

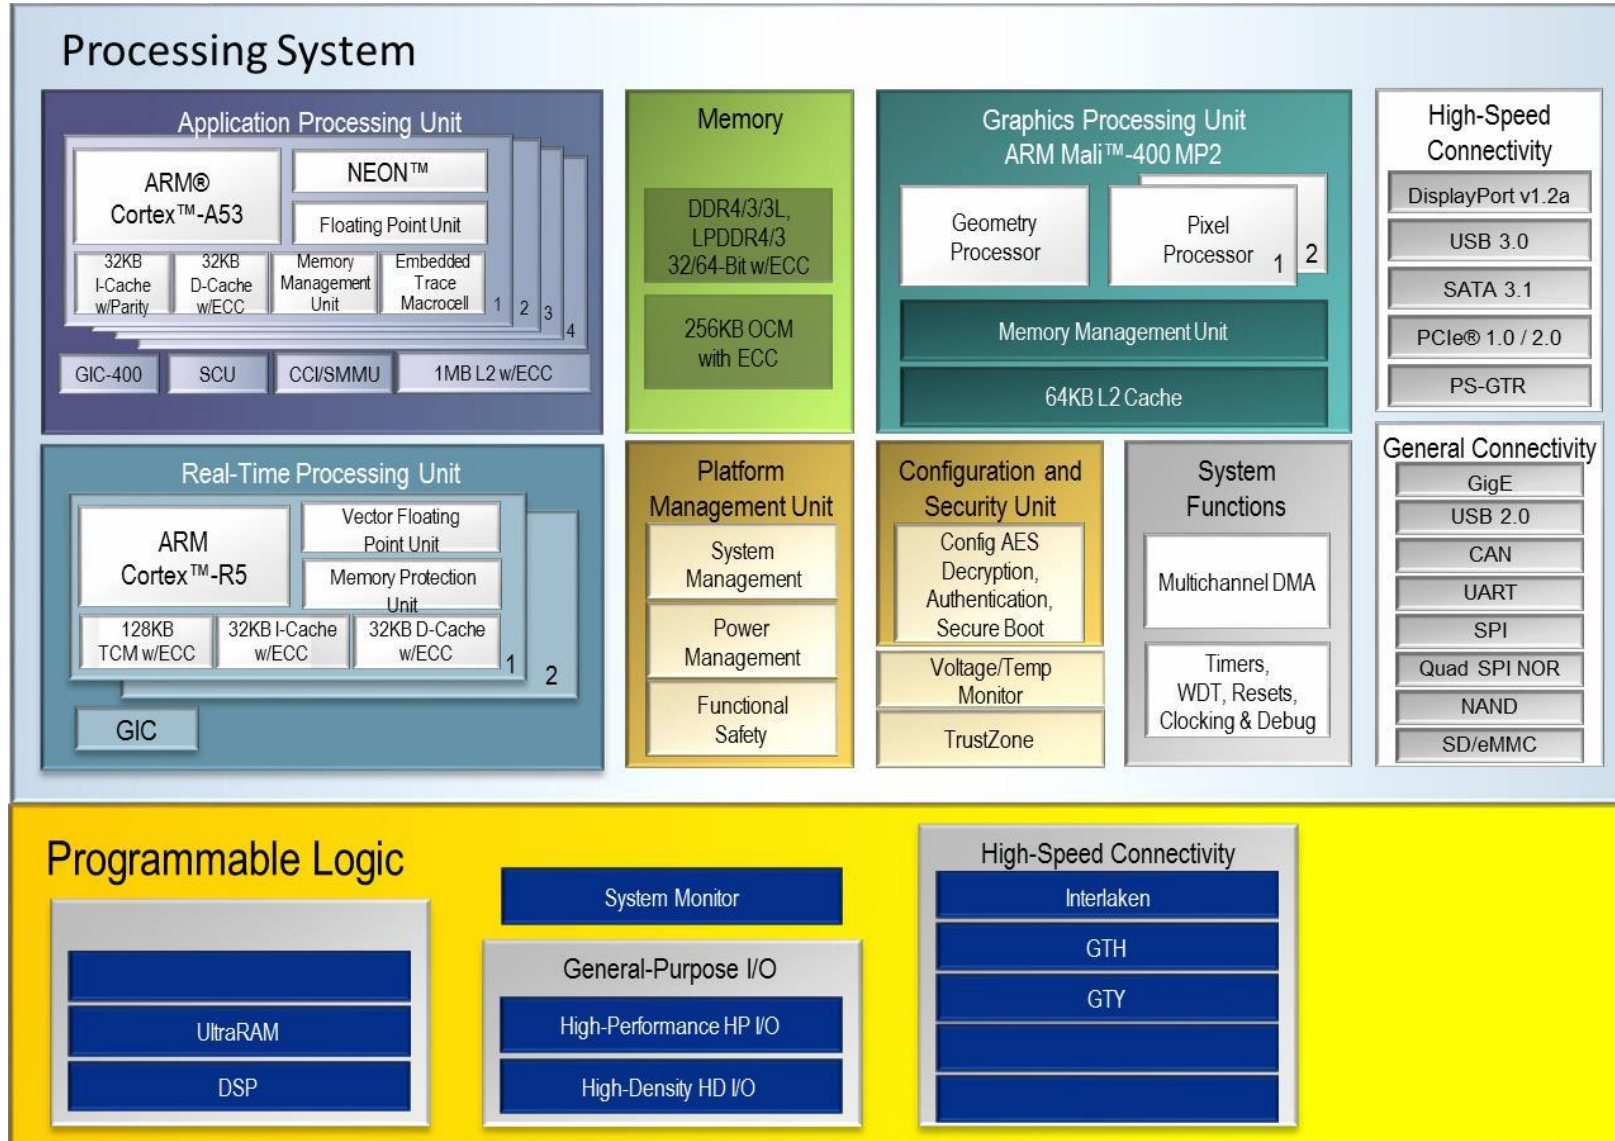
Static Partitioning with Xen

Dom0-less Xen configurations

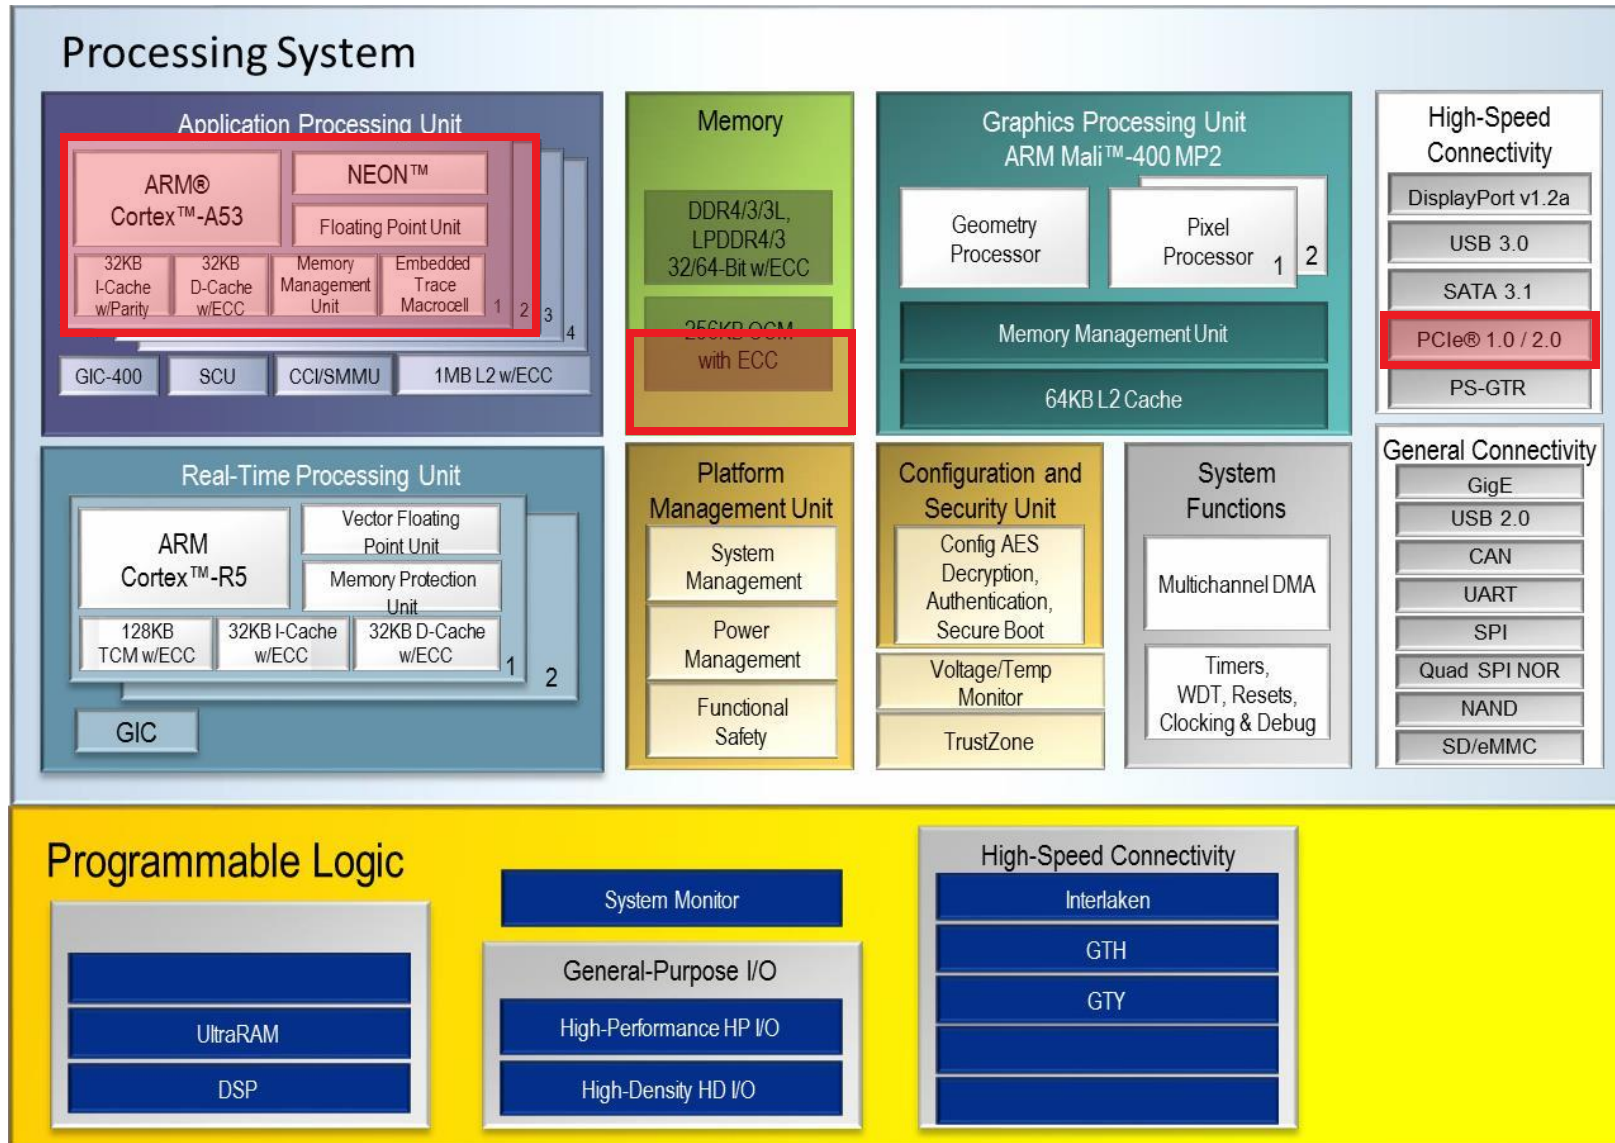
Stefano Stabellini – ELC August 2019



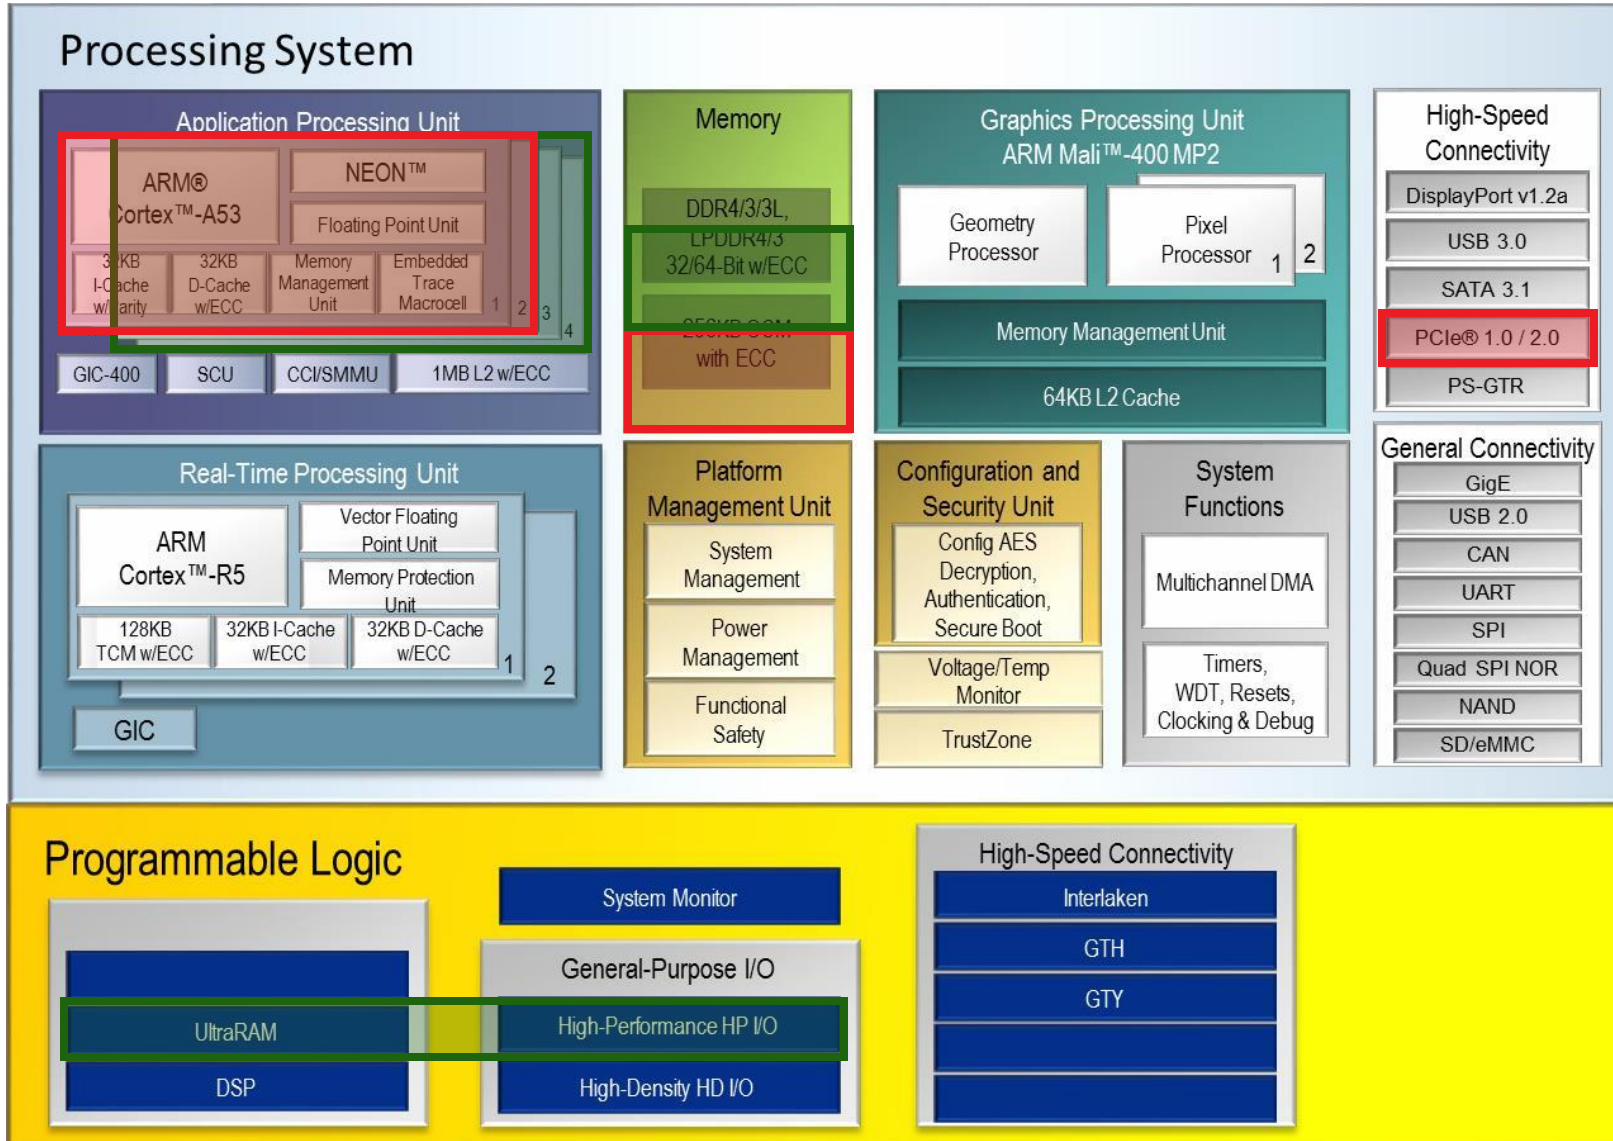
Static Partitioning



Static Partitioning



Static Partitioning

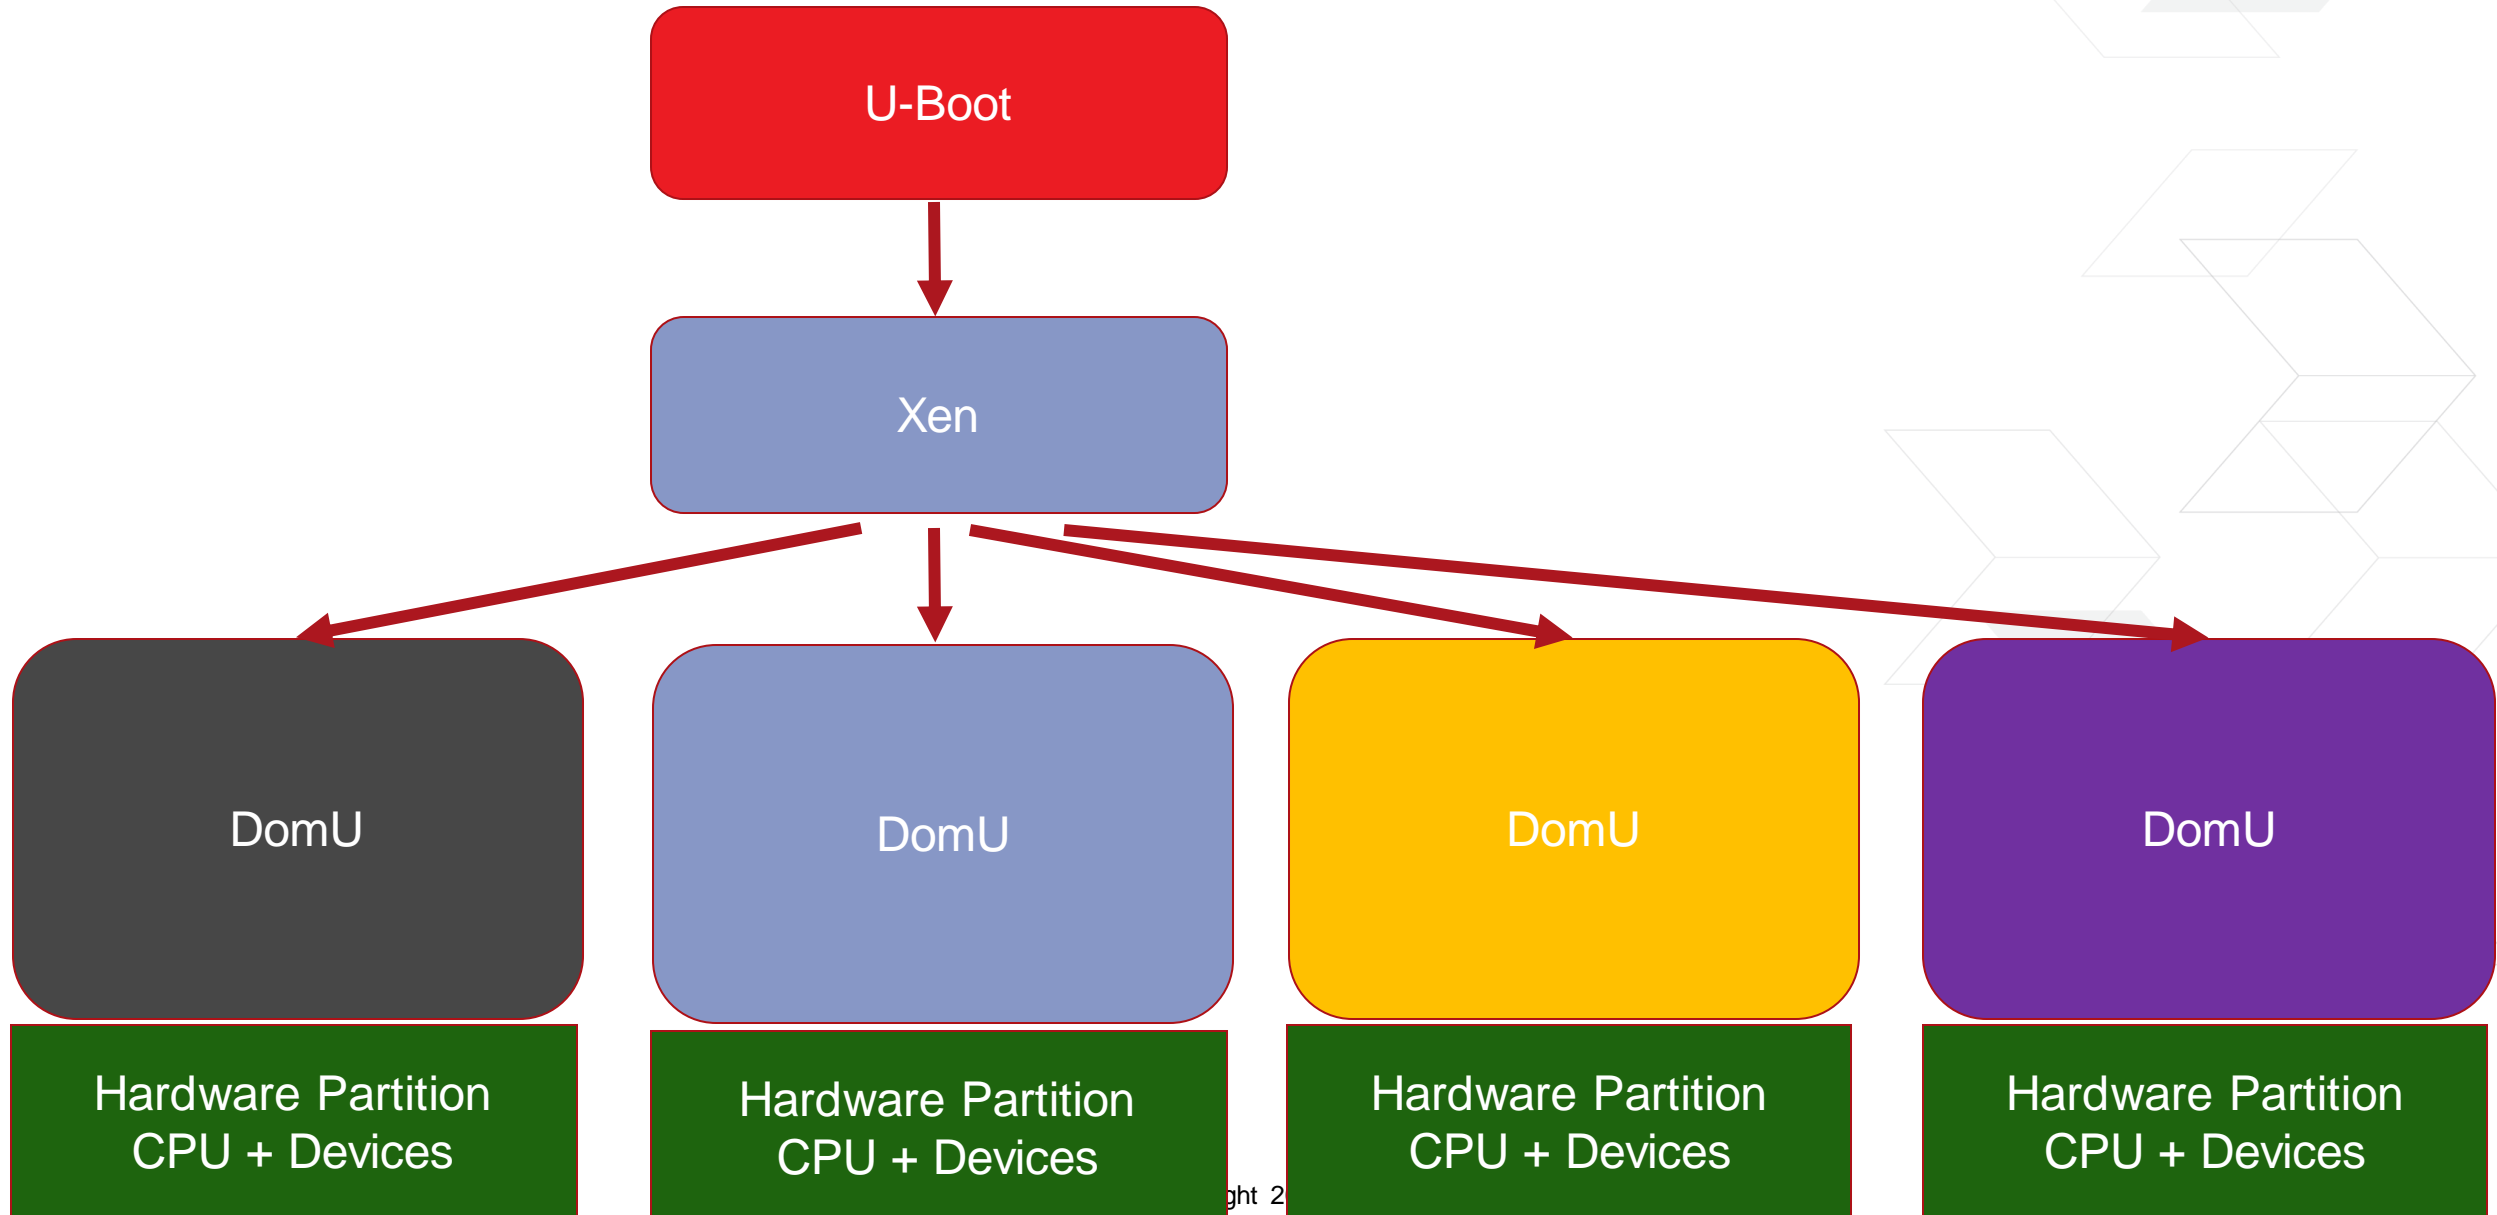


Static Partitioning: the use-case

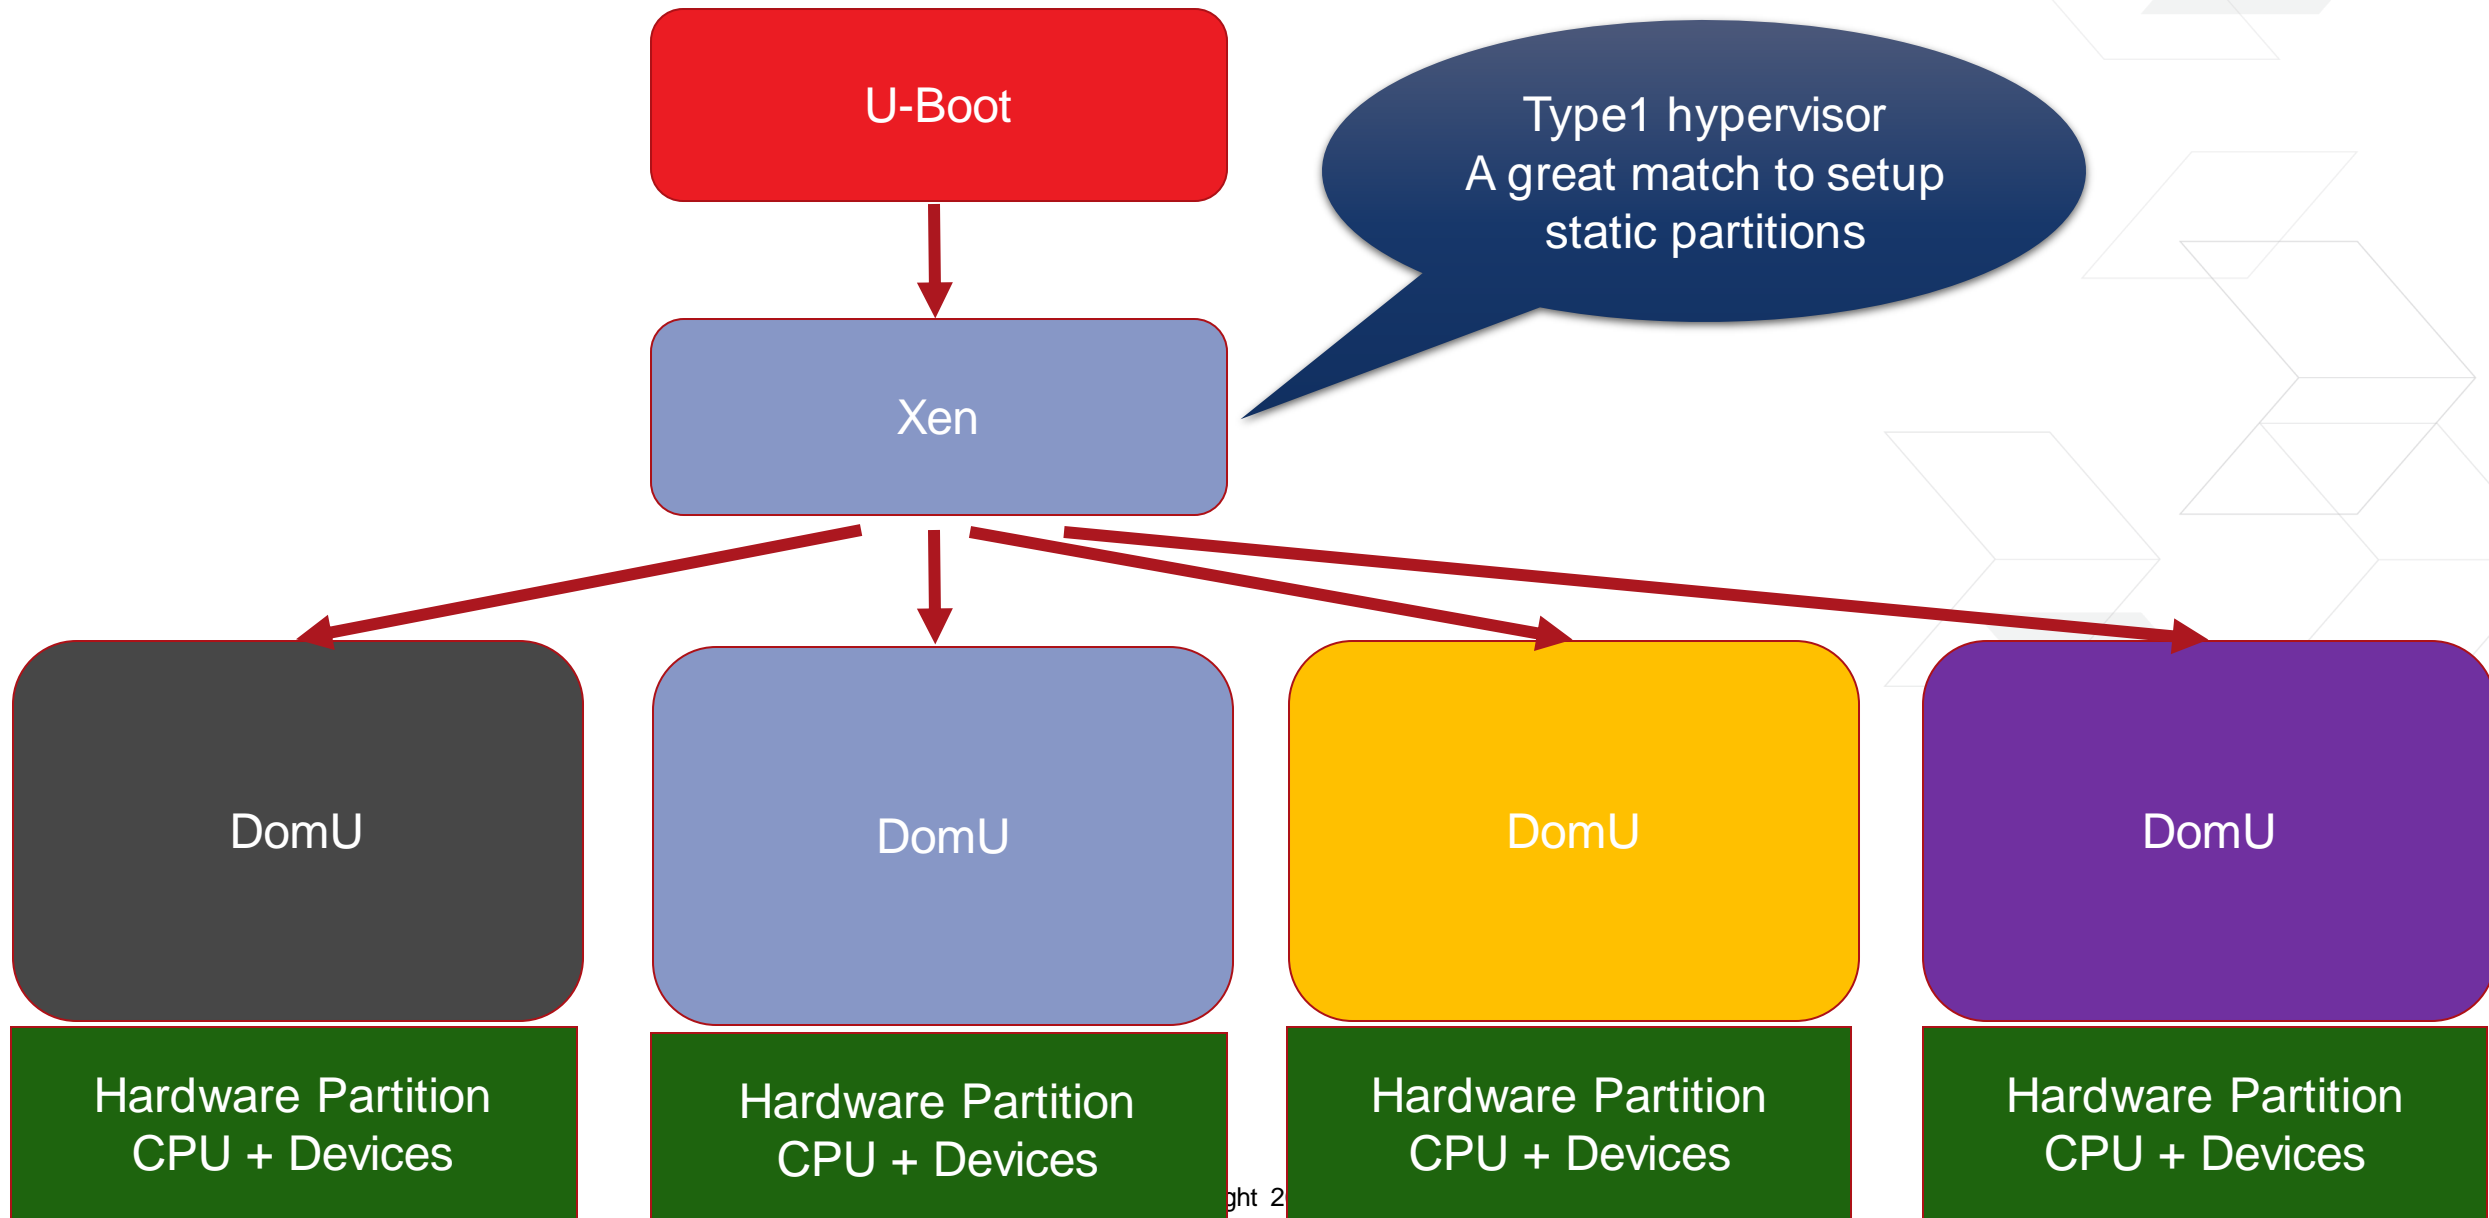
- > Mixed Criticality
- > Safety
- > Security
- > Real-Time
- > Fault isolation
- > Compartmentalization
- > Multiple OSes



Static Partitioning with Xen: the goal



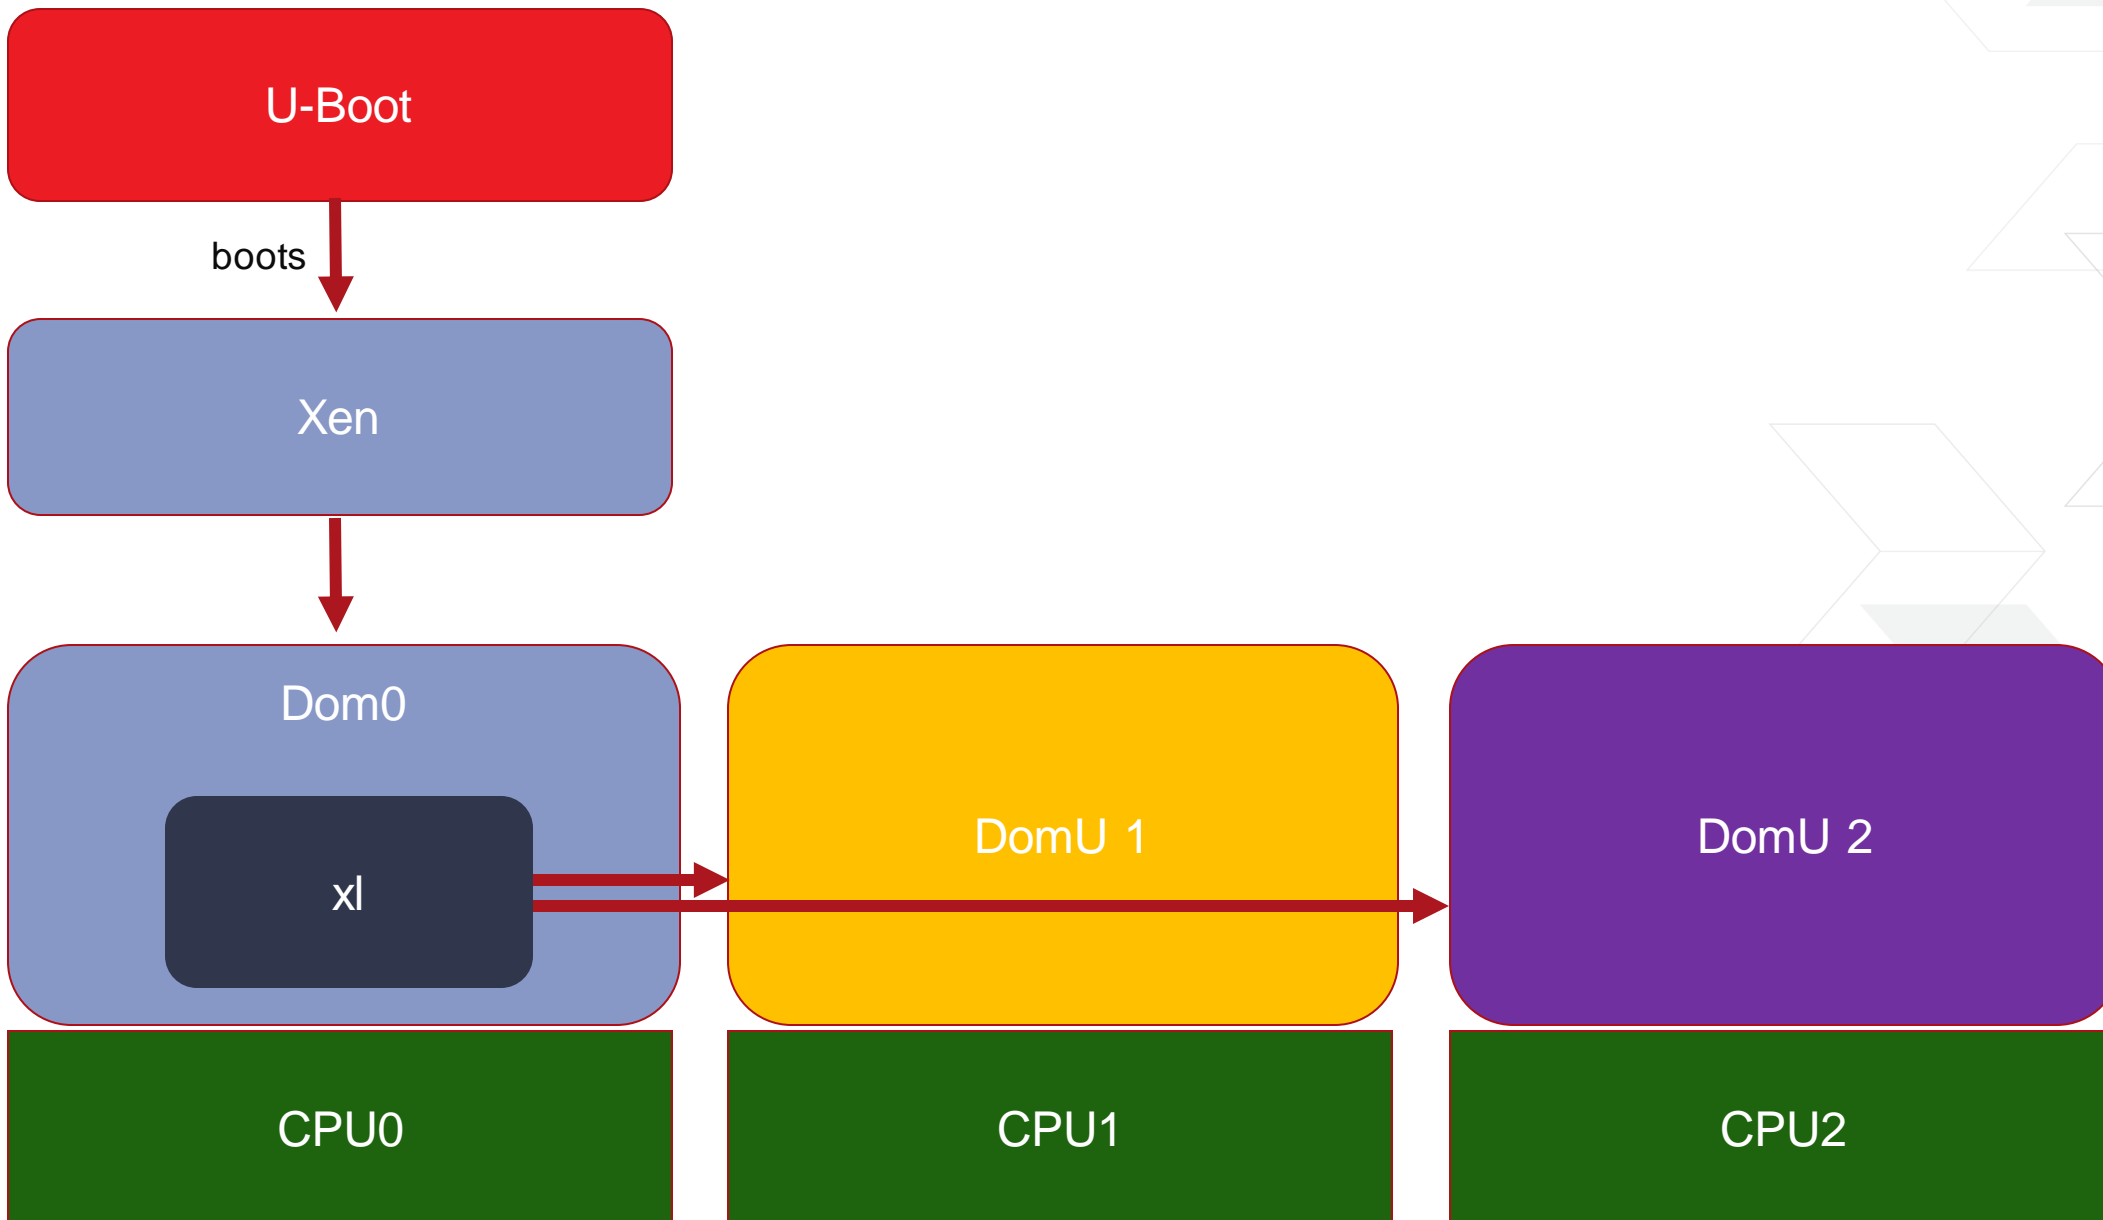
Static Partitioning with Xen: the goal





**REALITY
CHECK
AHEAD**

Traditional Xen System: Configuration and Boot



The problem

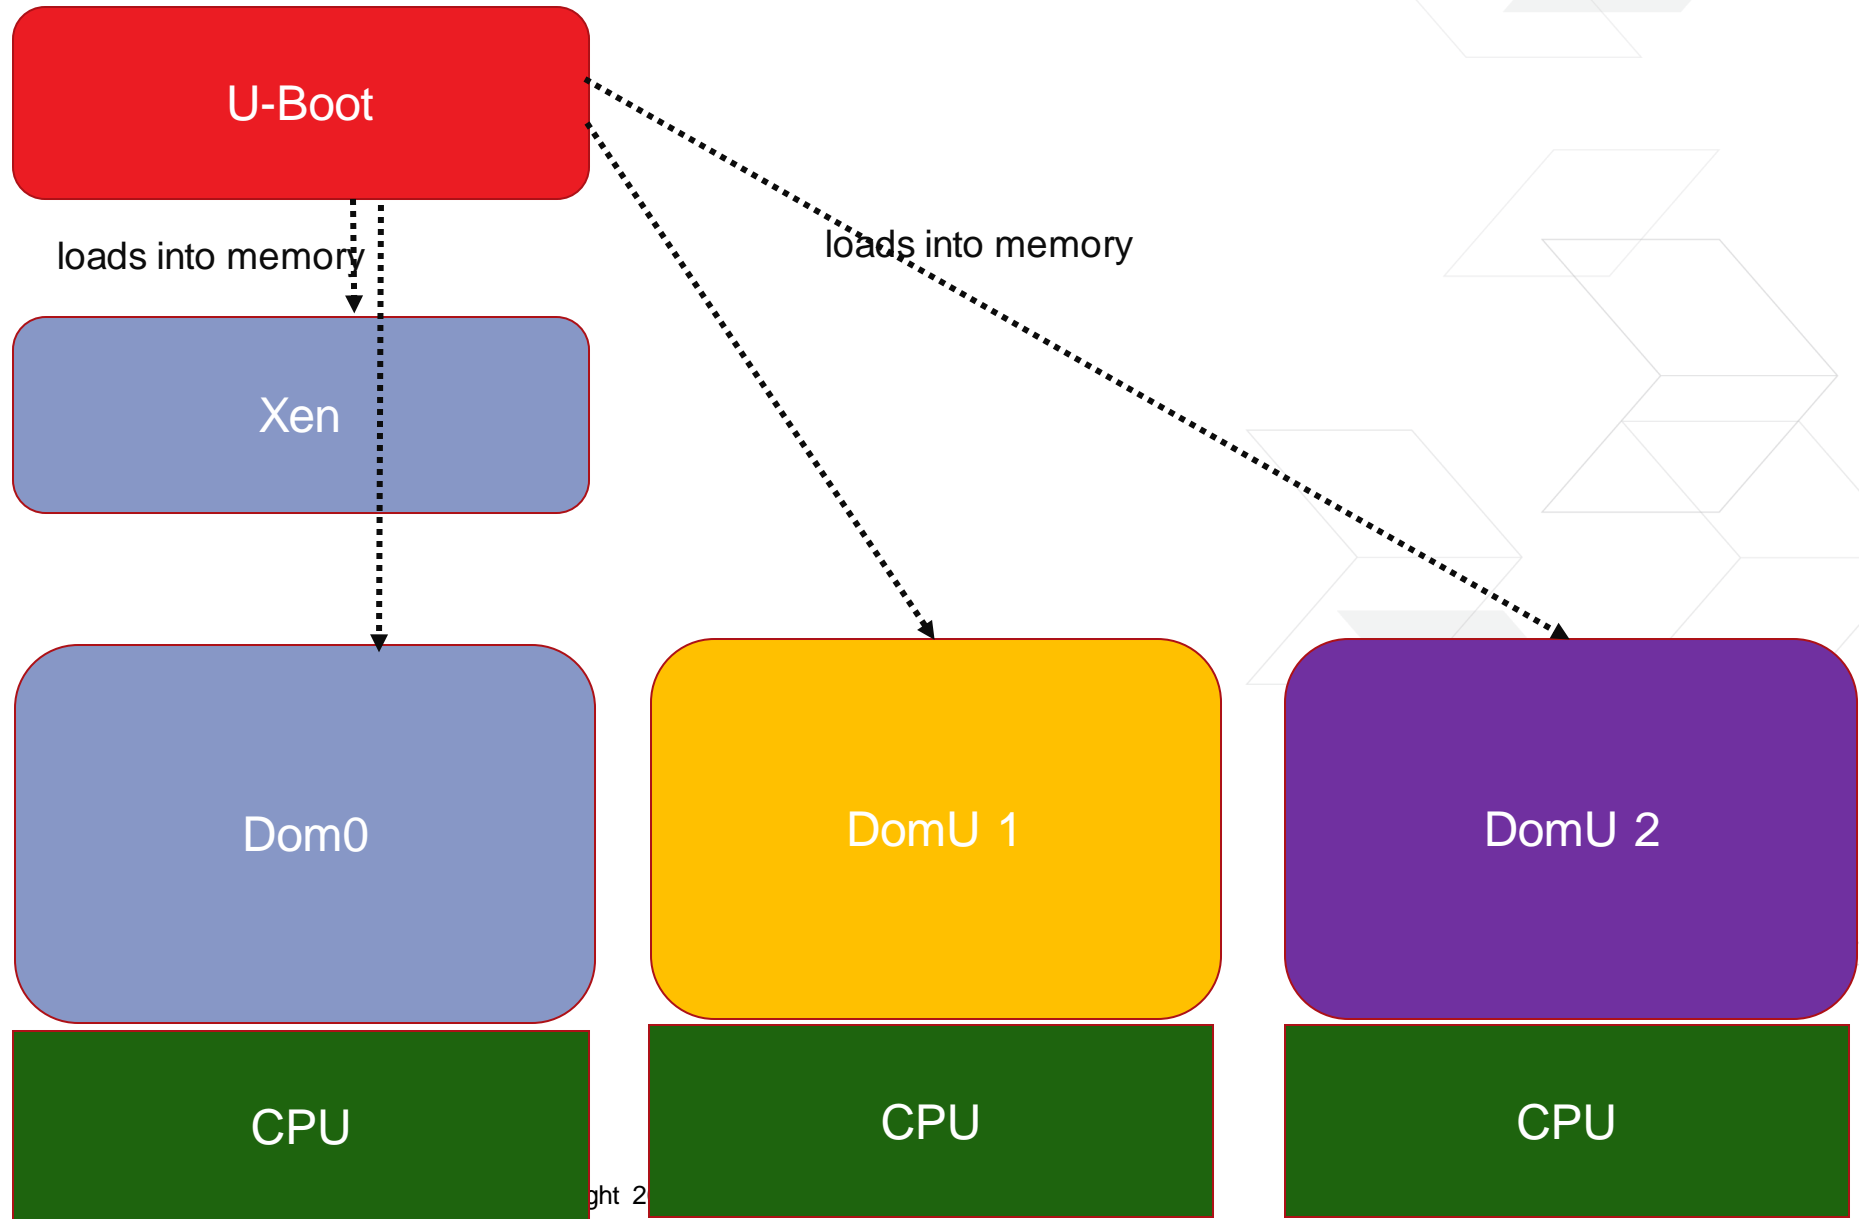
- > Requires Dom0
 - >> Dom0 is privileged
 - >> Dom0 is typically Linux
- > Boot Time
 - >> total > xen + dom0_kernel + dom0_user
- > Safety Certifications
 - >> non-Linux Dom0
 - >> exit Dom0 after boot
- > Complexity
 - >> build-time complexity
 - Yocto rootfs build
 - >> runtime flexibility
 - Monitoring
 - VMs restart



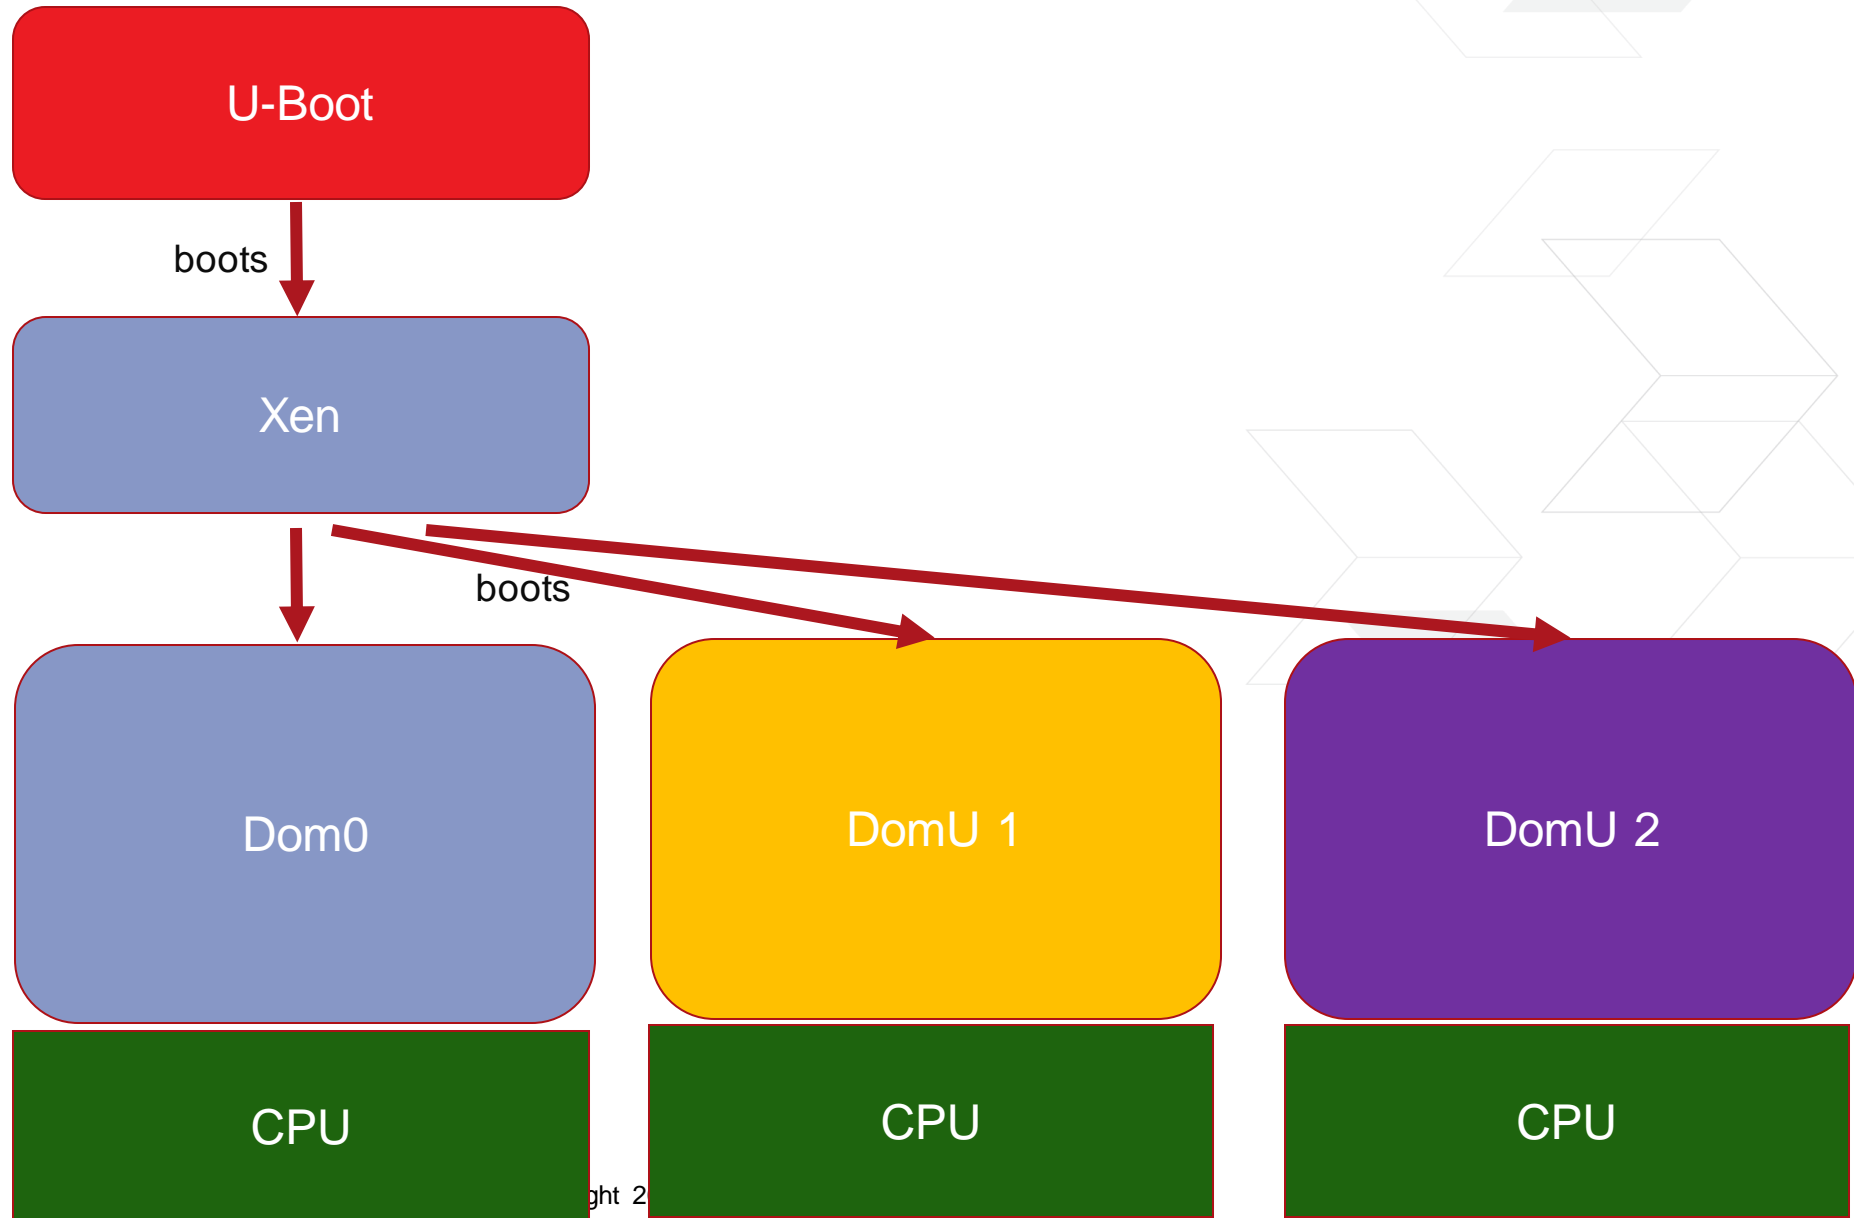
Introducing Dom0-less



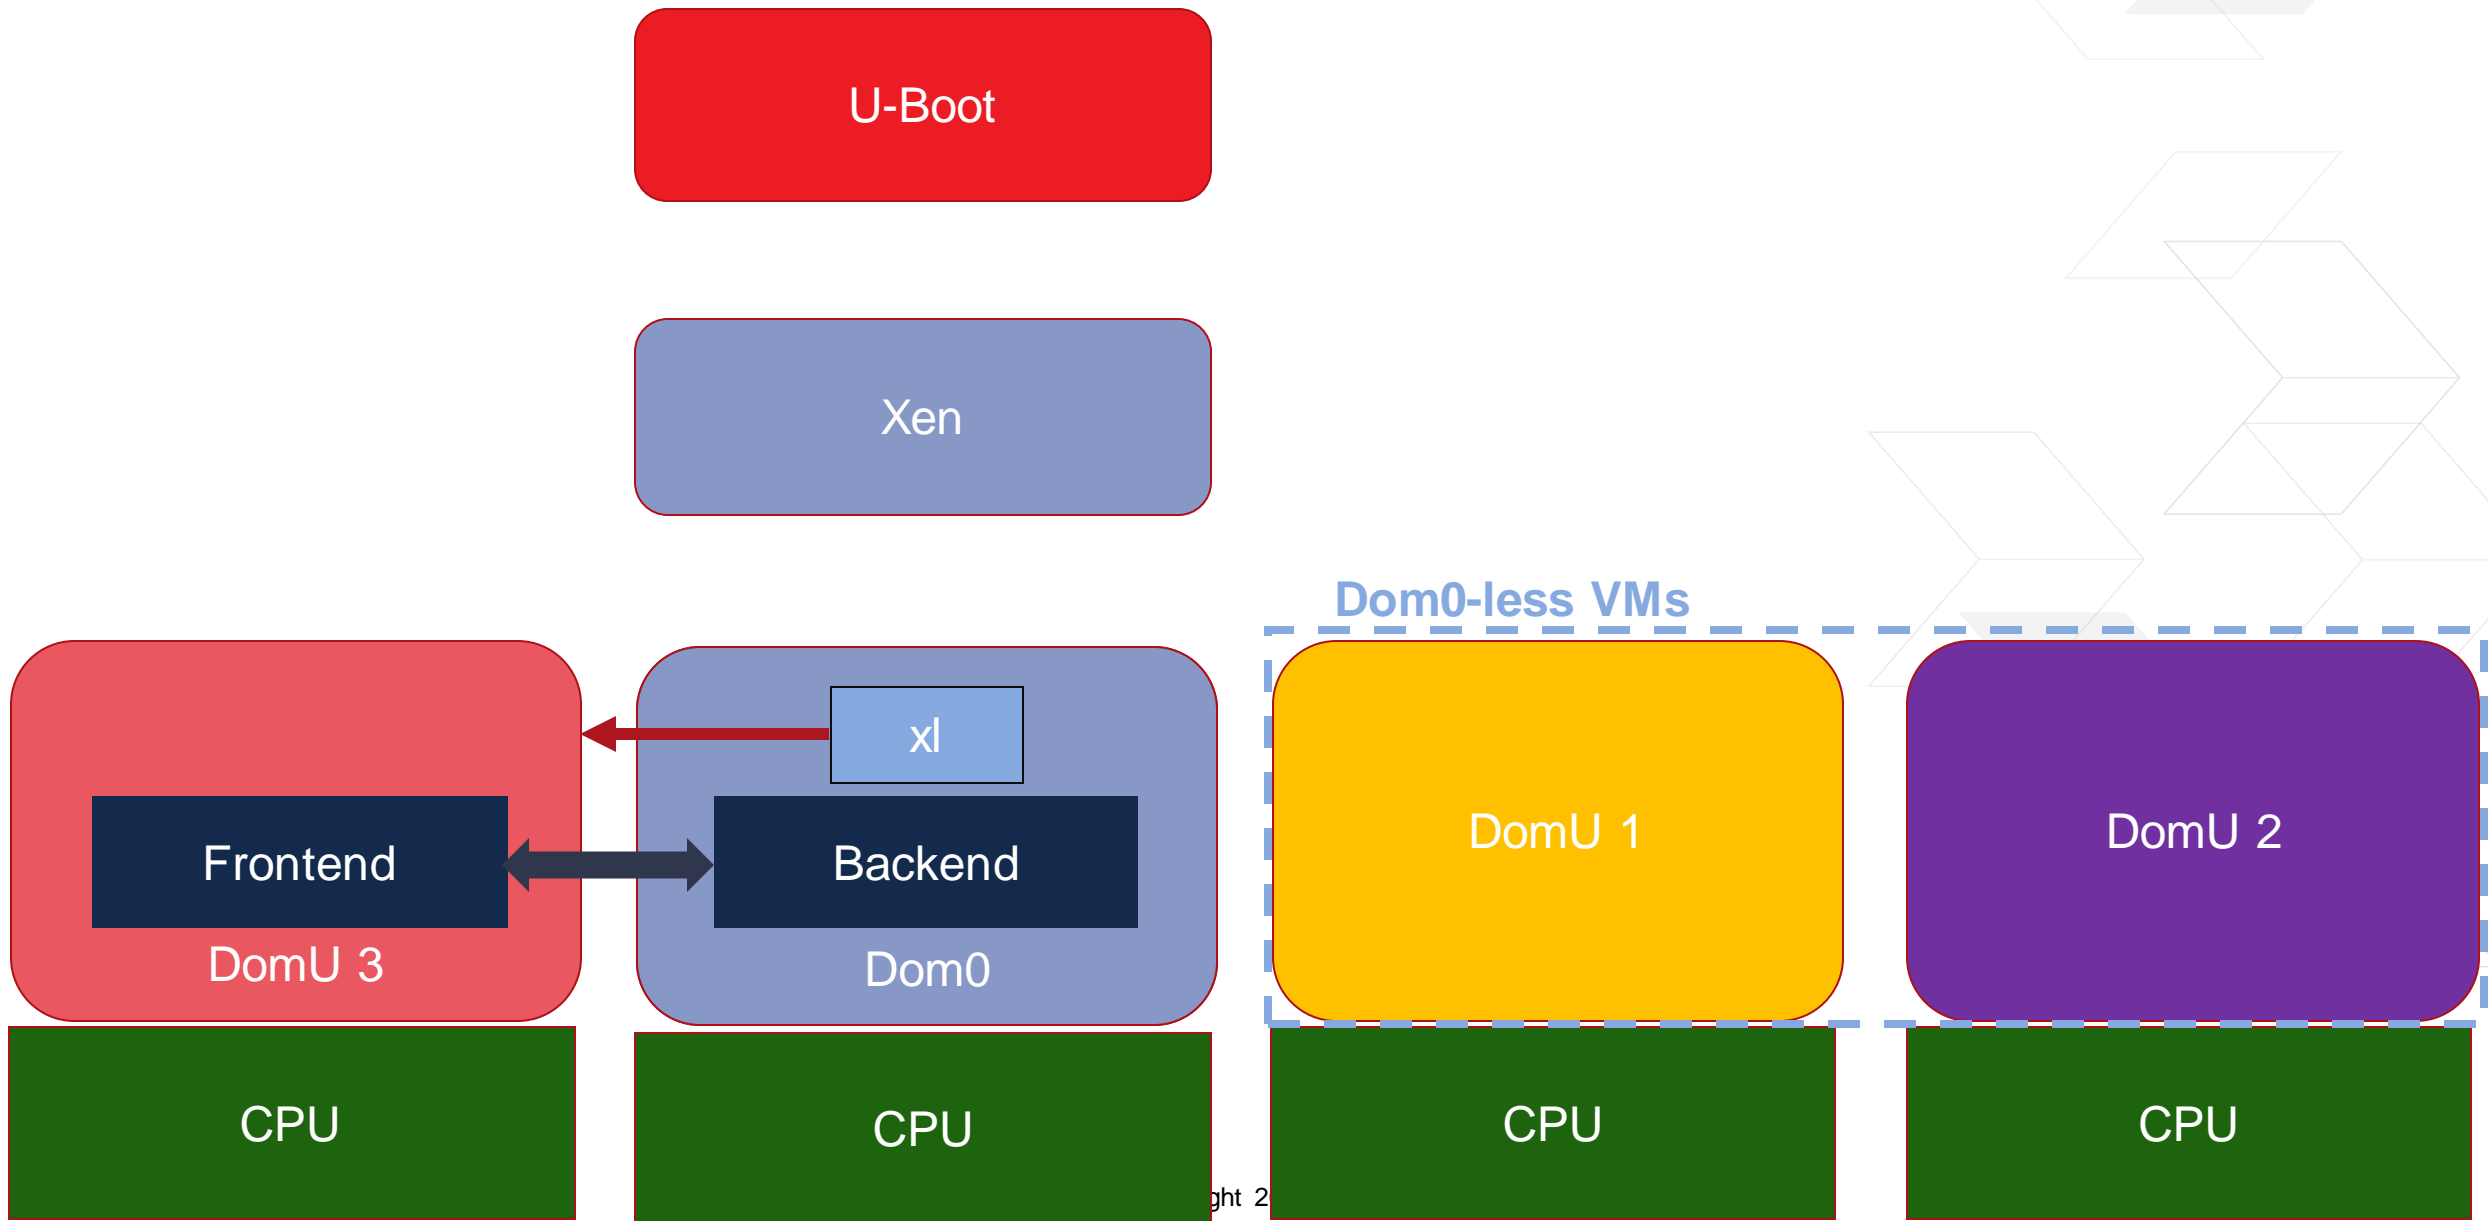
Dom0-less System Configuration and Boot



Dom0-less System Configuration and Boot



Dom0-less System Configuration and Boot



U-Boot + Device Tree protocol

- > Load all the required binaries via U-Boot commands

```
tftpb 0x4000000 dom0less/mpsoc.dtb
tftpb 0x80000 dom0less/Image-dom0
tftpb 0x5000000 dom0less/uXen
tftpb 0xd000000 dom0less/dom0-ramdisk.cpio.uboot

tftpb 0xa000000 dom0less/Image-domU
tftpb 0xb000000 dom0less/domU-ramdisk.cpio

bootm 0x5000000 0xd000000 0x4000000
```

U-Boot + Device Tree protocol

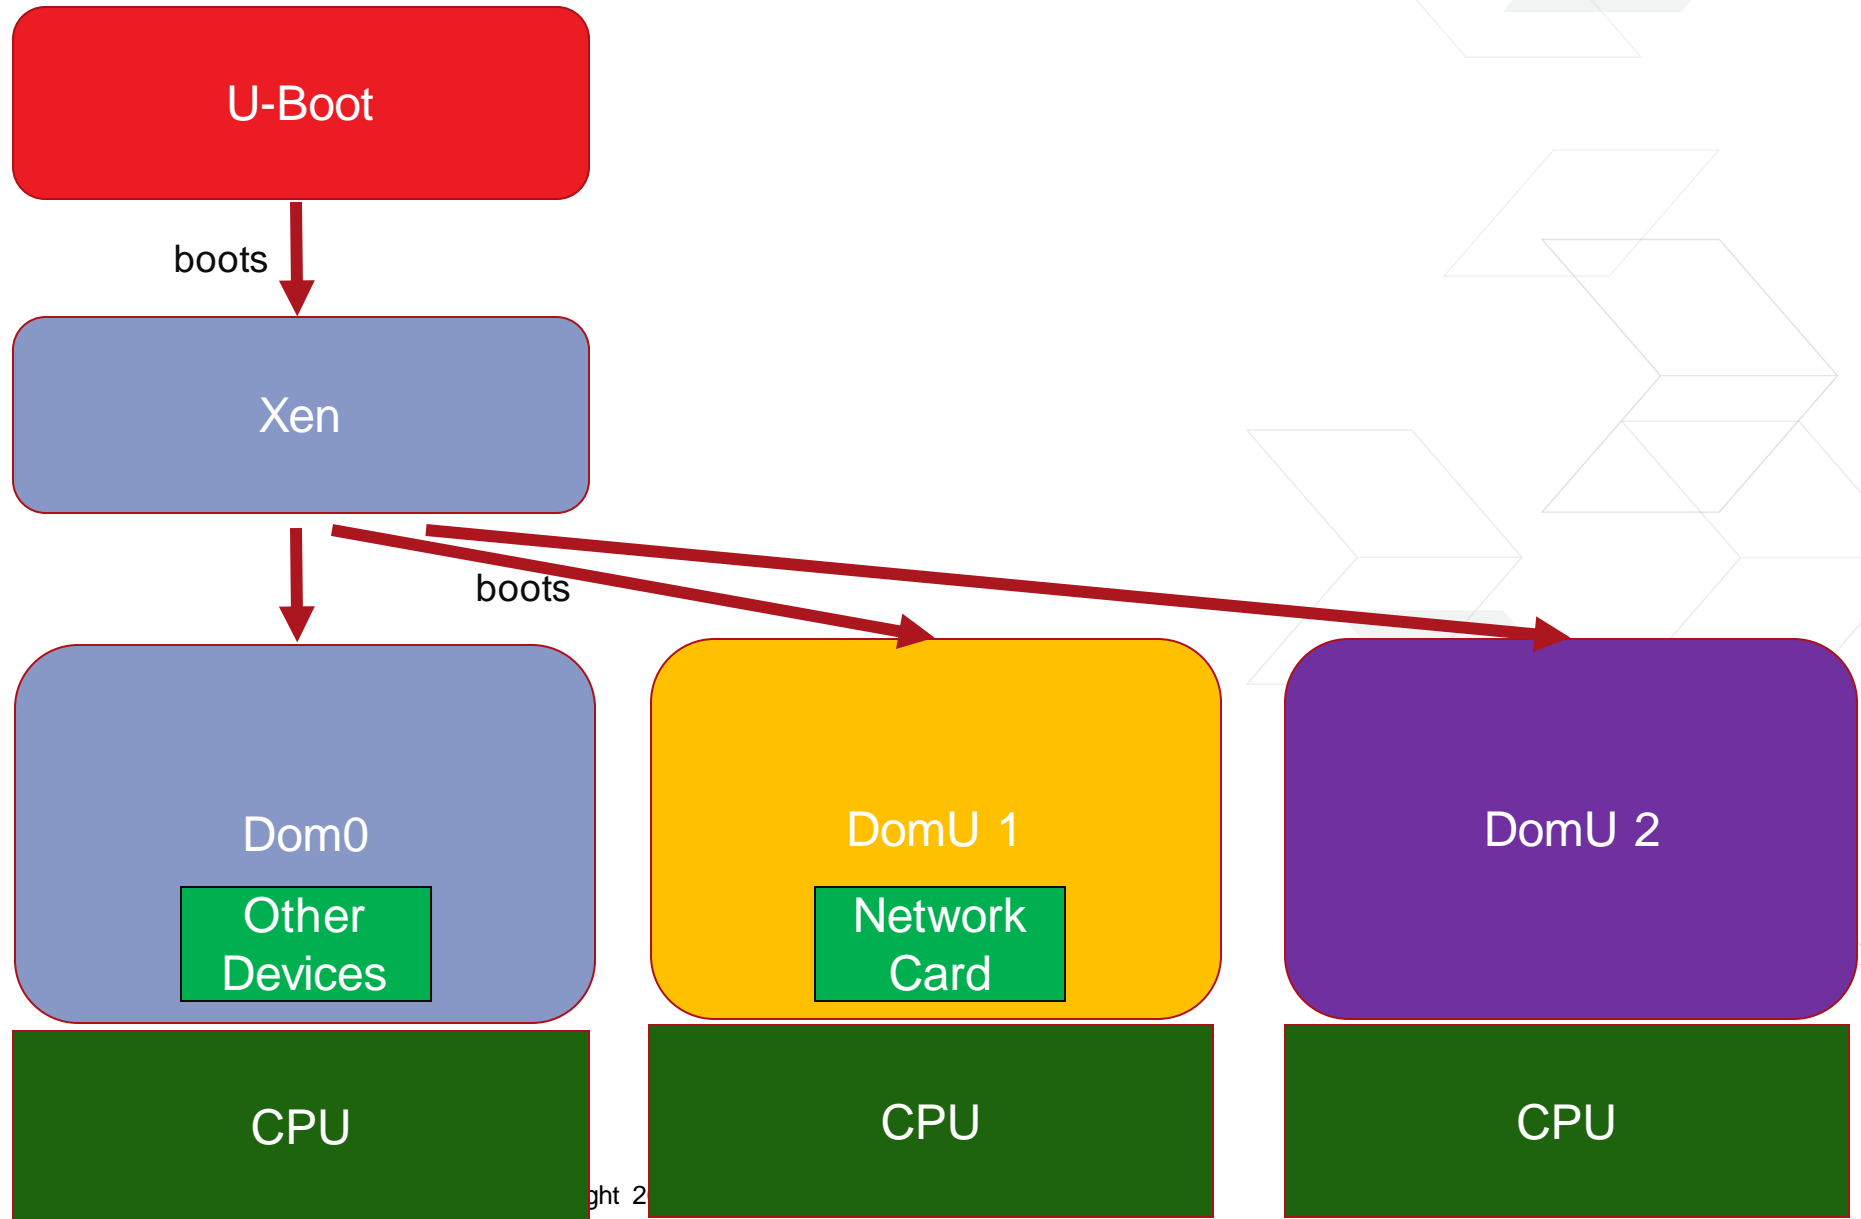
- > Advertise and configure Dom0-less VMs via Device Tree

```
domU1 {
    compatible = "xen,domain";
    memory = <0x0 0x20000>;
    cpus = 1;
    vpl011;

    module@a000000 {
        compatible = "multiboot,kernel", "multiboot,module";
        reg = <0xa000000 0xffffffff>;
        bootargs = "console=ttyAMA0";
    };

    module@b0000000 {
        compatible = "multiboot,ramdisk", "multiboot,module";
        reg = <0xb000000 0xffffffff>;
    };
};
```

Dom0-less Device Assignment



Dom0-less Device Assignment

- > Configured via a nested device tree snippet

```
domU1 {  
    [...]  
  
    module@a000000 {  
        compatible = "multiboot,kernel", "multiboot,module";  
        reg = <0xa000000 0xffffffff>;  
        bootargs = "console=ttyAMA0";  
    };  
  
    module@b0000000 {  
        compatible = "multiboot,ramdisk", "multiboot,module";  
        reg = <0xb000000 0xffffffff>;  
    };  
  
    module@c000000 {  
        compatible = "multiboot,device-tree", "multiboot,module";  
        reg = <0xc000000 0xffffffff>;  
    };  
};
```


Dom0-less Device Assignment

- > Configured via a nested device tree snippet
 - >> the device tree node of the device to assign
 - >> same as for regular DomUs
 - >> Special properties:
 - **xen,path**: path to the device node in the main DT
 - **xen,reg**: memory to remap


```
/dts-v1/;

/ {
    #address-cells = <0x2>;
    #size-cells = <0x1>;

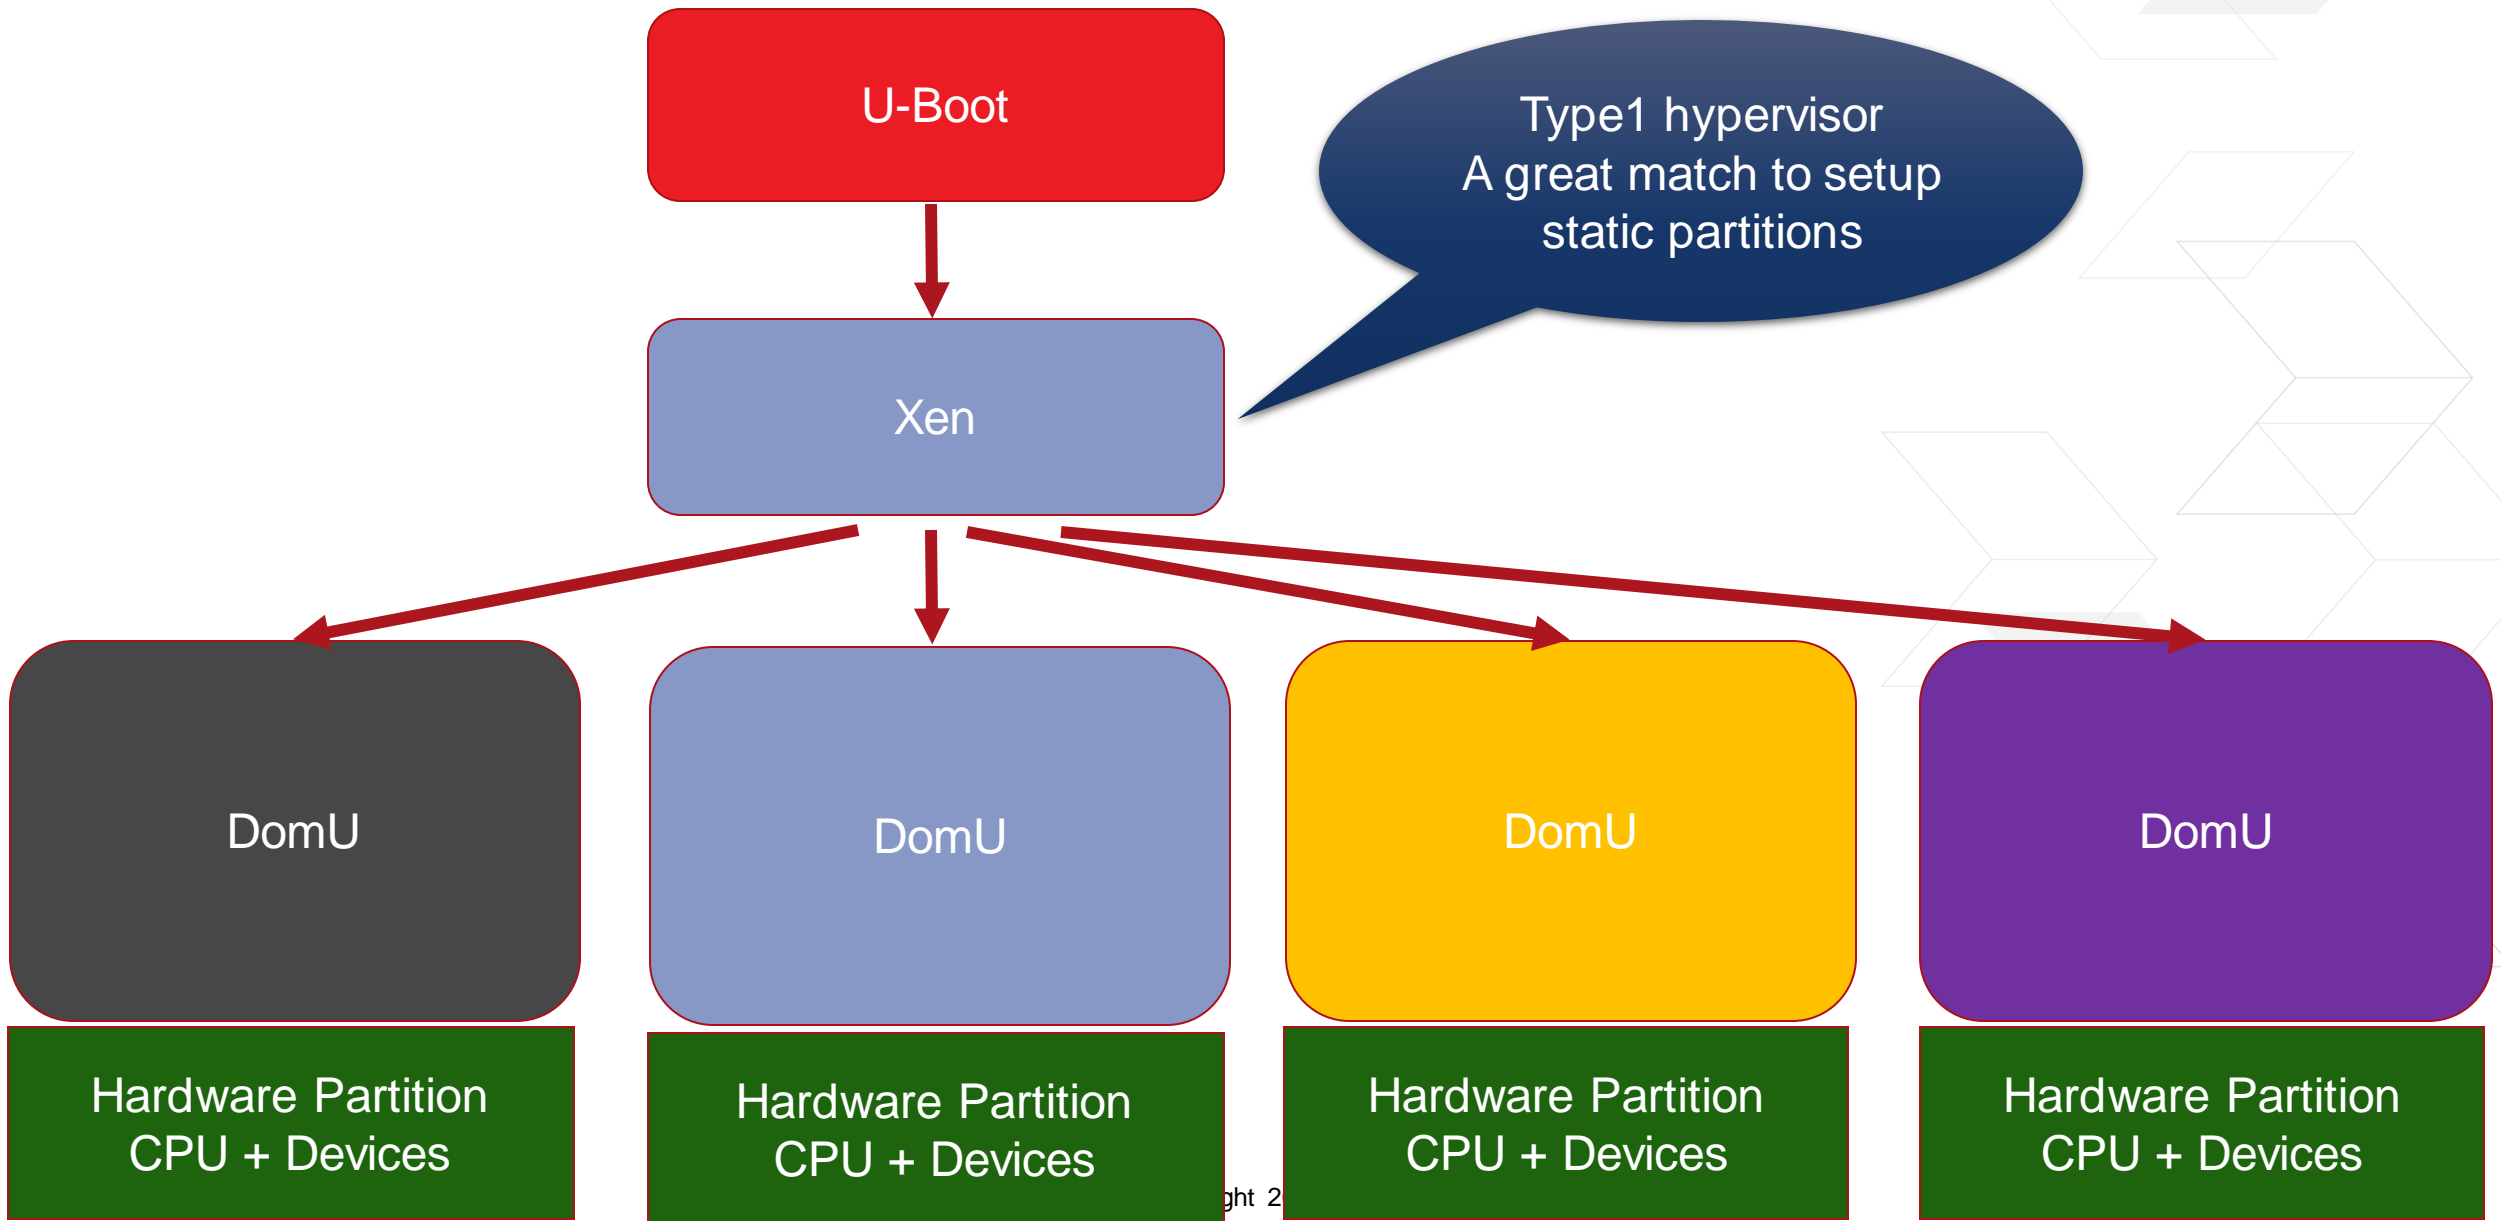
    passthrough {
        compatible = "simple-bus";
        ranges;
        #address-cells = <0x2>;
        #size-cells = <0x1>;

        ethernet@ff0e0000 {
            compatible = "cdns,zynqmp-gem";
            status = "okay";
            xen,path = "/amba/ethernet@ff0e0000";
            xen,reg = <0x0 0xff0e0000 0x1000 0x0 0xff0e0000>;
            reg = <0x0 0xff0e0000 0x1000>;
            clock-names = "pclk", "hclk", "tx_clk", "rx_clk";
            #address-cells = <0x1>;
            #size-cells = <0x0>;
        }
    }
}
```

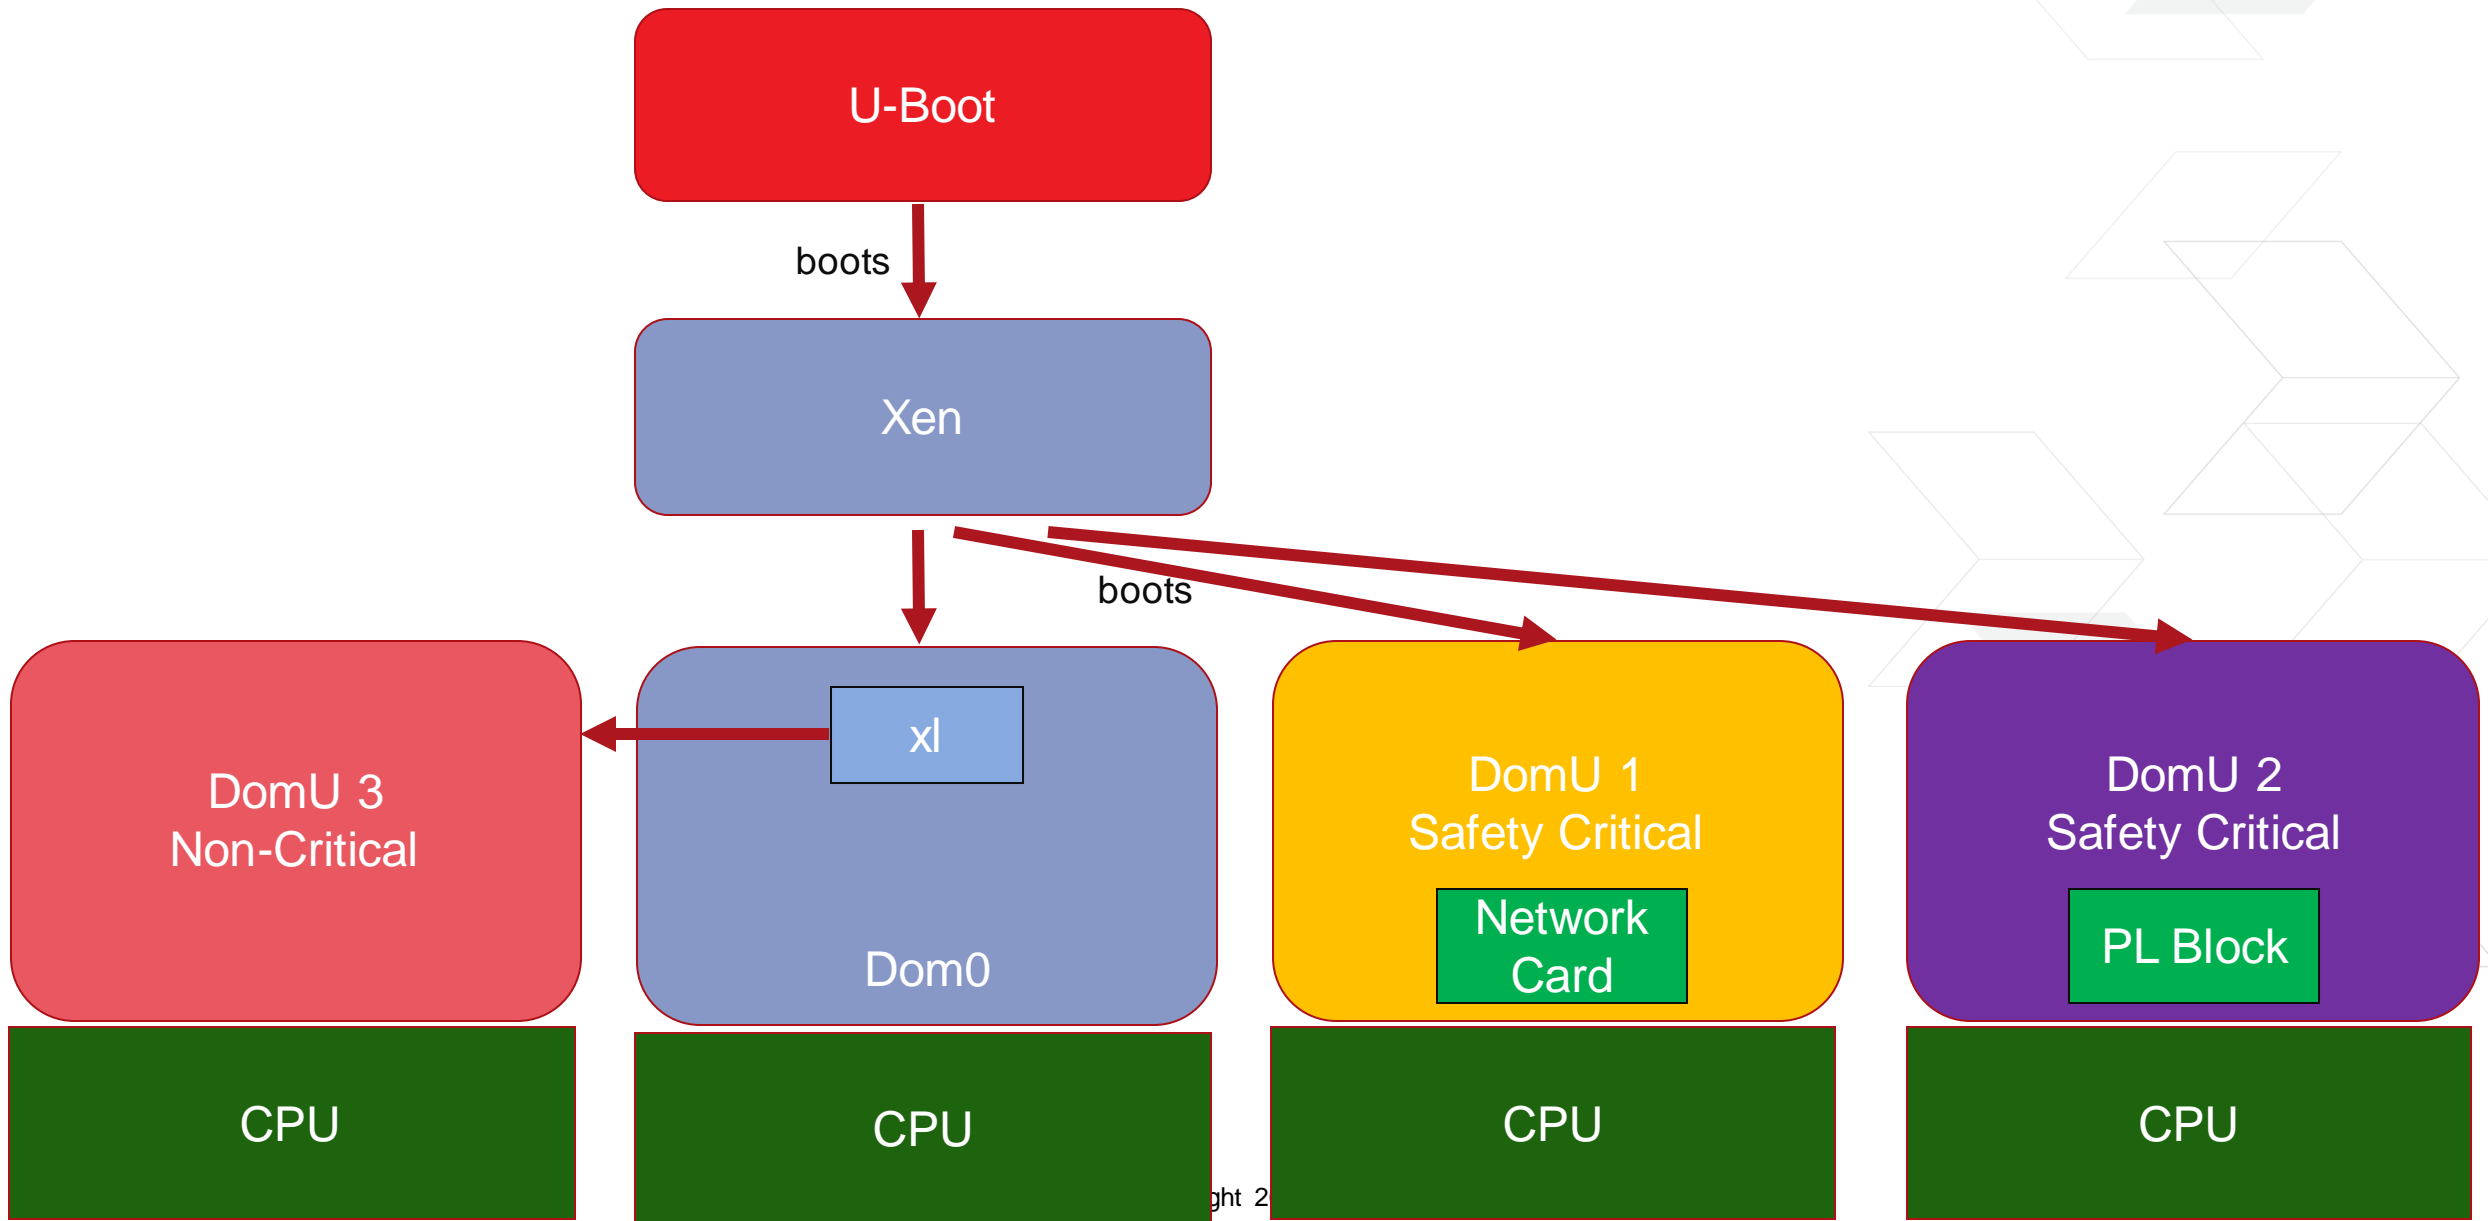
True Dom0-less



Static Partitioning with Xen



More than static partitioning



Dom0-less Pros & Cons

Pros:

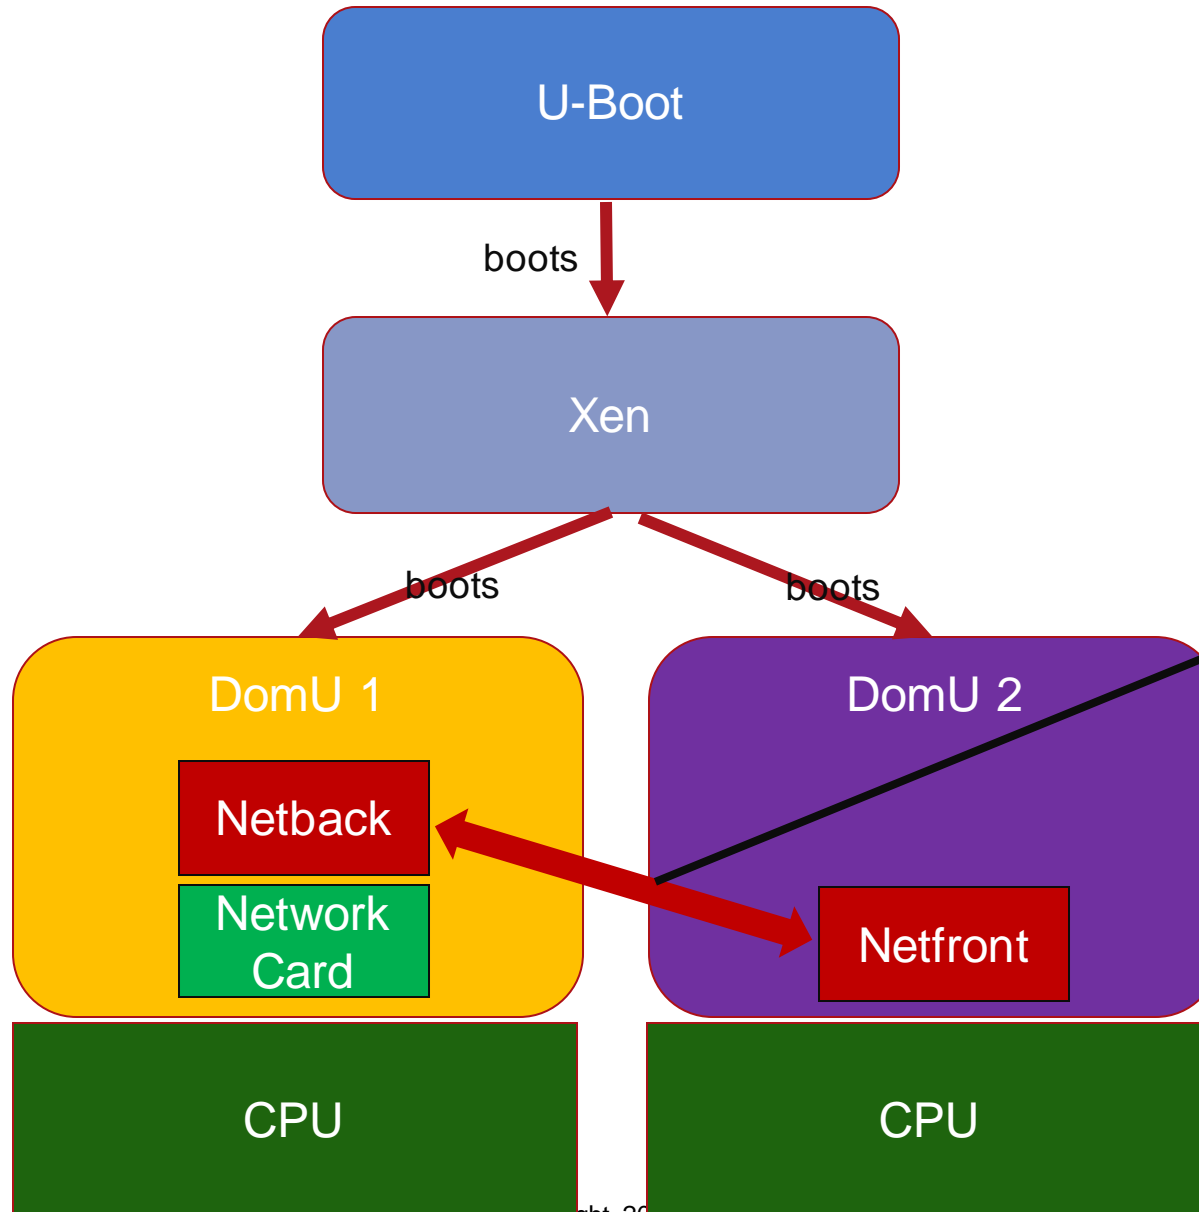
- > Much **faster** startup times
 - >> total \approx xen + domU
- > Enable static partitioning configurations
 - >> Excellent for small systems
 - >> Easier to certify
- > Lower Complexity
 - >> No need for the Xen tools
 - >> Does not require Yocto, just cross-build Xen
 - >> No need for Xen support in Dom0-less VMs, no need for CONFIG_XEN

Cons:

- > No monitoring and restarting DomUs without Dom0
- > No PV frontends/backends without Dom0



Dom0-less and PV drivers



Status & TODO

> DONE

- >> basic Dom0-less booting upstream in Xen 4.12
- >> device assignment implemented and sent to the list (not upstream)

> TODO


- >> True Dom0-less
- >> Shared memory and interrupts for VM-to-VM communications
- >> PV frontends/backends drivers for Dom0-less VMs



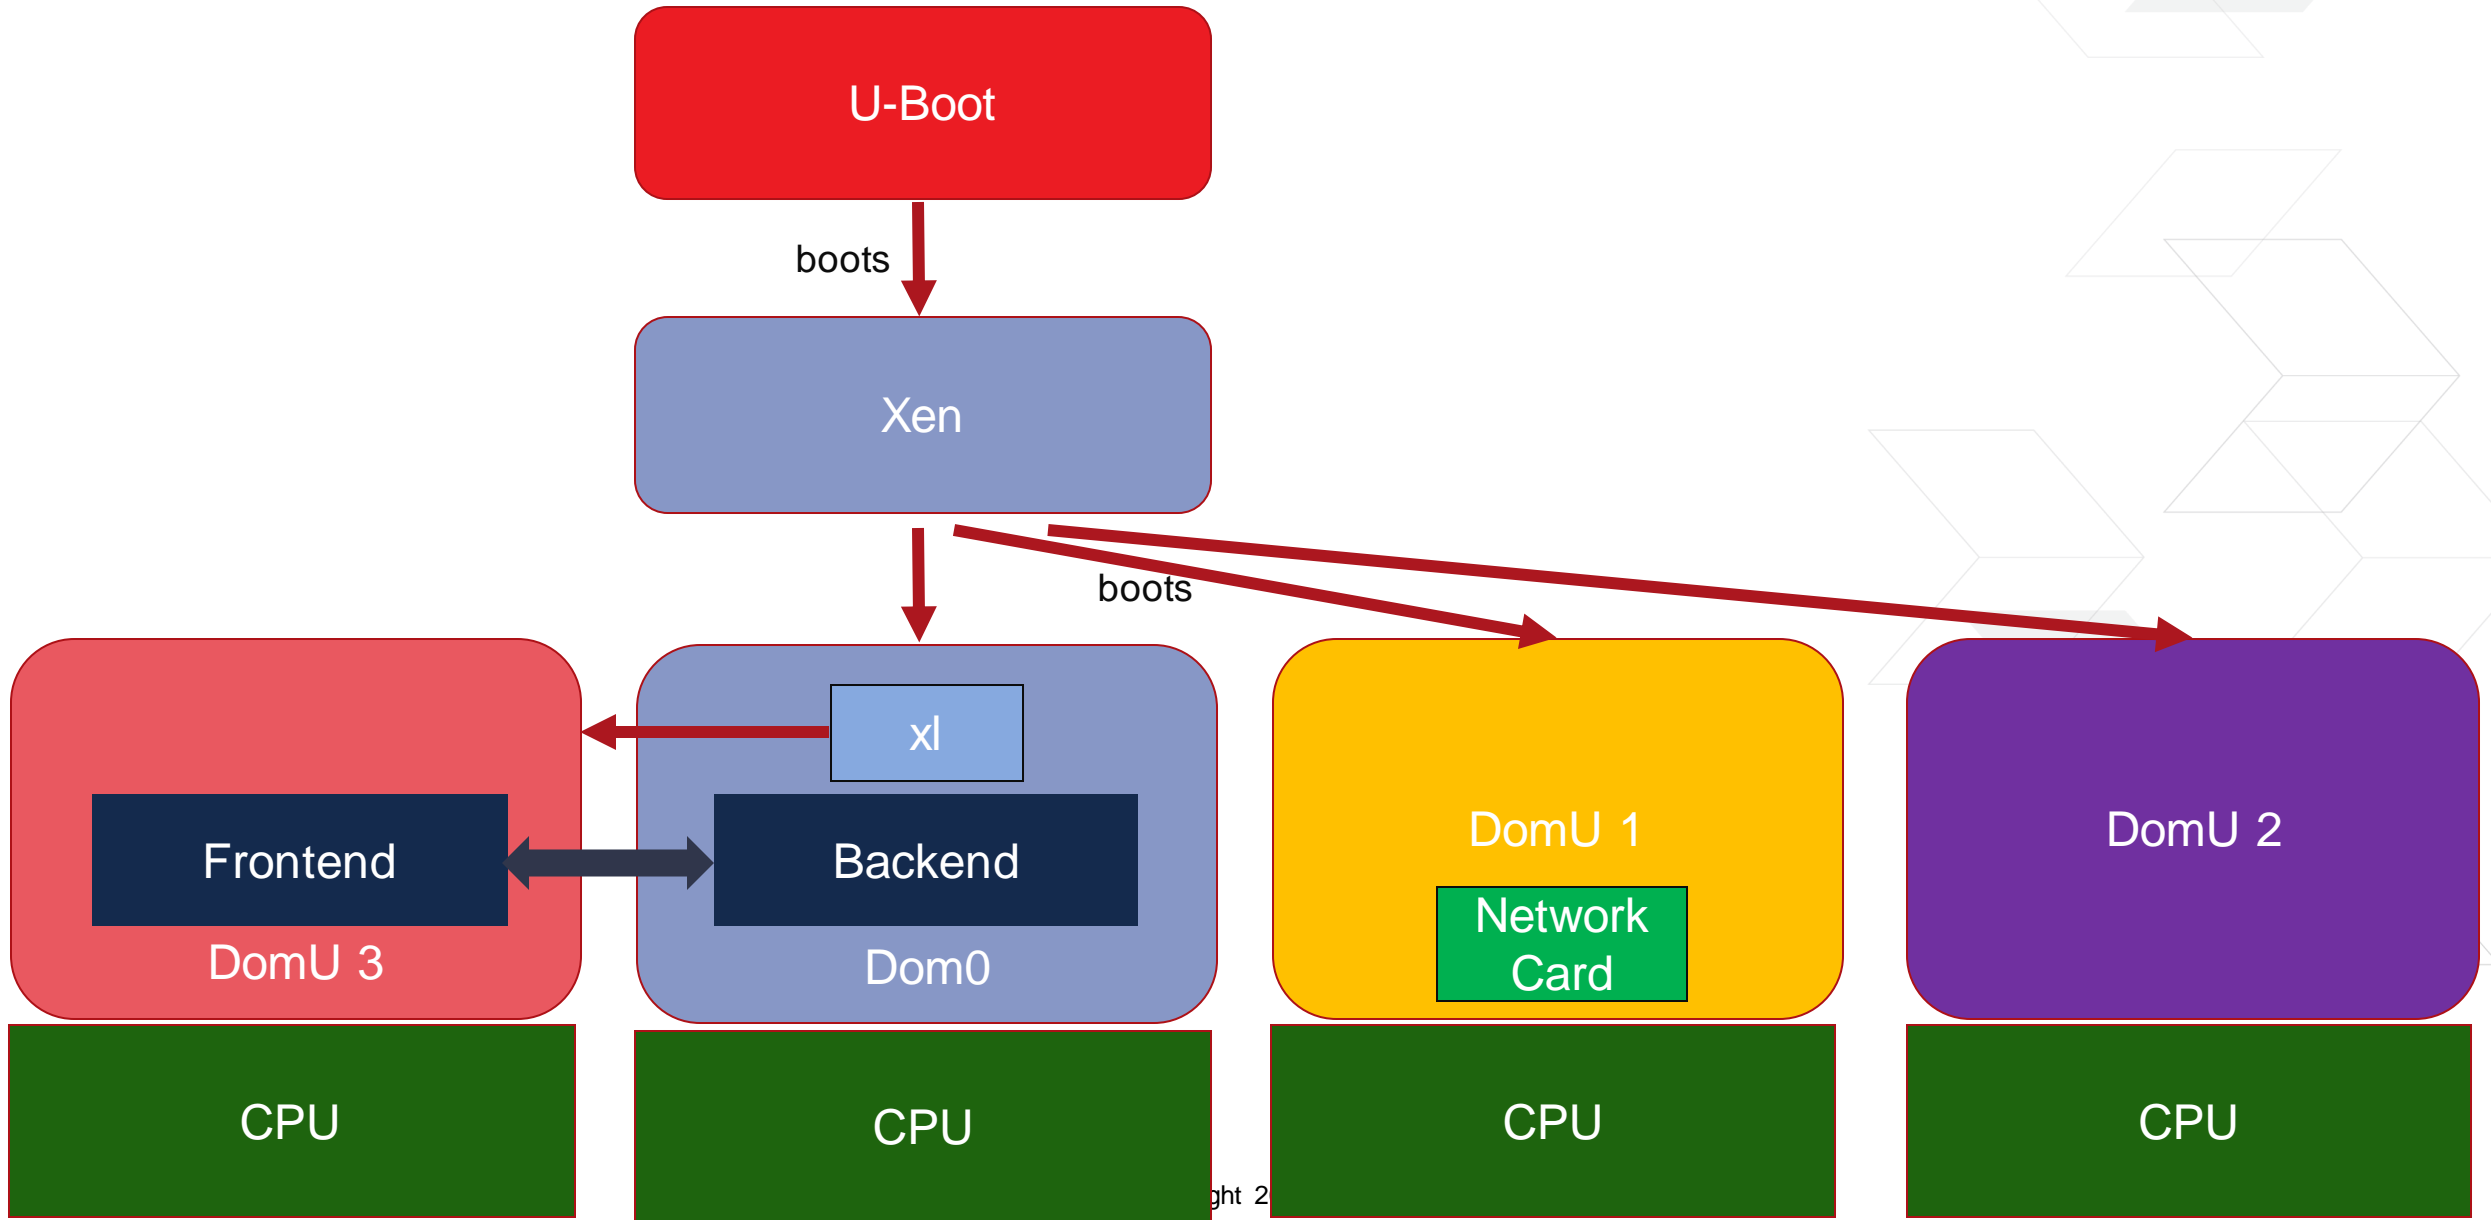
Demos




True Dom0-less



Dom0-less and PVCalls



Dom0-less and PVCalls



Questions?



Adaptable.
Intelligent.



Slides Parking Lot



Traditional Xen System Configuration and Boot

