

From RustVMM to Kata-Containers: THE DEVELOPMENT OF H EXT. BASED SOFTWARE ECOSYSTEM

Ruoqing He


RustVMM Infra Maintainer

heruoqing@iscas.ac.cn

BACKGROUND

The virtualization software has a wide audience nowadays, they have proven themselves a very important software stack, to study and industry.
Due to the absence of SoCs with H ext., these software are not properly ported to or verified on real hardware.

OBJECTIVES




HW VERIFIED

Work closely with hardware providers, all software are well tested



UPSTREAM FIRST

Focus on contributing to upstream, affecting and benefits all distributions

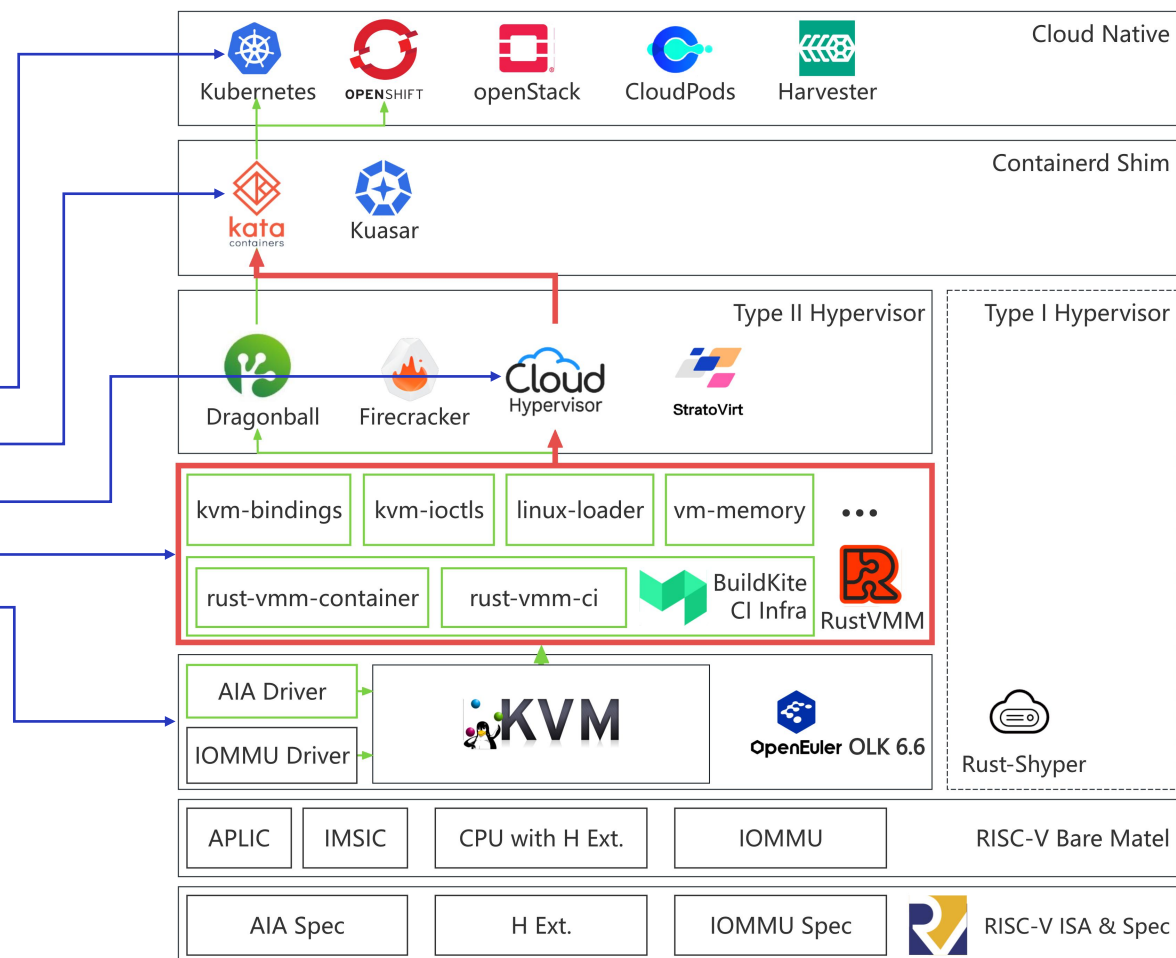


SECURE SOLUTION

Secure solution on RISC-V platform based on virtualization

RISC-V VIRTUALIZATION ROADMAP

Stage One
A bottom up, working and reliable software stack on RISC-V Orchestrated by Kubernetes, Secured by Kata-Containers, Speedup by StratoVirt, Powered by RustVMM, On openEuler RISC-V

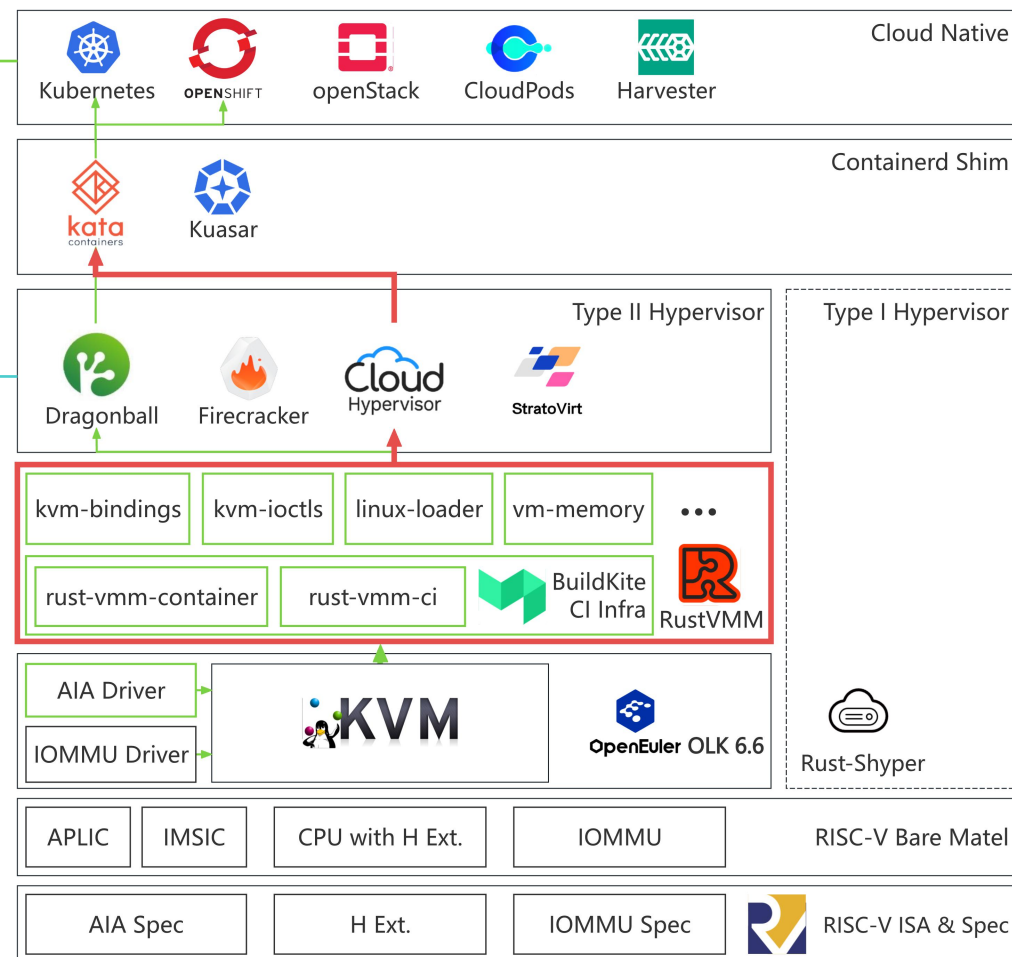


RISC-V VIRTUALIZATION ROADMAP


Stage One
A bottom up, working and reliable software stack of RISC-V

Stage Two
A fully supported, operational, prosperous virtualization software ecosystem of RISC-V

Enabling RISC-V to compete with x86 and ARM in Virtualization & CloudNative








RustVMM UPSTREAM STATUS



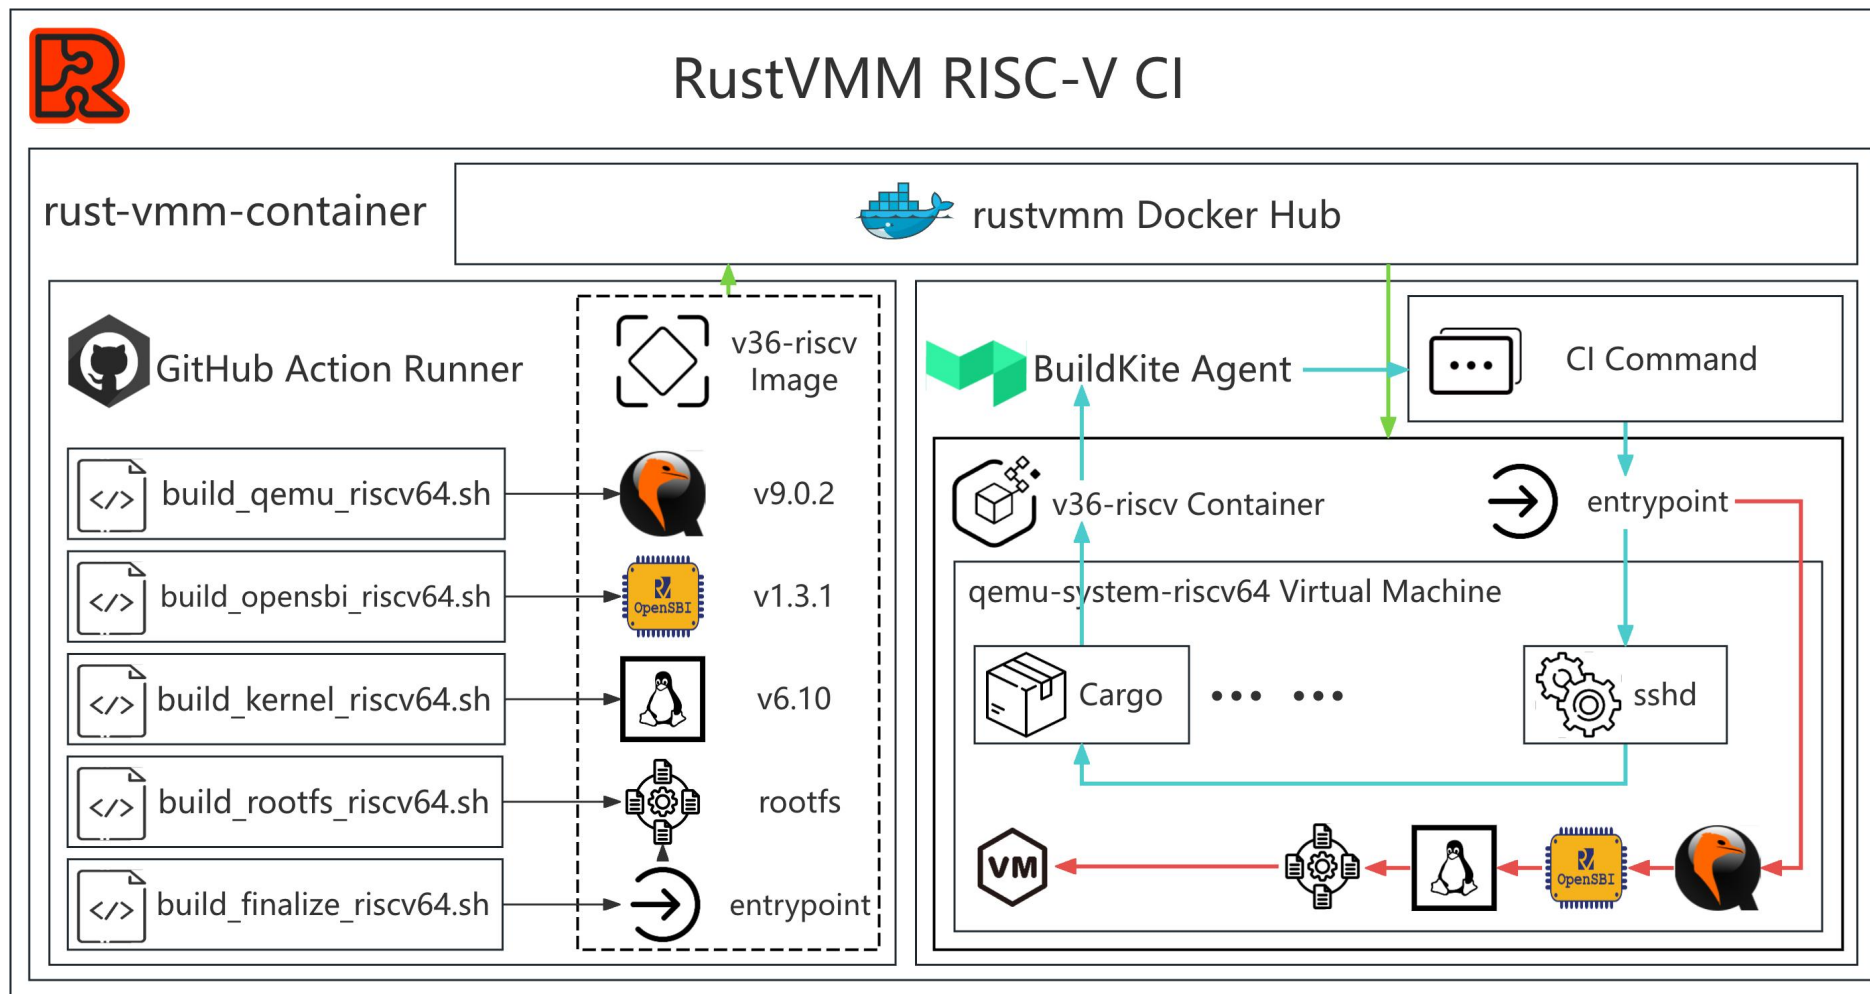
RustVMM CONTRIBUTIONS

In October 2024, after six months of research on RISC-V CI and RISC-V RustVMM counterpart, with a total of **90** commits and **107,490** lines of code changes, the RustVMM libraries officially announced complete support for RISC-V architecture. Which makes RISC-V the third officially supported architecture.

This significant breakthrough fills the gap in the RISC-V virtualization ecosystem, and provides developers with efficient, secure, and reliable virtualization solutions.

	TimePrinciple 51 commits 44,654 ++ 58,858 --	kvm	#3	...
	TimePrinciple 9 commits 121 ++ 94 --	vm-memory	#9	...
	TimePrinciple 14 commits 1,560 ++ 1,534 --	linux-loader	#5	...
	TimePrinciple 11 commits 425 ++ 142 --	rust-vmm-container	#2	...
	TimePrinciple 5 commits 85 ++ 17 --	rust-vmm-ci	#6	...

DESIGN OF RustVMM RISC-V CI



STRATOVIRT ON SPACEMIT FPGA

```
Serial-COM9 x
23.231414] usbhid: USB HID core driver
23.261859] riscv-pmu-sbi: SBI PMU extension is available
23.284954] riscv-pmu-sbi: 16 firmware and 18 hardware counters
23.452744] xt_time: kernel timezone is -0000
23.467965] IPVS: Registered protocols (TCP, UDP)
23.487760] IPVS: Connection hash table configured (size=4096, memory=32Kbytes)
23.529479] IPVS: ipvs loaded.
23.547724] IPVS: [rrr] scheduler registered.
23.589508] Initializing XFRM netlink socket
23.613998] NET: Registered PF_INET6 protocol family
23.762101] Segment Routing with IPv6
23.785935] In-situ OAM (IOAM) with IPv6
23.821529] sit: IPv6, IPv4 and MPLS over IPv4 tunneling driver
23.913769] NET: Registered PF_PACKET protocol family
23.941291] Bridge firewalling registered
23.959490] can: controller area network core
23.977849] NET: Registered PF_CAN protocol family
24.003476] can: raw protocol
24.015794] can: broadcast manager protocol
24.031790] can: netlink gateway - max_hops=1
24.054134] 8021q: 802.1Q VLAN Support v1.8
24.078046] Key type dns_resolver registered
26.302122] Loading compiled-in X.509 certificates
28.118556] Btrfs loaded, zoned=no, fsverity=no
28.212635] riscv-aplic 430500000.interrupt-controller: 255 interrupts forwarded to MSI base 0x0000000430600000
28.345558] printk: legacy console [ttyS0] disabled
28.416770] 418000000.uart0: ttyS0 at MMIO 0x418000000 (irq = 11, base_baud = 625000) is a 16550A
28.456798] printk: legacy console [ttyS0] enabled
28.456798] printk: legacy console [ttyS0] enabled
28.491590] printk: legacy bootconsole [sbi0] disabled
28.491590] printk: legacy bootconsole [sbi0] disabled
110.880464] clk: Not disabling unused clocks
110.896085] PM: genpd: Disabling unused power domains
287.609555] Initramfs unpacking failed: invalid magic at start of compressed archive
294.226524] Freeing initrd memory: 524288K
294.279360] Freeing unused kernel image (initmem) memory: 2320K
294.287958] Run /init as init process
294.292650] with arguments:
294.296564] /init
294.299730] with environment:
294.304628] HOME=/
294.307977] TERM=linux
Init from initramfs
Running sysctl: OK
Saving 256 bits of non-creditable seed for next boot
Mounting cgroupfs hierarchy: OK
Starting iptables: OK
Starting dhcpd...
dhcpd-9.4.1 starting
[ 309.216253] random: crng init done
dhcp_vendor: Invalid argument
forked to background, child pid 148
no interfaces have a carrier

Welcome to Buildroot
buildroot login: █
```

硬件环境：进迭时空高性能
服务器芯片 FPGA 原型

10








openEuler RISC-V VIRT WG

Cloud Hypervisor

Maintained by: openEuler RISC-V VIRT WG

RISE PROJECT

Home / ... / Kernel and Virtualizat... / Kernel and Virtualization Projects						     	
LK_03_039	QEMU-KVM Message Proxy support	LK_02_030					
LK_03_040	RustVMM rust-vmm-container RISC-V support		ISCAS	Ruoqing He			
LK_03_041	RustVMM rust-vmm-ci RISC-V support	LK_03_040	ISCAS	Ruoqing He			
LK_03_042	RustVMM kvm-bindings introduce RISC-V bindings	LK_03_041	ISCAS	Ruoqing He			
LK_03_043	RustVMM kvm-bindings fam-wrappers & serde support	LK_03_041 , LK_03_042	ISCAS	Ruoqing He			
LK_03_044	RustVMM kvm-ioctls RISC-V support	LK_03_041 , LK_03_043	ISCAS	Ruoqing He			
LK_03_045	RustVMM linux-loader RISC-V support	LK_03_041	ISCAS	Ruoqing He			
LK_03_046	RustVMM vm-memory RISC-V support	LK_03_041	ISCAS	Ruoqing He			
LK_03_047	Cloud-Hypervisor RISC-V support	LK_03_041 , LK_03_042 , LK_03_043 , LK_03_044 , LK_03_045 , LK_03_046	ISCAS	Ruoqing He			
LK_03_048	Dragonball RISC-V support	LK_03_041 , LK_03_042 , LK_03_043 , LK_03_044 , LK_03_045 , LK_03_046	ISCAS	Ruoqing He			

openEuler RISC-V VIRT WG



Honghao Cheng
CHonghaohao

- Rust-Shyper RISC-V
AIA



Ruqing He
TimePrincipl

- RustVMM
- StratoVirt
- Dragonball
- CloudHypervisor
- Firecracker
- KataContainers



Xinhao Yi
TexasOct

- Kubernetes
component images
- Kubeadm
deployment
support

THANKS