



LTD20-200

Heterogeneous Hardware-Accelerated OCI
Containers & Secure OCI Containers

Wesley Skeffington
Stefano Stabellini



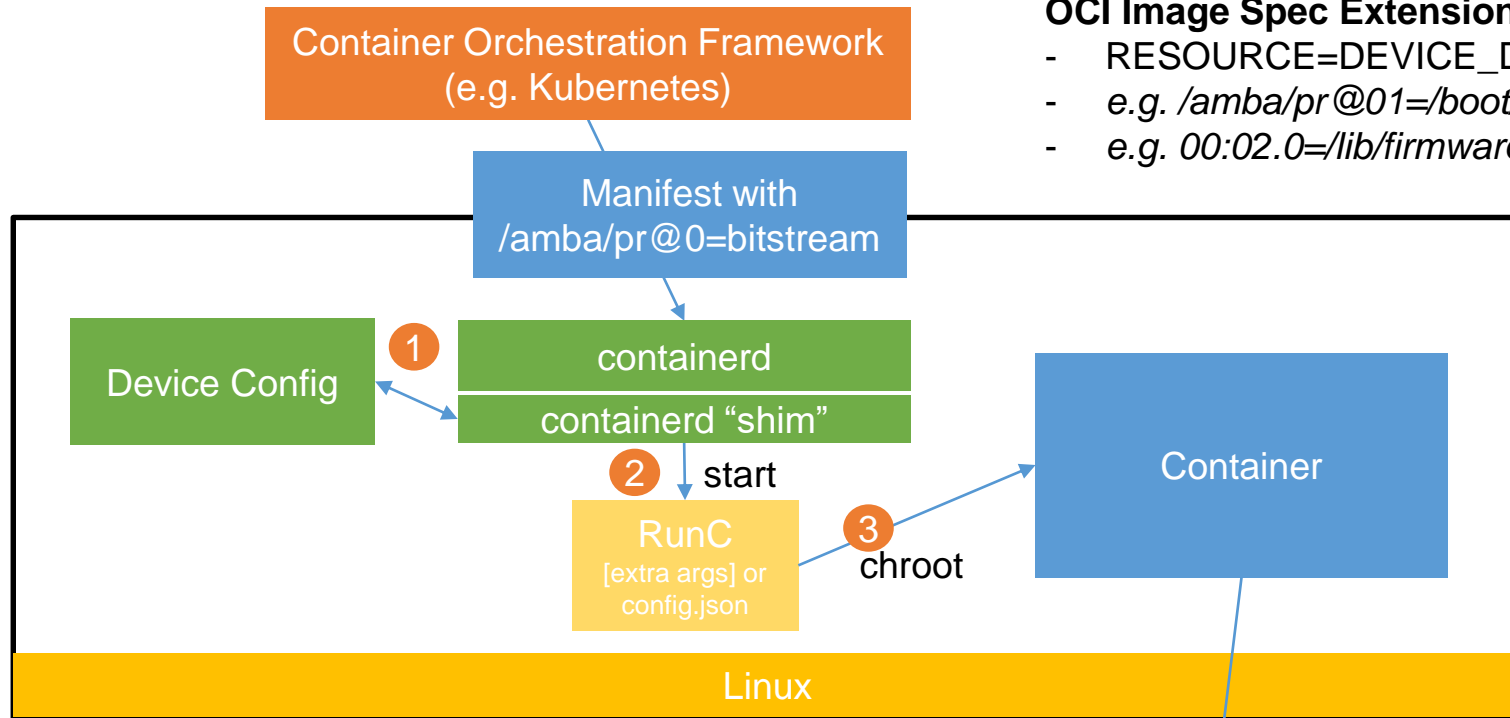
Heterogenous HW Aware Containers

- **Goal:** Extend OCI specifications to support the utilization of more heterogeneous hardware capabilities within the containers ecosystem
 - Heterogeneous accelerator (FPGA, GPU, ...) deployment & access
 - Support for Cortex-R and Cortex-M co-processors
 - Kernel deployment (Xen hypervisor implementation as reference)
- **Motivation:** OCI is a quickly growing application packaging & deployment framework in cloud and edge targets. Simultaneously application performance growth is ever more dependent on mapping to heterogeneous platform targets. Therefore compute industry has a necessity to bridge this gap between container software ecosystem and increasingly domain-specific hardware architectures.

RunC + Heterogenous HW Access – “Startup”

OCI Image Spec Extensions:

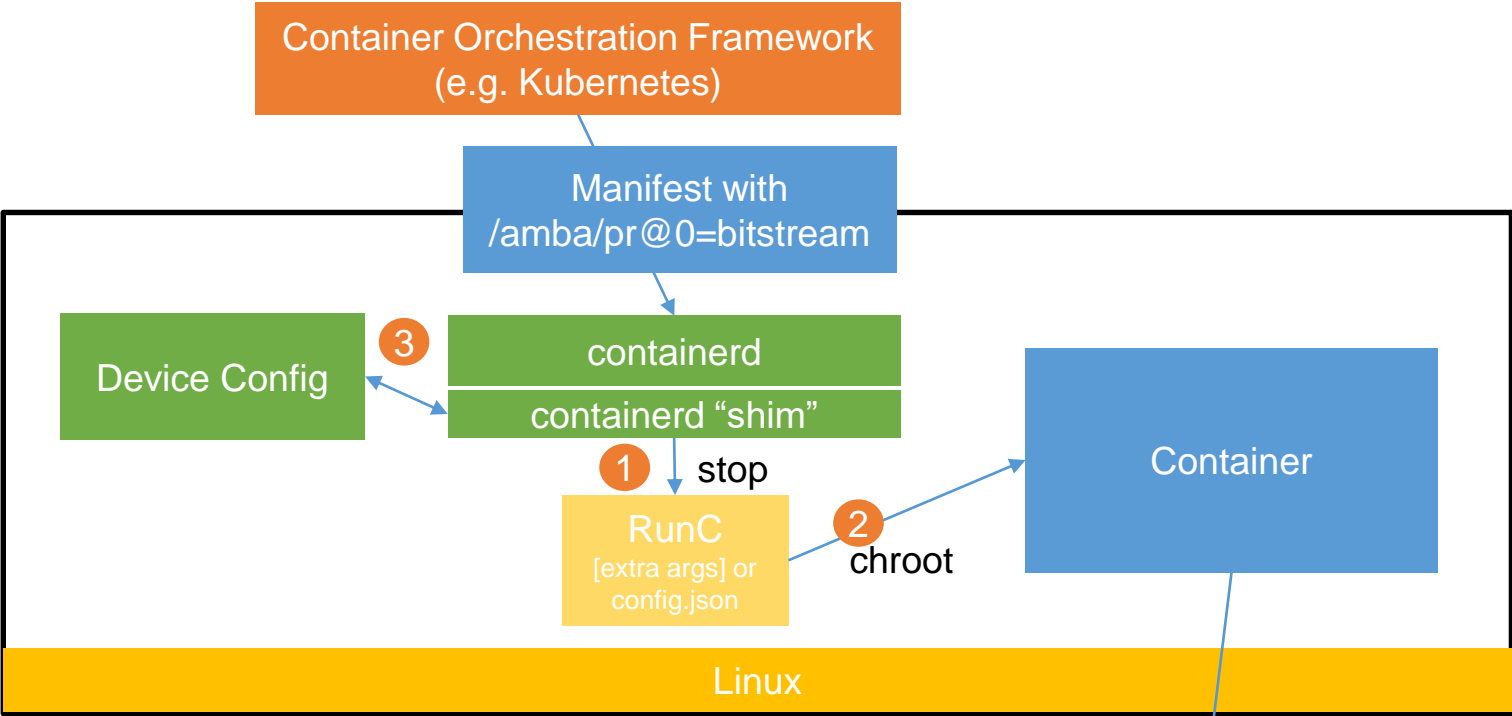
- RESOURCE=DEVICE_DATA
- e.g. /amba/pr@01=/boot/bitstream
- e.g. 00:02.0=/lib/firmware/ipu3-fw.bin



1

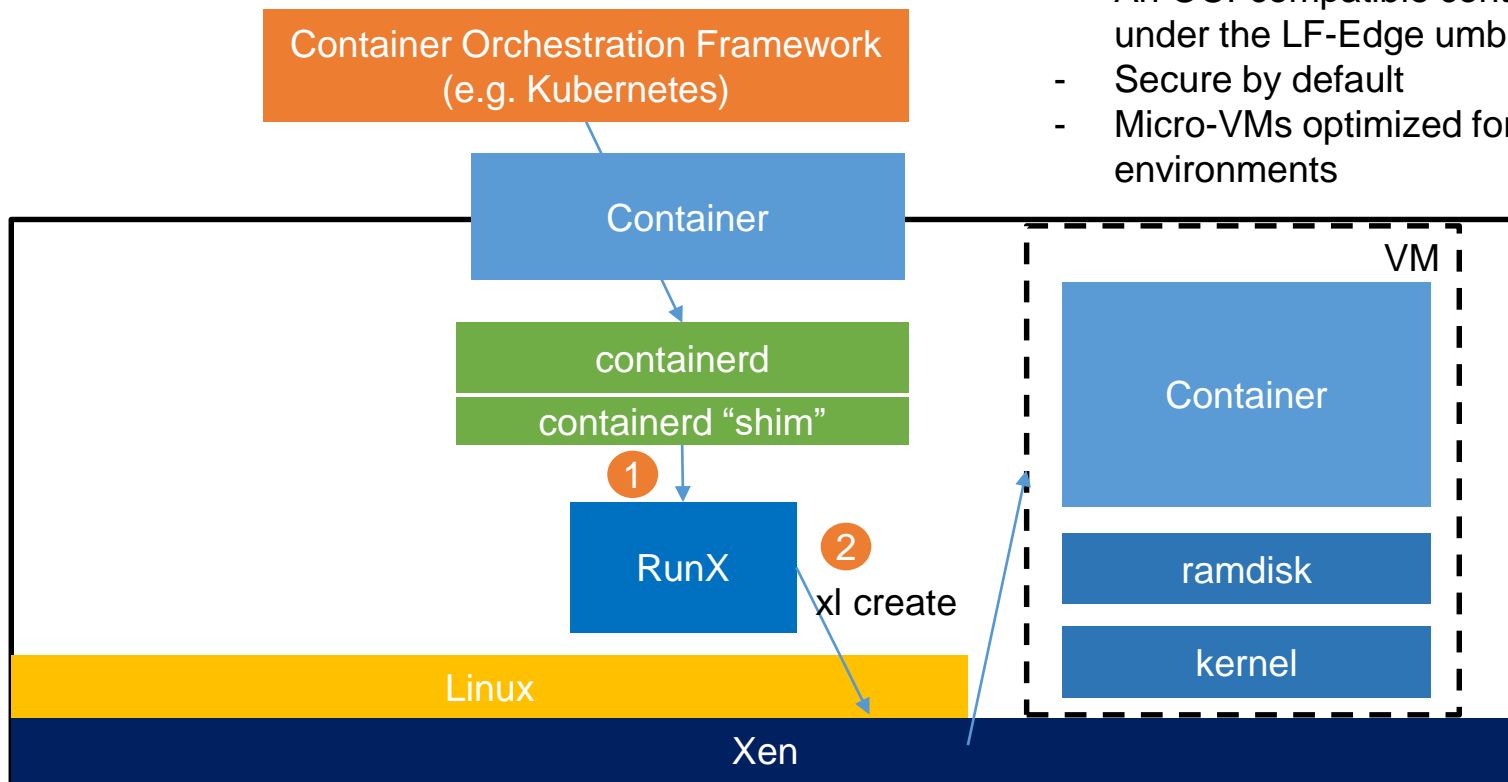
Container pre-start hook
HWConfig(resource, bitstream)
-> dev name & info

RunC + Heterogenous HW Access – “Shutdown”



3
Container post-stop hook
HWConfig(resource)
-> Free resources

Secure Containers w/ Xen



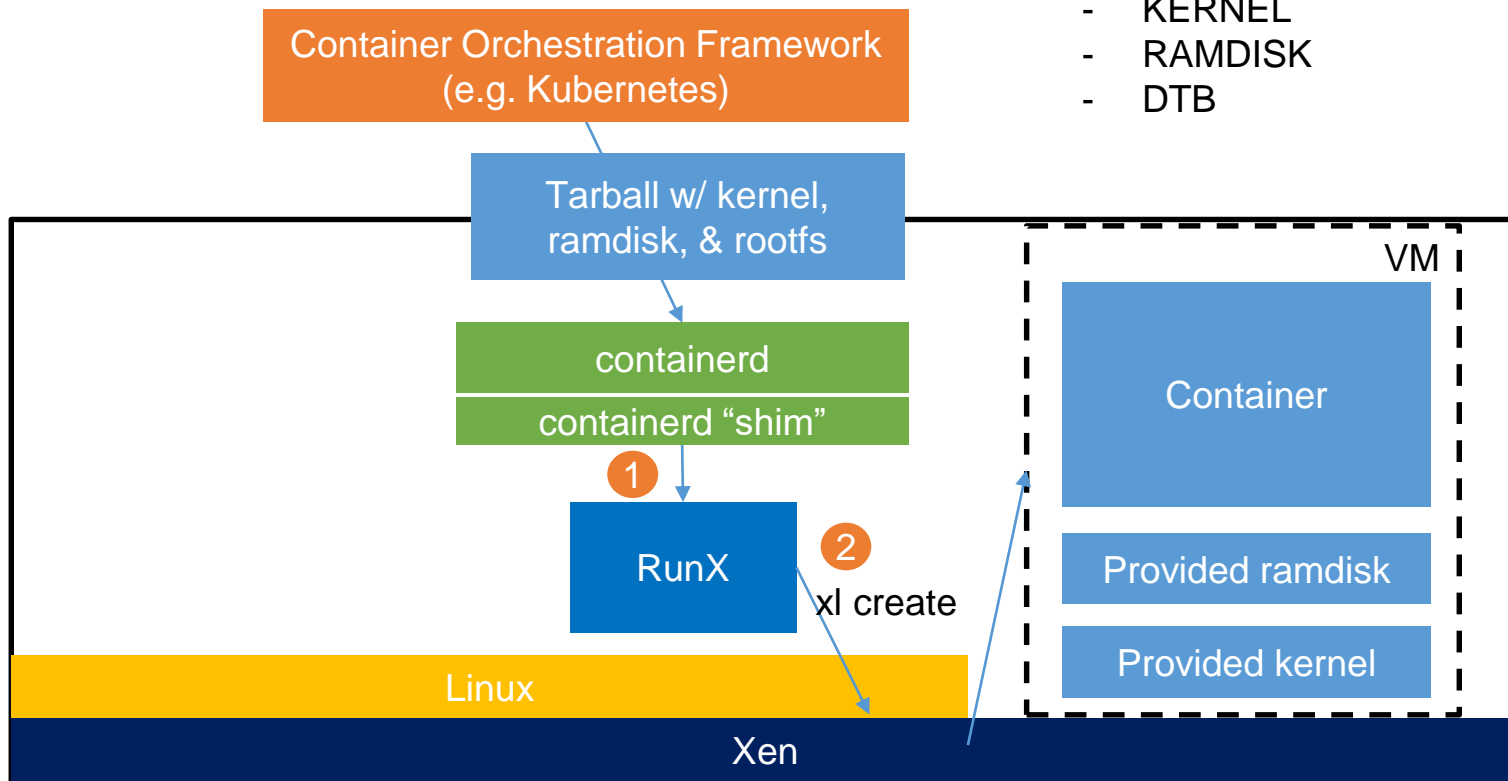
RunX:

- An OCI-compatible containers runtime under the LF-Edge umbrella
- Secure by default
- Micro-VMs optimized for constrained environments

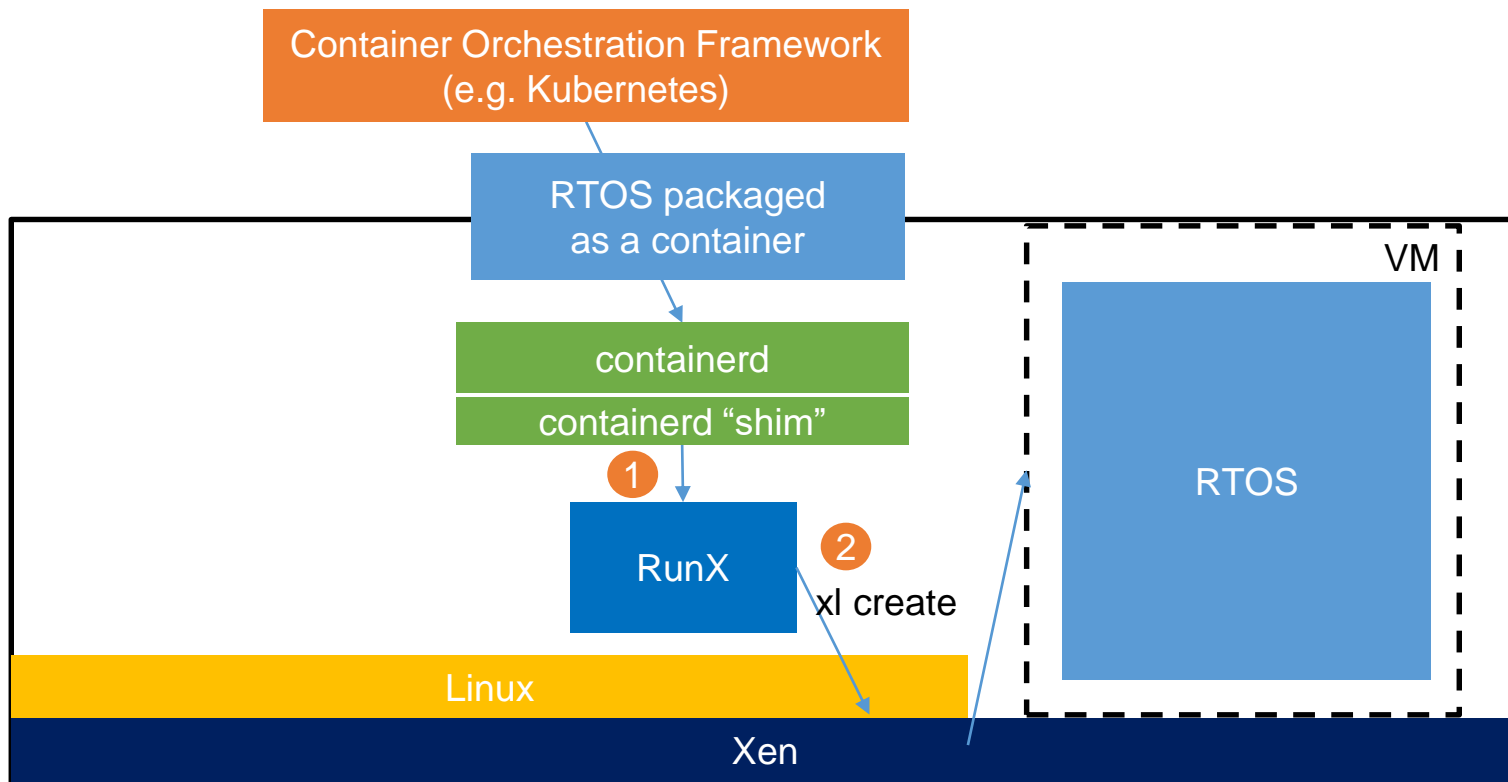
Secure Containers w/ Xen

OCI Image Spec Extensions:

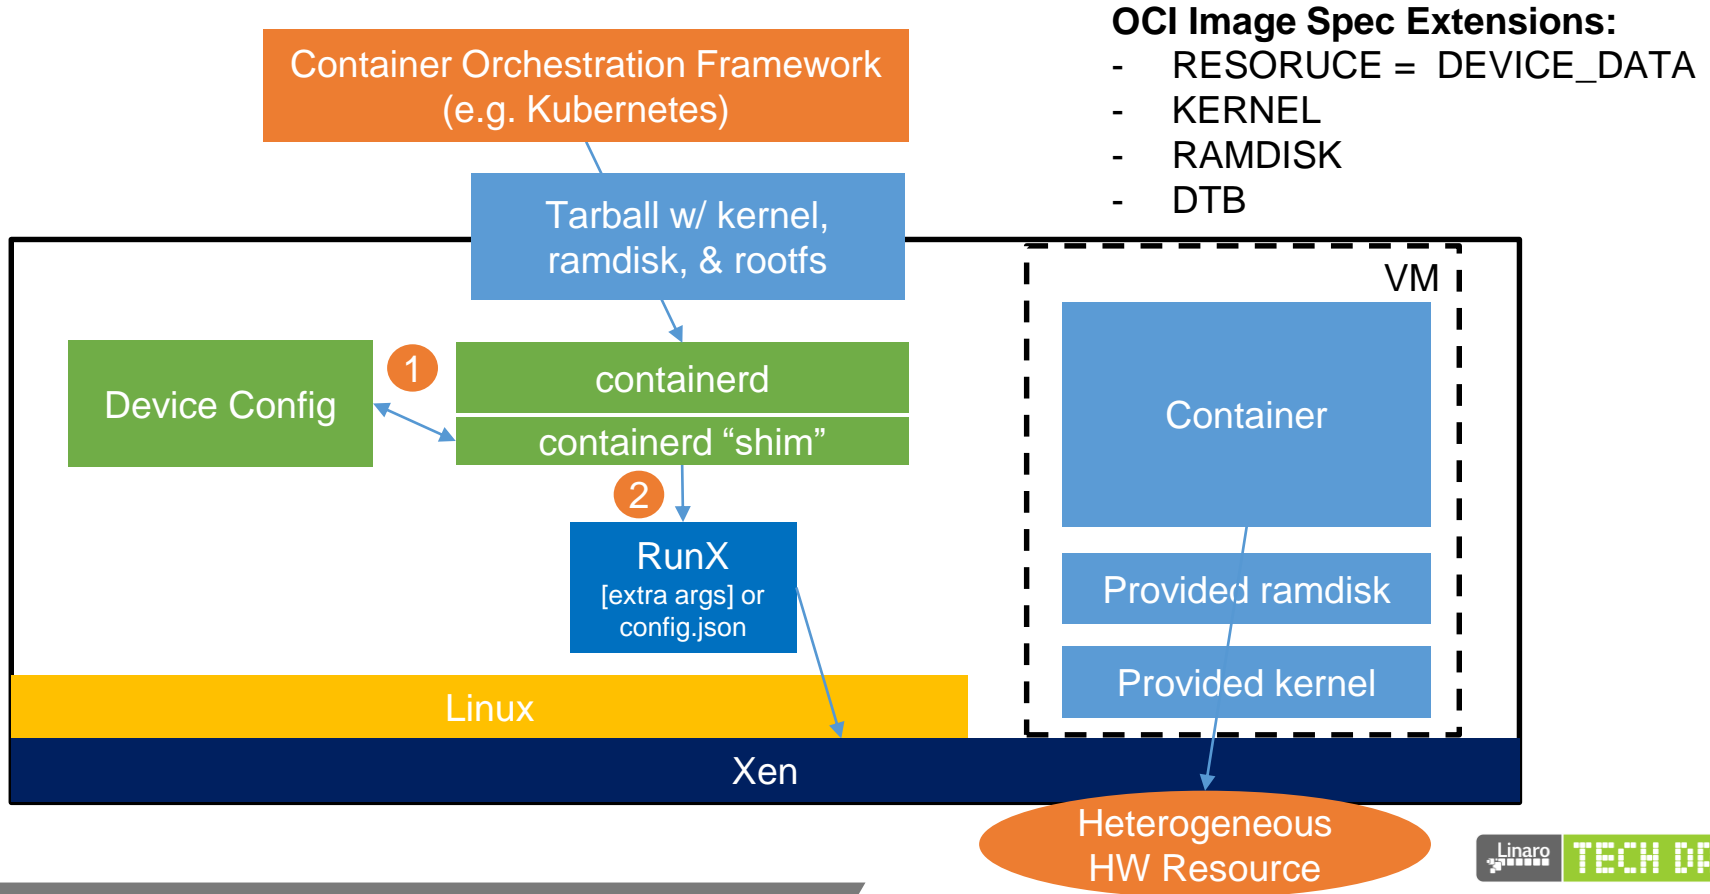
- KERNEL
- RAMDISK
- DTB



Secure Containers w/ Xen



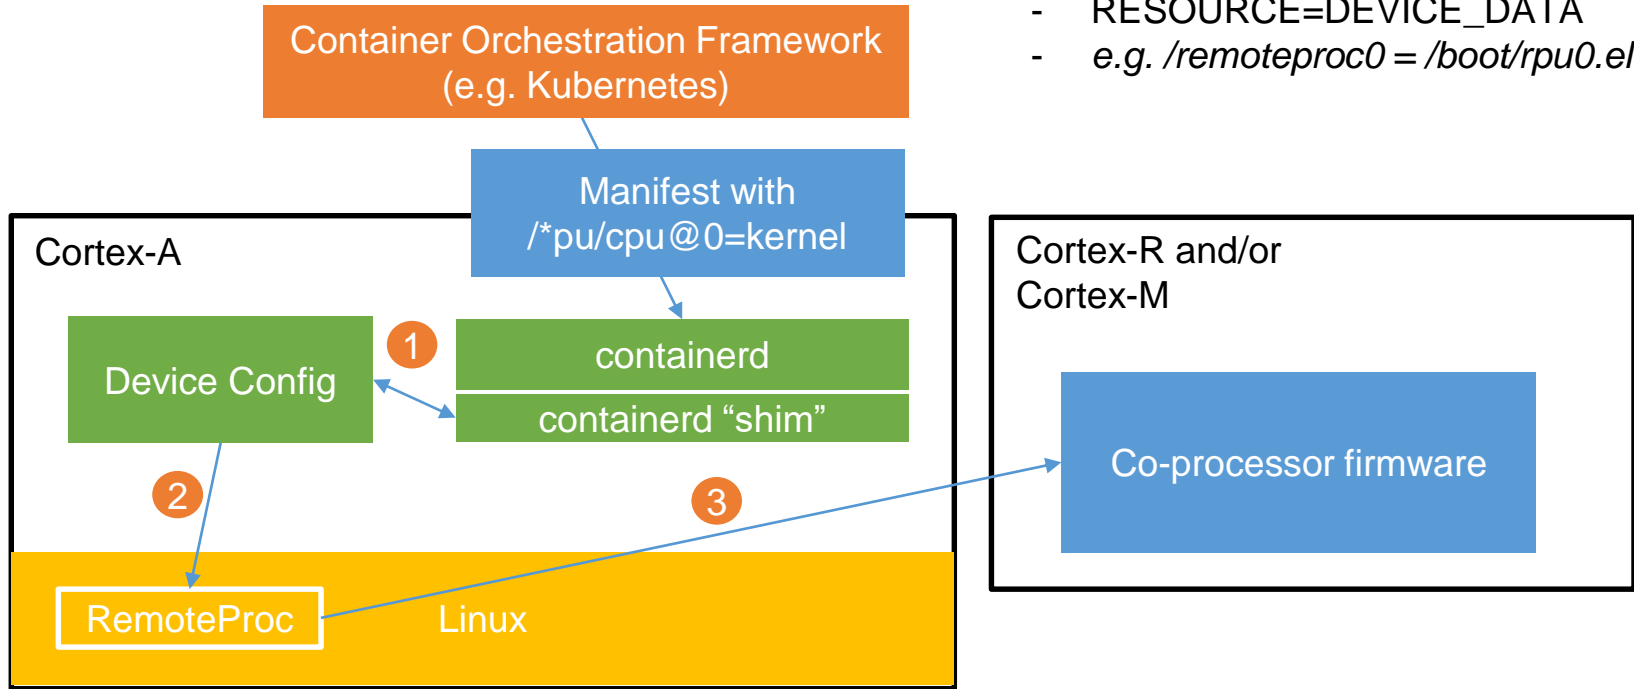
Secure Containers w/ Xen + Heterogenous HW



OCI Deployment of OpenAMP Resources

OCI Image Spec Extensions:

- RESOURCE=DEVICE_DATA
- e.g. `/remoteproc0 = /boot/rpu0.elf`



OCI Extension Topics

- OCI image specification
 - Key-value pairs for heterogenous accelerators
 - Key-value pairs for secure containers (RunX as reference implementation)
 - Device configuration data packaging within container image
- OCI runtime specification
 - Ensure containerd does not remove/filter relevant image spec data required for device configuration
 - Standardize container device configuration data passing between container & host
 - Standardize method for exporting of dynamic HW w/ container

Summary & Collaboration Request

- Xilinx development teams focused on implementing heterogeneous hardware container access architecture definition and prototyping. Requesting additional platform inputs, development collaboration, & OCI specification change advocates.
- Present collaborations underway with Project EVE within LF-Edge
 - Collaboration focused on proposed OCI changes
 - <https://wiki.lfedge.org/display/EVE/OCI+Image+Spec+for+VMs>
 - Collaboration with team at Zededa
- Proposal summary & collaboration request
 - Extend OCI specifications to support the utilization of more heterogeneous hardware capabilities within the containers ecosystem
 - Heterogeneous accelerator (FPGA, GPU, ...) deployment & access
 - OpenAMP integration for Cortex-R and Cortex-M co-processors
 - Kernel deployment (RunX as reference implementation)
 - Value is in having an integrated standards based infrastructure that can be used for deployment & orchestration of heterogeneous platform edge to cloud
 - If interested please contact Wes Skeffington (wesley.skeffington@xilinx.com) and Stefano Stabellini (stefano.stabellini@xilinx.com)



Thank you

Accelerating deployment in the Arm Ecosystem

