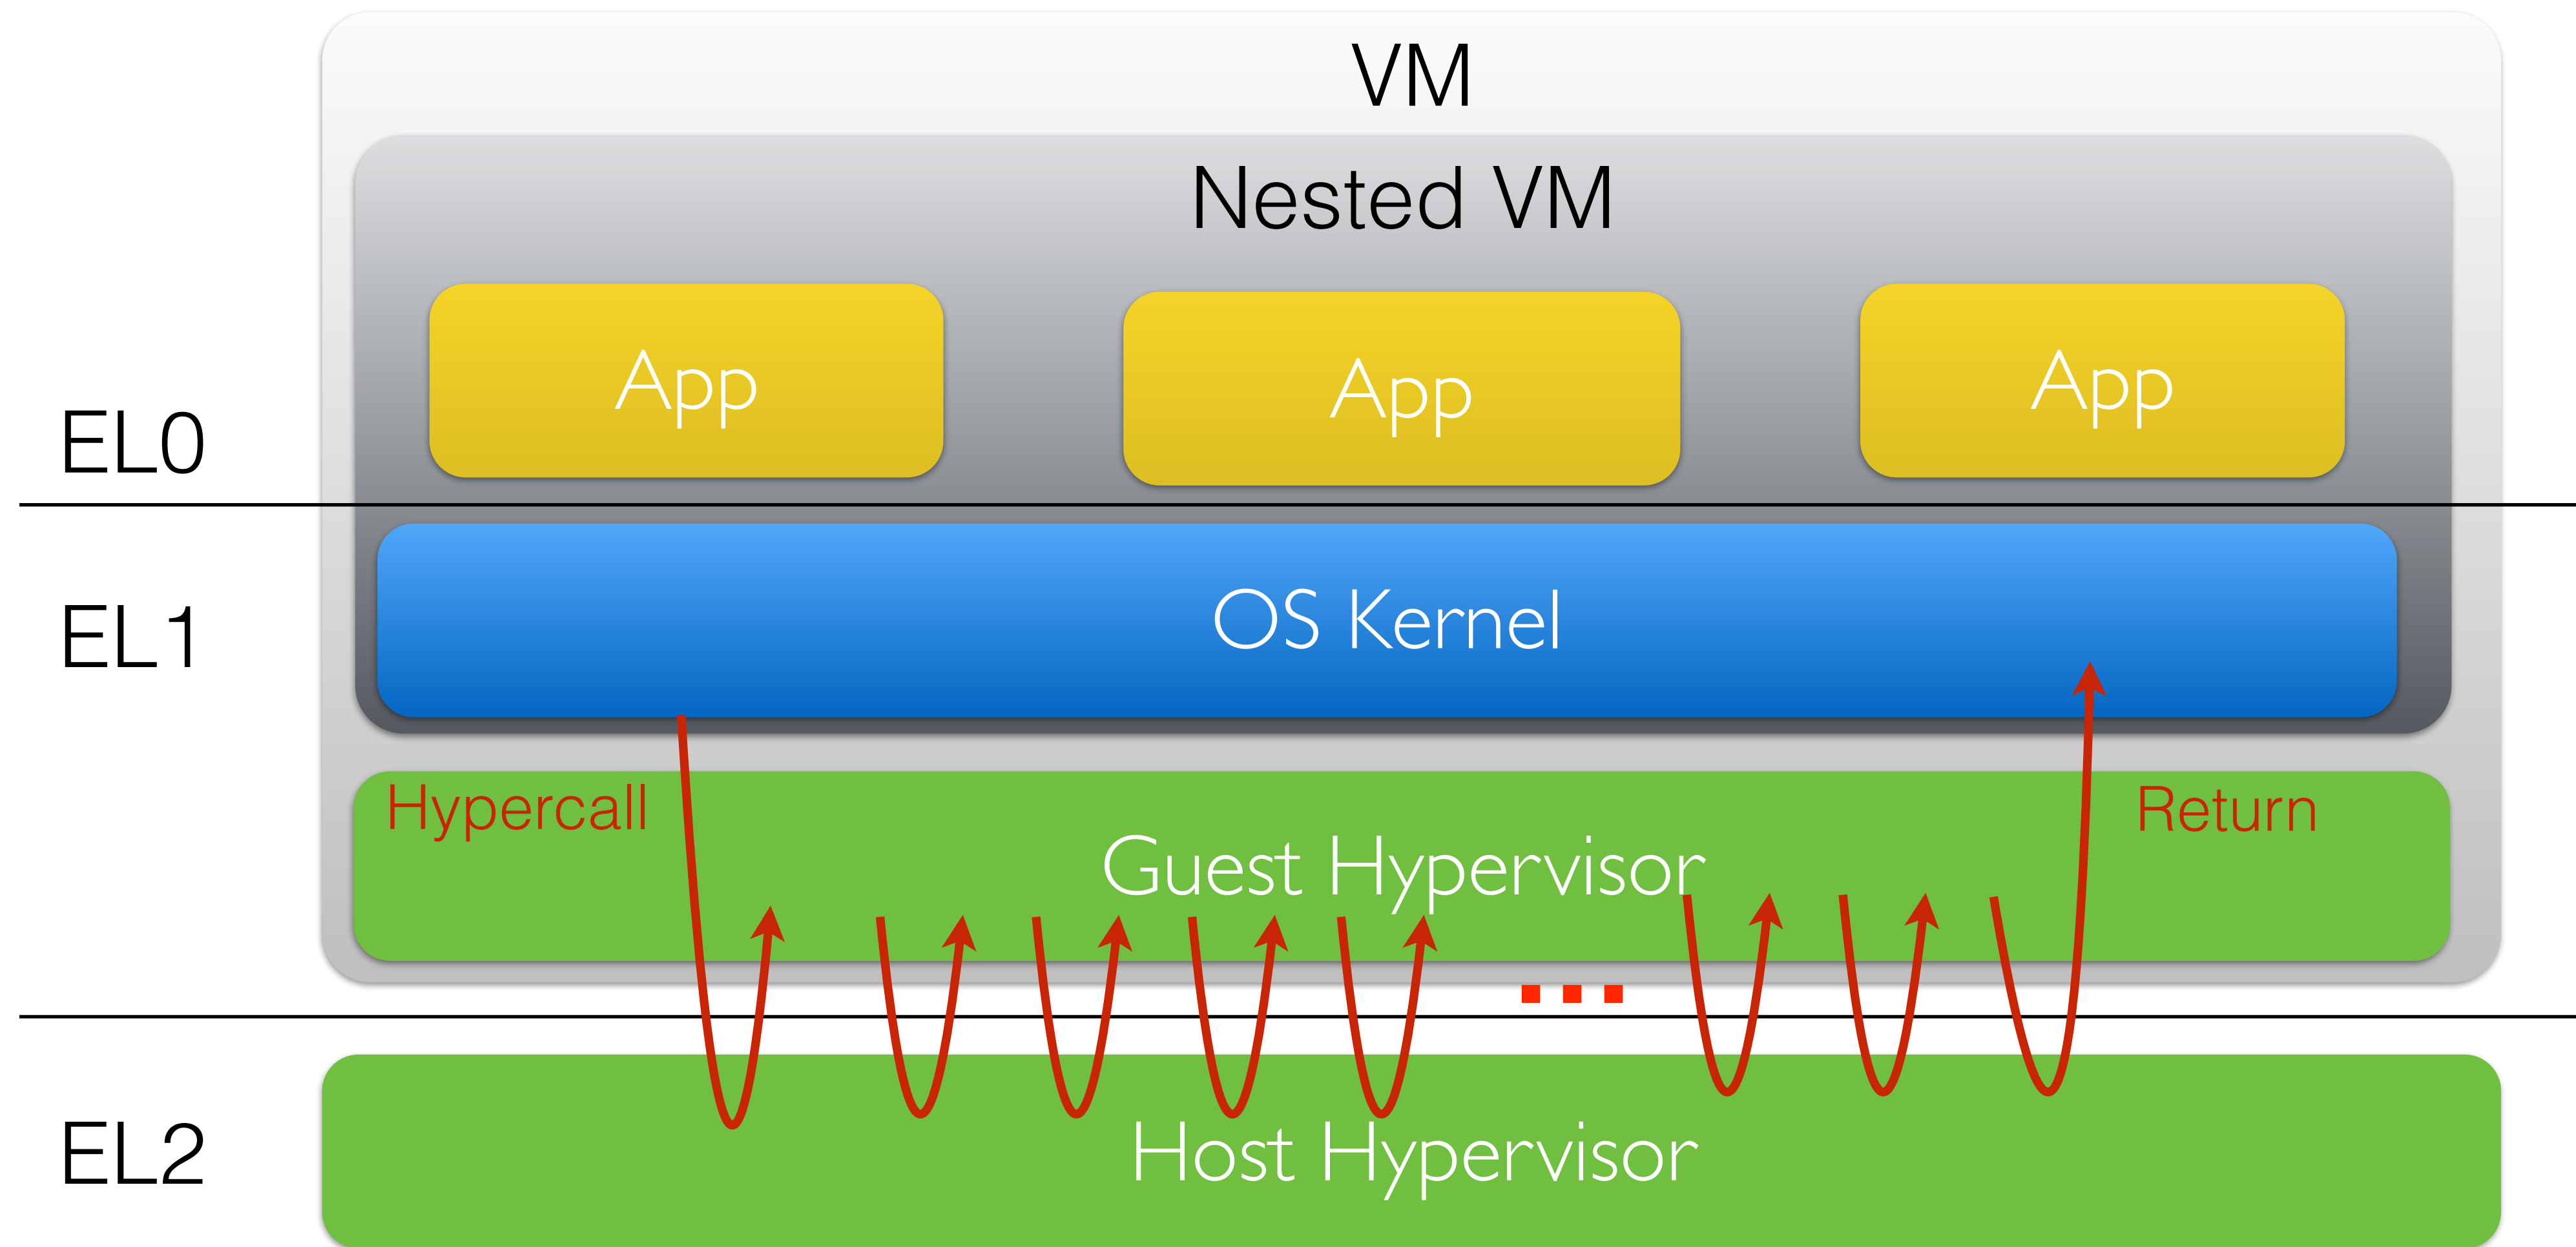


Hypercall MicroBenchmark

	ARMv8.3	
	VM	Nested VM
Cycle counts	2,729	422,720
Ratio to VM	1	155x
Trap counts	1	126



All Traps are Necessary?



NEVE: NEsted Virtualization Extensions for ARM

- Improves performance of nested virtualization
- Key Mechanisms
 1. Redirects register accesses to memory
 2. Redirects register accesses to other registers

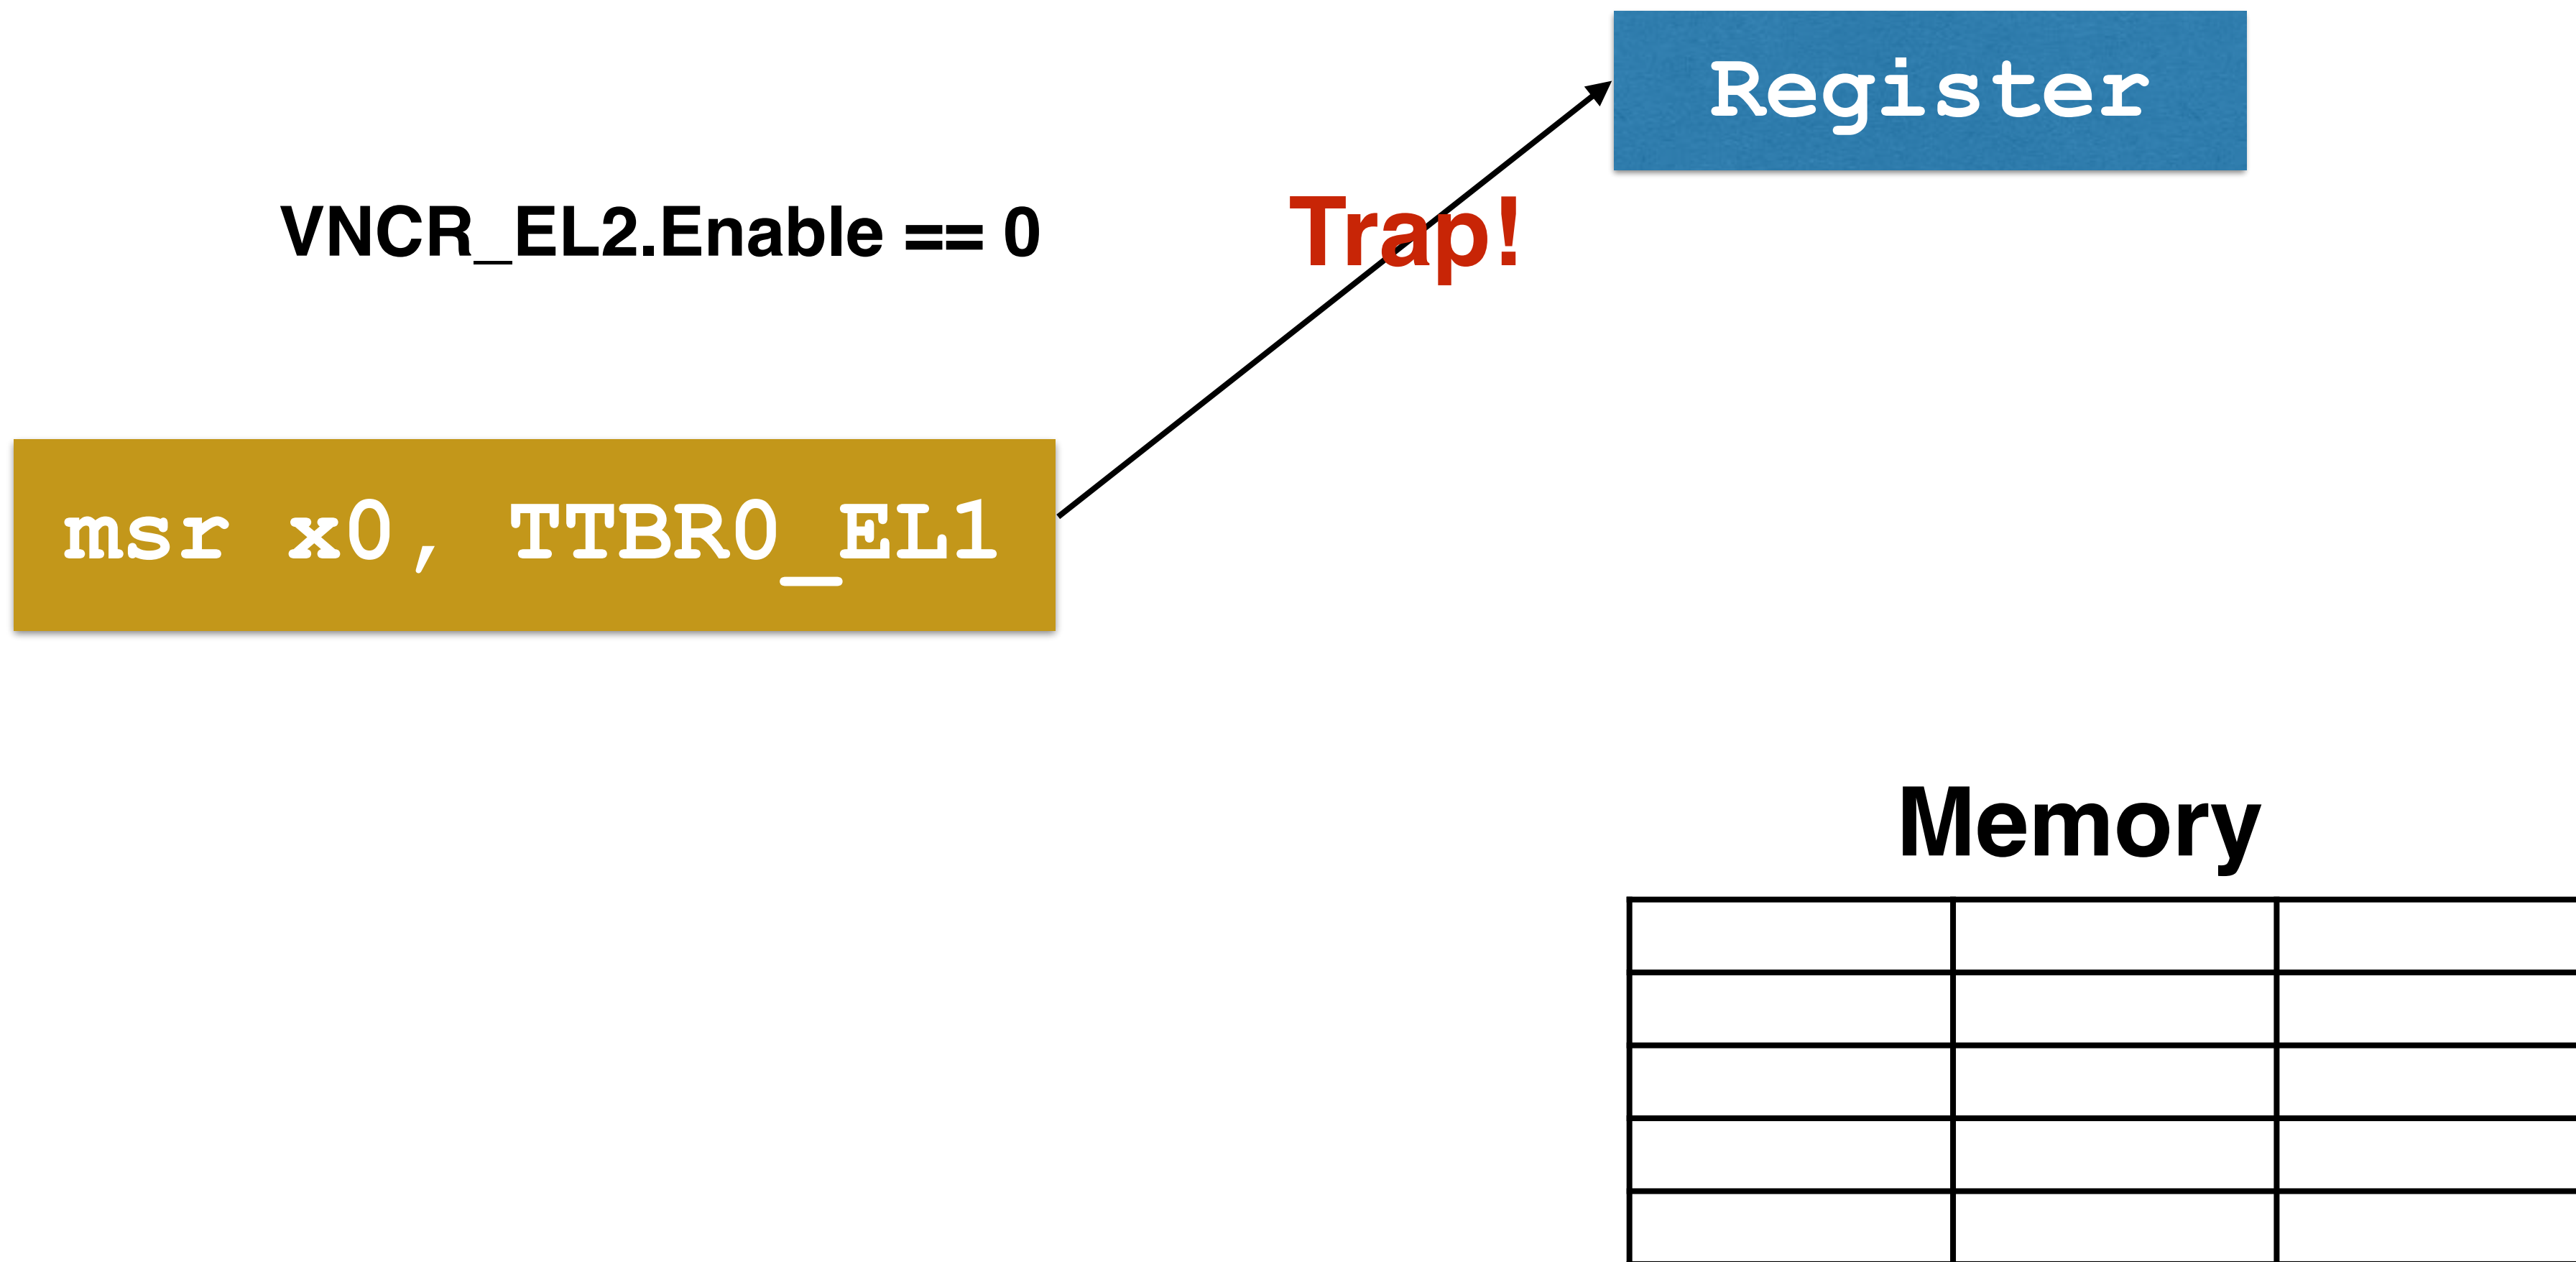


VNCR_EL2

- Virtual Nested Control Register
 - Enable bit
 - BADDR (Base Address)



#1 Redirection to Memory



#1 Redirection to Memory

Register

```
msr x0, TTBR0_EL1
```

VNCR_EL2.Enable == 1


VNCR_EL2.BADDR

Memory

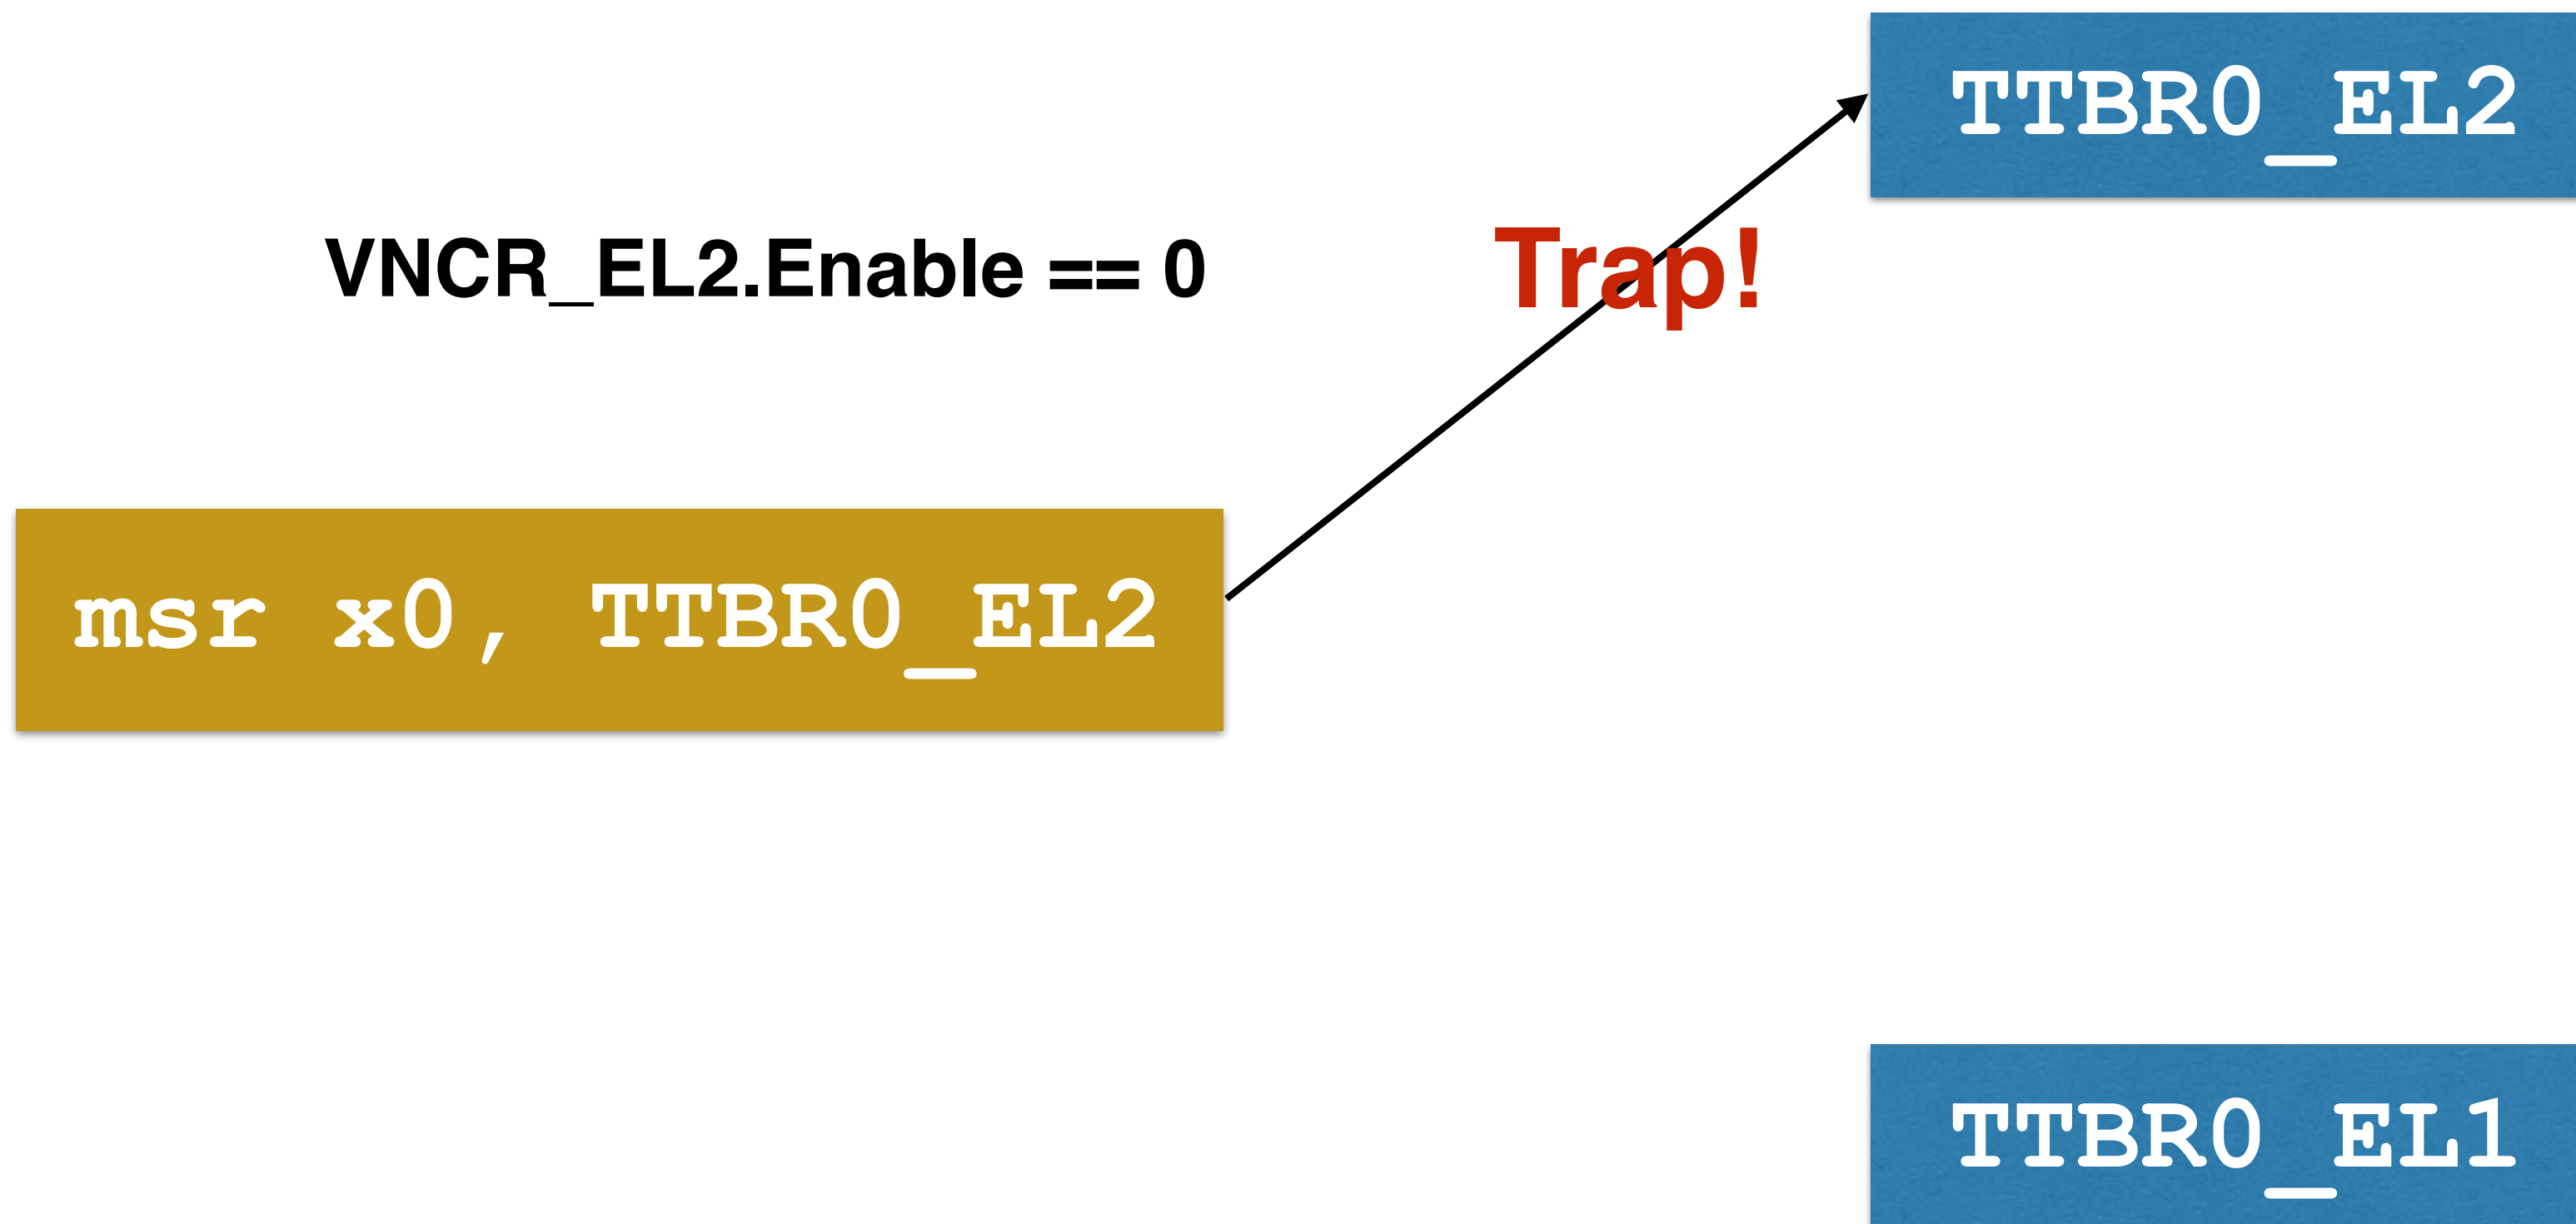
	TTBR0_EL1	



Accessing EL2 State



#2 Redirection to Register




#2 Redirection to Register

VNCR_EL2.Enable == 1

TTBR0_EL2

```
msr x0, TTBR0_EL2
```

TTBR0_EL1



Hypercall MicroBenchmark

	ARMv8.3		NEVE
	VM	Nested VM	Nested VM
Cycle counts	2,729	422,720	92,385
Ratio to VM		155x	34x
Trap counts	1	126	15



Application Workloads

Application	Description	Application	Description
Kernbench	Kernel compile	Netperf TCP_RR	Network performance
Hackbench	Scheduler stress	Netperf TCP STREAM	Network performance
SPECjvm2008	Java Runtime	Netperf TCP MAERTS	Network performance
MySQL	Database management	Apache	Web server stress
Memcached	Key-Value store	Nginx	Web server stress



Experimental Setup

- Hardware

- APM X-Gene (ARMv8.0)
- 8-way SMP
- 64 GB RAM
- 10 Gb Ethernet

- Native/VM/Nested VM Setup

- 4-way SMP
- 12 GB RAM
- Virt I/O (VM/nested VM)

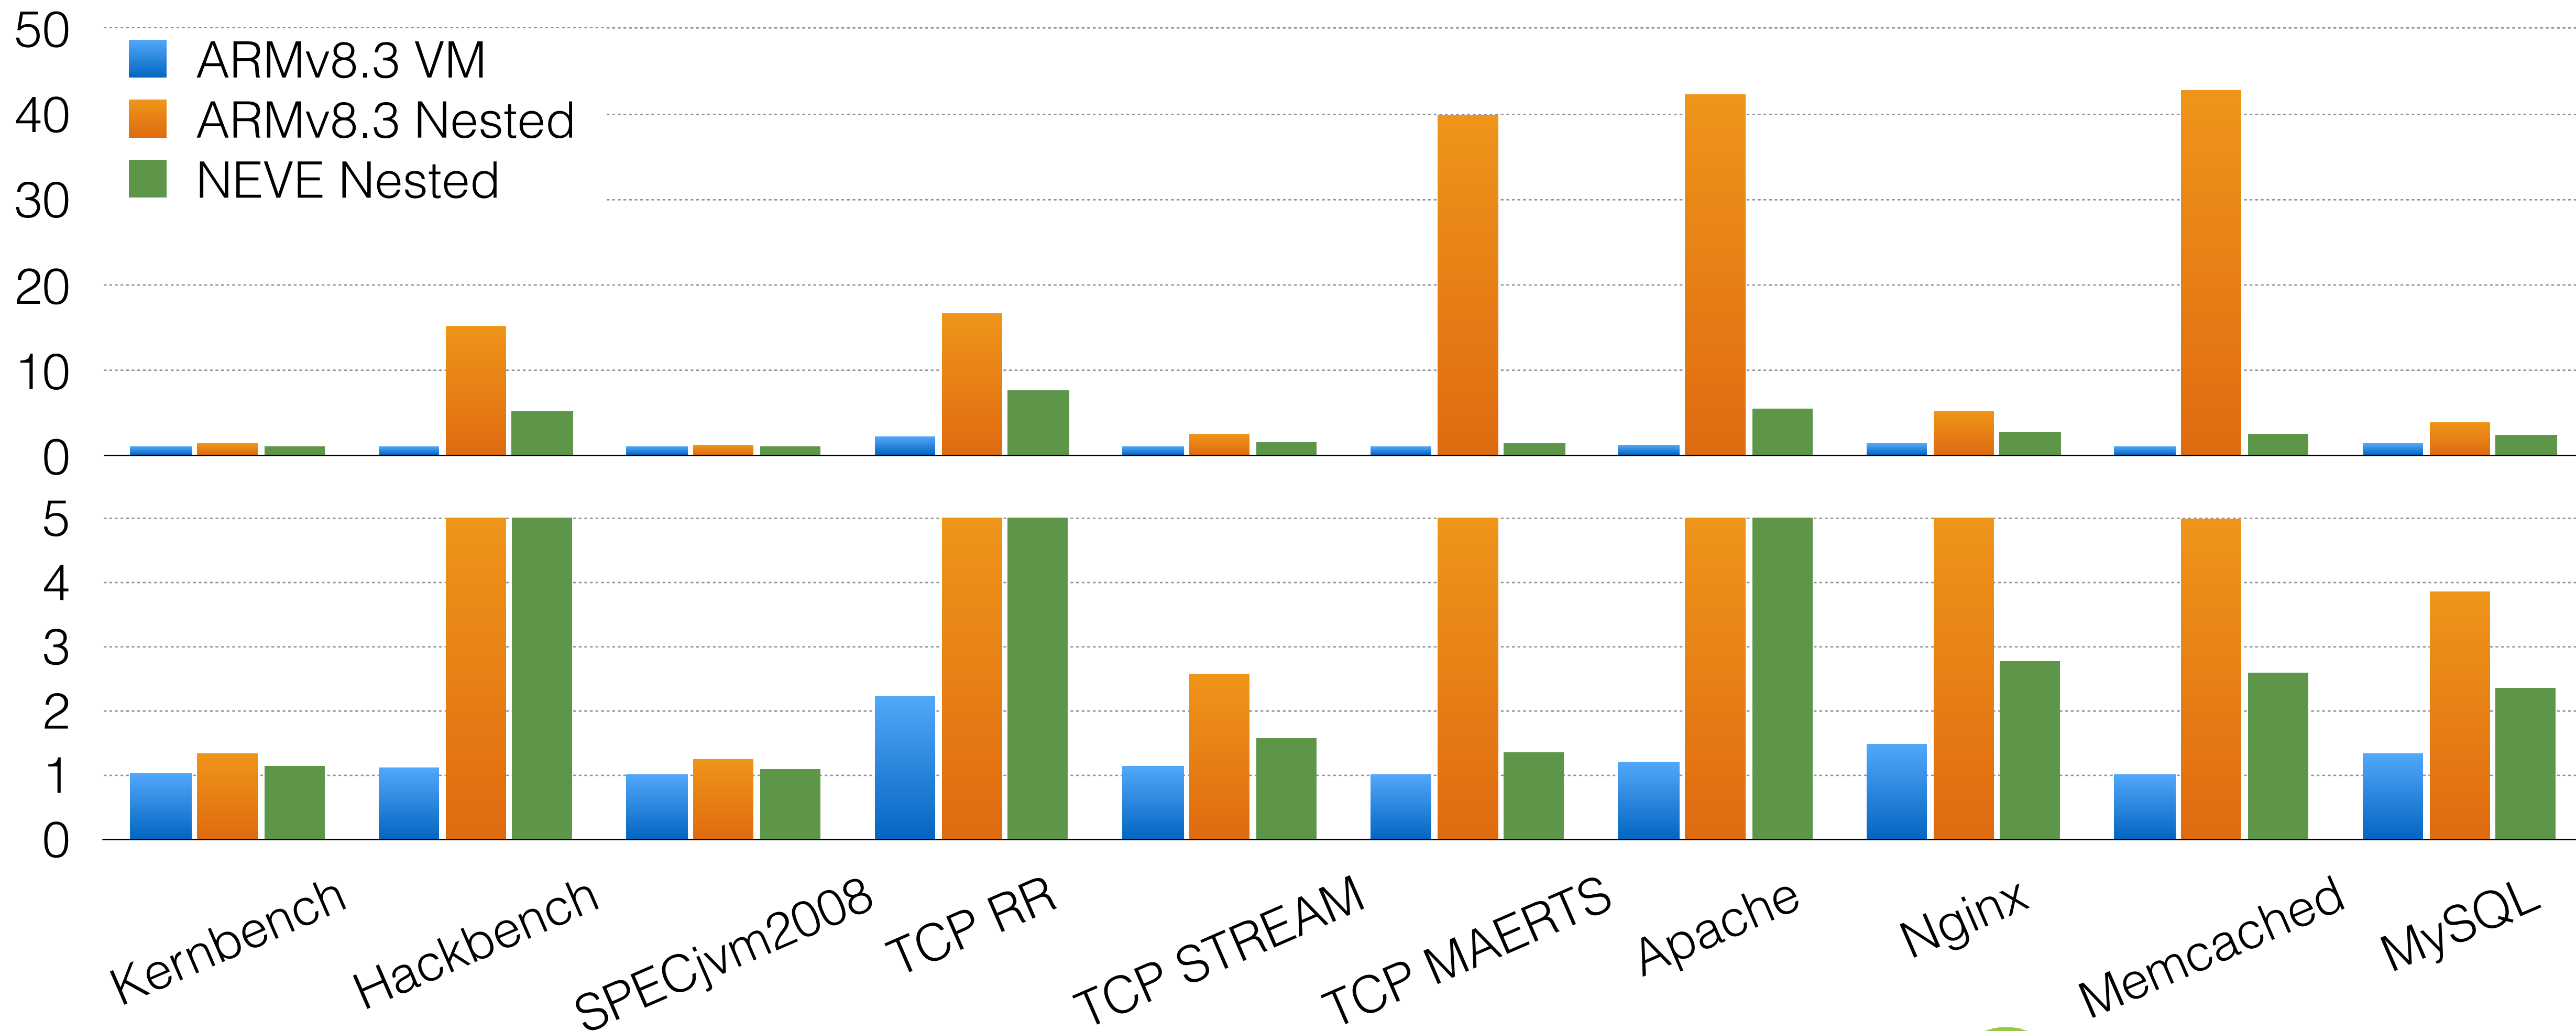
- Software

- KVM on KVM
- v4.10




Application Benchmarks

Normalized overhead
(lower is better)



Application Benchmarks

Normalized overhead
(lower is better)



Conclusion

- Nested virtualization performance on ARMv8.3 incurs high overhead
- Due to the exit multiplication problem
- NEVE enhances performance significantly by reducing number of traps
- NEVE is included in ARMv8.4

