



Red Hat OpenShift

KVM 基座能力来支持
OpenShift Virtualization


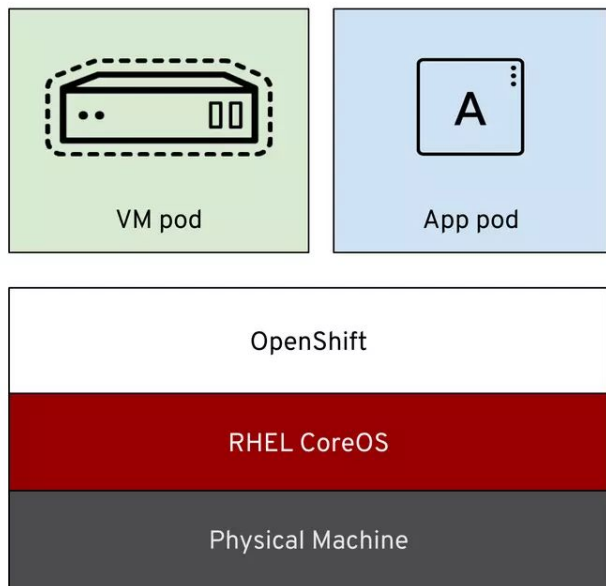
Wenkang Ji <wji@redhat.com>
KVM-QE Team

目录

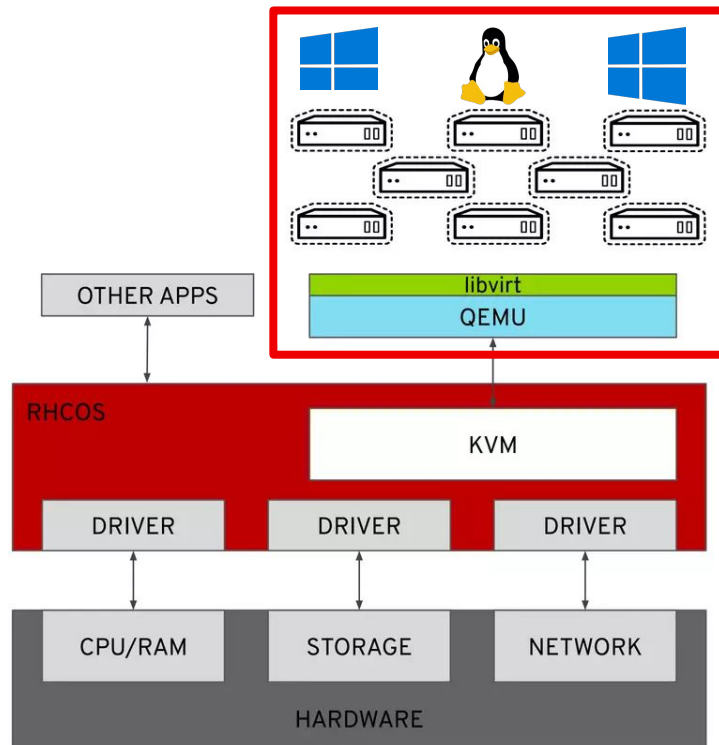
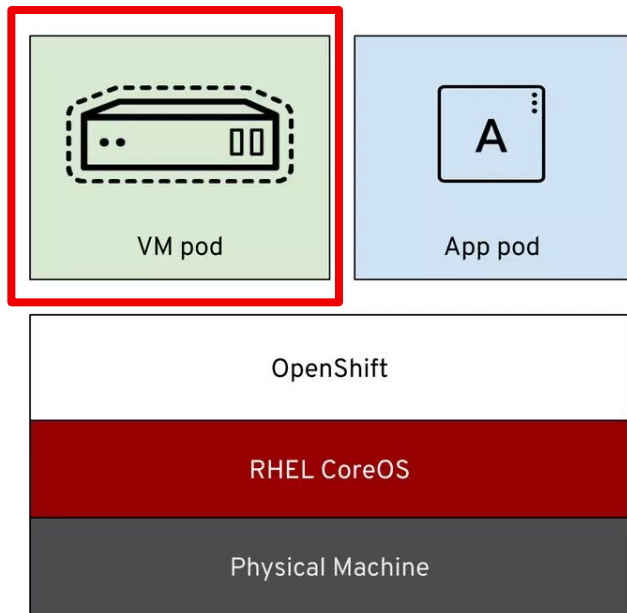
- ❑ **OpenShift Virtualization 的基座 KVM**
- ❑ OpenShift Virtualization 对比 Qemu-kvm cmdline
- ❑ **Virtio-Win 组是做什么的？**
- ❑ OpenShift Virtualization 从社区到客户
- ❑ **Qemu-kvm 底层的虚拟机指标收集工具**
- ❑ Qemu-kvm 底层的 CPU 内存的热插拔
- ❑ Qemu-kvm 底层的 容灾备份
- ❑ **Qemu-kvm 底层的热迁移**
- ❑ **Industry Edge: OpenShift Virtualization Real Time**

OpenShift Virtualization 的基座 KVM

OpenShift 介绍



OpenShift 介绍



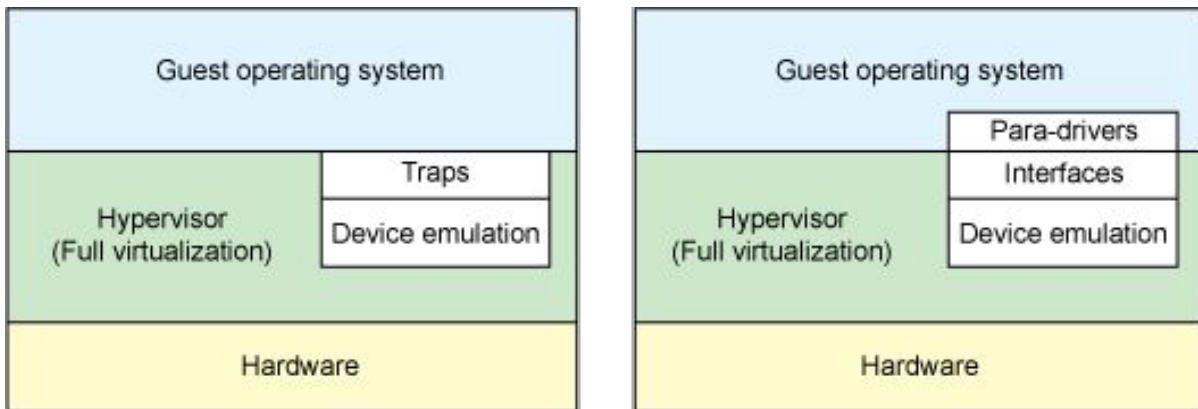
Qemu-kvm 虚拟化 <https://www.qemu.org/>

虚拟化由三个部分组成：KVM, Qemu 和 LibVirt

- Qemu 是虚拟化的底层

只有我们 Qemu 底层支持了， LibVirt 作为上层才可以调用。

- KVM 是一个内核模块。我们可以使用 `--enable-kvm` 开启全虚拟化支持，否则单纯的 Qemu 采用二进制翻译的方式进行虚拟化，性能较为低下。



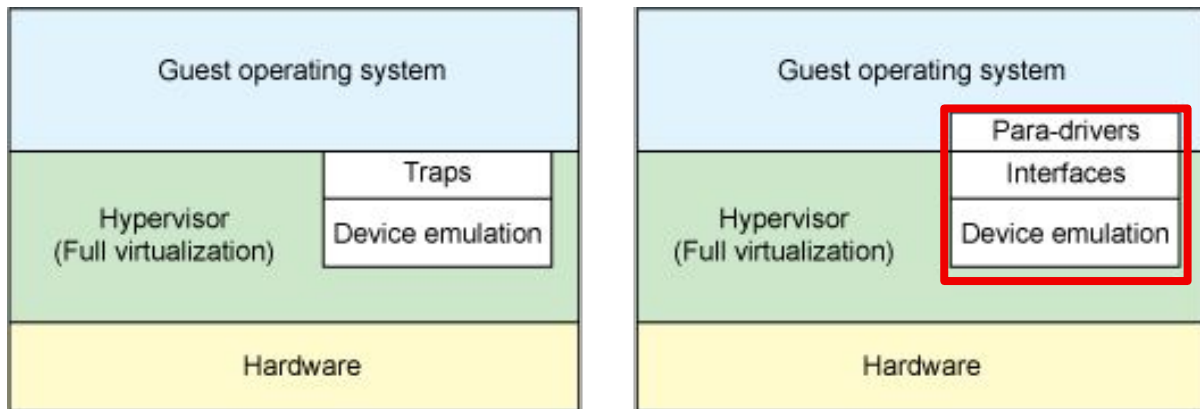
Qemu-kvm 虚拟化 <https://www.qemu.org/>

虚拟化由三个部分组成：KVM, Qemu 和 LibVirt

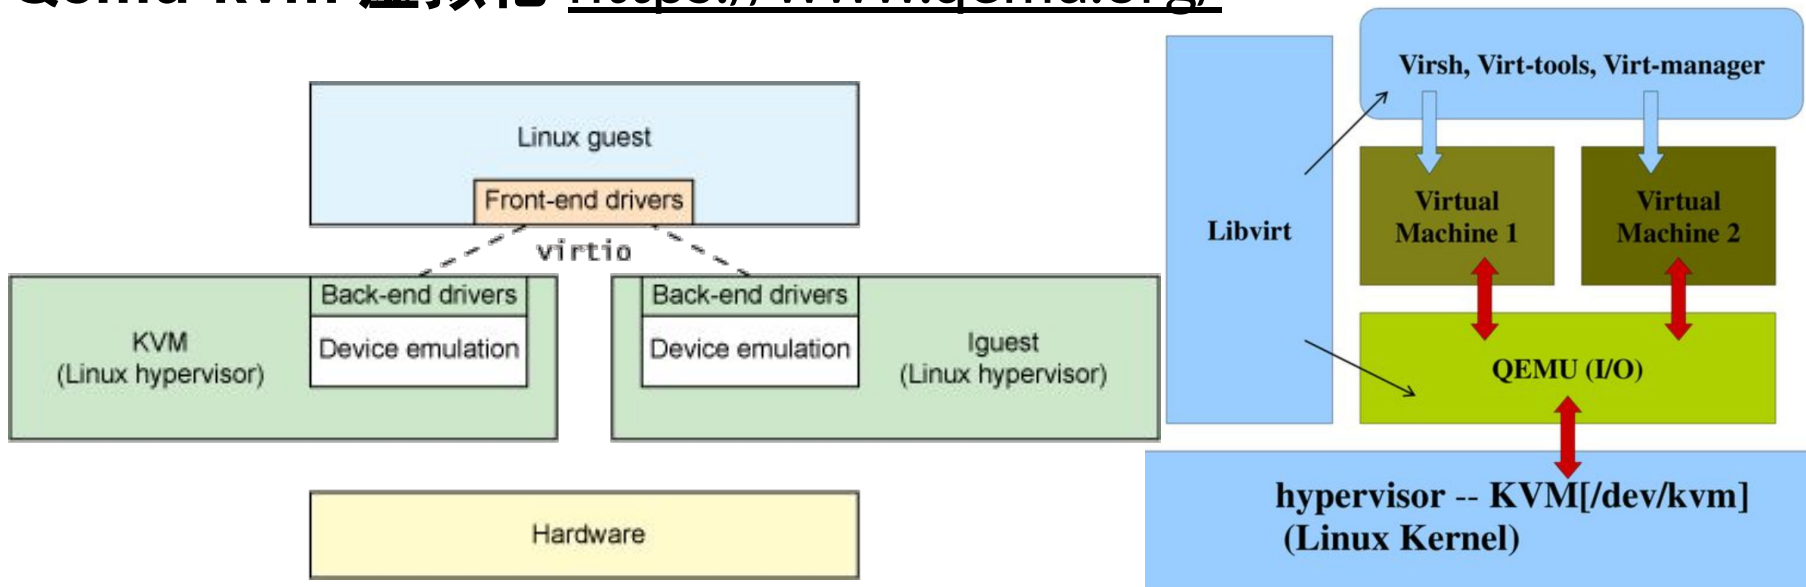
- Qemu 是虚拟化的底层

只有我们 Qemu 底层支持了， LibVirt 作为上层才可以调用。

- KVM 是一个内核模块。我们可以使用 `--enable-kvm` 开启全虚拟化支持，否则单纯的 Qemu 采用二进制翻译的方式进行虚拟化，性能较为低下。




Qemu-kvm 虚拟化 <https://www.qemu.org/>



VirtIO 是一种**半虚拟化手段**。通过让虚拟机知道自己是虚拟机的方式，将其部分虚拟化处理的转为物理机处理以提高整体性能。虚拟机可以直接访问宿主机上的物理设备资源，从而避免了在虚拟机内部进行大量的模拟计算。

OpenShift VM 对比 Qemu-kvm cmdline

OpenShift Virtualization 介绍



Red Hat OpenShift

管理员

项目: 所有项目

Pod > Pod详情

virt-launcher-wji-win2k22-fxqkr Running

详情 指标 YAML 环境 日志 事件 终端

连接到 compute

```
sh-5.1$ whoami
qemu
sh-5.1$ virsh list
Authorization not available. Check if polkit service is running or see debug message for more information.
  Id Name State
-----
  1 default_wji-win2k22 running

sh-5.1$ ps aux | grep qemu-kvm
qemu      101  4.2  1.0 6870972 4326240 ?        Ss
secret", "id": "masterKey0", "format": "raw", "file": "/usr/share/OVMF/OVMF_CODE.secboot.flash0-format", "read-only": true, "driver": "raw", "nvram/wji-win2k22_VARS.fd", "node-name": "libvirt-pflash1-storage", "driver": "raw", "file": "libvirt-pflash1-storage-format,pflash1=libvirt-pflash1-format -accel kvm rdpid=on,fsrm=on,md-clear=on,stibp=on,arch-capability=on,pschange-mc-no=on,tsx-ctrl=on,hle=off,rtm=off,rtx=on,hv-runtime=on,hv-sync=on,hv-stimer=on,hv-stimulash01,property=secure,value=on -m 4096 -object {"
```

```
sh-5.2$ whoami
```

```
wji
```

```
sh-5.2$ oc project
```

Using project "default" on server

```
"https://api.kvm-ocp.wenkangji.com:6443".
```

```
sh-5.2$ oc get vm
```

```
NAME      AGE STATUS READY
```

```
rhel9-first 56d Running True
```

```
wji-win2k22 8h Running True
```

```
sh-5.2$ oc get pods
```

```
NAME                                READY STATUS RESTARTS AGE
```

```
virt-launcher-rhel9-first-mpdn8 1/1 Running 0 55d
```

```
virt-launcher-wji-win2k22-fxqkr 2/2 Running 0 8h
```

```
sh-5.2$ oc rsh virt-launcher-wji-win2k22-fxqkr
```

```
sh-5.1$ whoami
```

```
qemu
```

```
sh-5.1$ virsh list
```

Authorization not available. Check if polkit service is running or see debug message for more information.

```
Id Name State
```

```
-----
```

```
1 default_wji-win2k22 running
```

```
sh-5.1$ ps aux | grep qemu-kvm
```

```
qemu      101  4.3  1.0 6870972 4326228 ?        Ss 05:35 23:42
```

```
/usr/libexec/qemu-kvm -name
```

```
quest=default_wji-win2k22,debug-threads=on
```



```

/usr/libexec/qemu-kvm
-name guest=default_wji-win2k22,debug-threads=on
-s
-object {"qom-type":"secret","id":"masterKey0","format":"raw","file":"/var/run/kubevirt-private/libvirt/qemu/lib/domain-1-default_wji-win2k22/master-key.aes"}
-blockdev {"driver":"file","filename":"/usr/share/OVMF/OVMF_CODE.seeboot.fd","node-name":"libvirt-pflash0-storage","auto-read-only":true,"discard":"unmap"}
-blockdev {"node-name":"libvirt-pflash0-format","read-only":true,"driver":"raw","file":"libvirt-pflash0-storage"}
-blockdev {"driver":"file","filename":"/var/run/kubevirt-private/libvirt/qemu/nvram/wji-win2k22_VARS.fd","node-name":"libvirt-pflash1-storage","auto-read-only":true,"discard":"unmap"}
-blockdev {"node-name":"libvirt-pflash1-format","read-only":false,"driver":"raw","file":"libvirt-pflash1-storage"}
-machine pc-q35-rhel9.2.0,usb=off,smm=on,dump-guest-core=off,memory-backend=pc.ram,pflash0=libvirt-pflash0-format,pflash1=libvirt-pflash1-format
-accel kvm
-cpu
!clacka-Server ds=on,ss=on,dtes64=on,vmx=on,pdcm=on,hypervisor=on,tsc-adjust=on,avx512fma=on,sha-ni=on,rdpid=on,fsrm=on,md-clear=on,stlbp=on,arch-capabilities=on,xsaves=on,ibpb=on,ibrs=on,amd-stlbp=on,amd-sbds=on,rdctl-no=on,ibrs-all=on,skip-ldlfi-vmentry=on,mds-no=on,pschange-mc-no=on,tsc-ctrl=on,hle=off,rtm=off,mxcsr=off,intel-pt=off,hv-time=on,tsc-frequency=2394374000,hv-relaxed=on,hv-vapic=on,hv-spinlocks=0x1fff,hv-vpindex=on,hv-runtime=on,hv-synic=on,hv-stimer=on,hv-stimer-direct=on,hv-reset=on,hv-frequencies=on,hv-reenlightenment=on,hv-tlbfush=on,hv-ipi=on
-global driver=cfi.pflash01,property=secure,value=on
-m 4096
-object {"qom-type":"memory-backend-ram","id":"pc.ram","size":4294967296}
-overcommit mem-lock=off
-smp 1,sockets=1,dies=1,cores=1,threads=1
-object {"qom-type":"iothread","id":"iothread1"}
-uuid 423d1c64-26aa-5e21-ad1f-e4ad2265ff9
-smbios type=1,manufacturer=Red Hat,product=OpenShift Virtualization,version=4.15.1,uuid=423d1c64-26aa-5e21-ad1f-e4ad2265ff9,family=Red Hat
-no-user-config
-nodefaults
-chardev socket,id=charmonitor,fd=18,server=on,wait=off
-mon chardev=charmonitor,id=monitor,mode=control
-rtc base=utc,driftfix=slew
-global kvm-pit.lost_tick_policy=delay
-no-hpet
-no-shutdown
-boot strict=on
-device {"driver":"pcie-root-port","port":0,"chassis":0,"id":"pci.0","bus":"pcie.0","addr":"0x2.0x3"}
-device {"driver":"pcie-root-port","port":1,"chassis":1,"id":"pci.1","bus":"pcie.0","addr":"0x2.0x4"}
-device {"driver":"pcie-root-port","port":2,"chassis":2,"id":"pci.2","bus":"pcie.0","addr":"0x2.0x5"}
-device {"driver":"pcie-root-port","port":3,"chassis":3,"id":"pci.3","bus":"pcie.0","addr":"0x2.0x6"}
-device {"driver":"pcie-root-port","port":4,"chassis":4,"id":"pci.4","bus":"pcie.0","addr":"0x2.0x7"}
-device {"driver":"pcie-root-port","port":5,"chassis":5,"id":"pci.5","bus":"pcie.0","addr":"0x2.0x8"}
-device {"driver":"pcie-root-port","port":6,"chassis":6,"id":"pci.6","bus":"pcie.0","addr":"0x2.0x9"}
-device {"driver":"pcie-root-port","port":7,"chassis":7,"id":"pci.7","bus":"pcie.0","addr":"0x2.0xA"}
-device {"driver":"pcie-root-port","port":8,"chassis":8,"id":"pci.8","bus":"pcie.0","addr":"0x2.0xB"}
-device {"driver":"pcie-root-port","port":9,"chassis":9,"id":"pci.9","bus":"pcie.0","addr":"0x3"}
-device {"driver":"qemu-xhci","id":"usb","bus":"pci.5","addr":"0x0"}
-device {"driver":"virtio-scsi-pci-non-transitional","id":"scsi0","bus":"pci.6","addr":"0x0"}
-device {"driver":"virtio-serial-pci-non-transitional","id":"virtio-serial0","bus":"pci.7","addr":"0x0"}
-blockdev {"driver":"host_device","filename":"/dev/rootdisk","aio":"native","node-name":"libvirt-4-storage","cache":{"direct":true,"no-flush":false},"auto-read-only":true,"discard":"unmap"}
-blockdev {"node-name":"libvirt-4-format","read-only":false,"discard":"unmap","cache":{"direct":true,"no-flush":false},"driver":"raw","file":"libvirt-4-storage"}
-device {"driver":"ide-hd","bus":"ide.0","drive":"libvirt-4-format","read-only":true,"discard":"unmap","cache":{"direct":true,"no-flush":false},"driver":"raw","file":"libvirt-4-storage"}
-blockdev {"driver":"file","filename":"/var/run/kubevirt/container-disks/disk_1.img","node-name":"libvirt-3-storage","cache":{"direct":true,"no-flush":false},"auto-read-only":true,"discard":"unmap"}
-blockdev {"node-name":"libvirt-3-format","read-only":true,"discard":"unmap","cache":{"direct":true,"no-flush":false},"driver":"raw","file":"libvirt-3-storage"}
-blockdev {"driver":"file","filename":"/var/run/kubevirt/ephemeral-disks/disk-data/windows-drivers-disk/qcow2","node-name":"libvirt-2-storage","cache":{"direct":true,"no-flush":false},"auto-read-only":true,"discard":"unmap"}
-blockdev {"node-name":"libvirt-2-format","read-only":true,"discard":"unmap","cache":{"direct":true,"no-flush":false},"driver":"qcow2","file":"libvirt-2-storage","backing":"libvirt-3-format"}
-device {"driver":"ide-cd","bus":"ide.1","drive":"libvirt-2-format","id":"ua-windows-drivers-disk","bootindex":3,"write-cache":"on","werror":"stop","rerror":"stop"}
-blockdev {"driver":"host_device","filename":"/dev/installation-cdrom","aio":"native","node-name":"libvirt-1-storage","cache":{"direct":true,"no-flush":false},"auto-read-only":true,"discard":"unmap"}
-blockdev {"node-name":"libvirt-1-format","read-only":true,"discard":"unmap","cache":{"direct":true,"no-flush":false},"driver":"raw","file":"libvirt-1-storage"}
-device {"driver":"ide-cd","bus":"ide.2","drive":"libvirt-1-format","id":"ua-installation-cdrom","bootindex":4,"write-cache":"on","werror":"stop","rerror":"stop"}
-netdev {"type":"tap","fd":"19","id":"hostua-default"}
-device {"driver":"e1000e","netdev":"hostua-default","id":"ua-default","mac":"02:e1:81:00:00:0c","bus":"pci.1","addr":"0x0","romfile":""}
-chardev socket,id=charserial0,fd=16,server=on,wait=off
-device {"driver":"isa-serial","chardev":"charserial0","id":"serial0","index":0}
-chardev socket,id=charchannel0,fd=17,server=on,wait=off
-device {"driver":"virtserialport","bus":"virtio-serial0.0","nr":1,"chardev":"charchannel0","id":"channel0","name":"org.qemu.guest_agent.0"}
-chardev socket,id=chrtpm,path=/var/run/libvirt/qemu/run/swtspm/1-default_wji-win2k22-swtspm.sock
-tpmdev emulator,id=tpm-tpm0,chardev=chrtpm
-device {"driver":"tpm-tis","tpmdev":"tpm-tpm0","id":"tpm0"}
-device {"driver":"usb-tablet","id":"ua-tablet","bus":"usb.0","port":"1"}
-audiodev {"id":"audiol1","driver":"none"}
-vnc vnc-unix:/var/run/kubevirt-private/5d79f1c9-5ba3-4f4a-a5e8-1c6858f413a4/virt-vnc,audiodev=audiol1

```

Too hard for user to use the qemu-kvm DIRECTLY!!!



kind: VirtualMachine

cmdline: qemu-kvm

cpu:

cores: 1

sockets: 1

threads: 1

-cpu

lcelake-Server,ds=on,ss=on,dtes64=on,vmx=on,pdcm=on,hypervisor=on,tsc-adjust=on,avx512ifma=on,sha-ni=on,rdpid=on,fsrm=on,md-clear=on,stibp=on,arch-capabilities=on,xsaves=on,ibpb=on,ibrs=on,amd-stibp=on,amd-ssbd=on,rdctl-no=on,ibrs-all=on,skip-lldfl-vmentry=on,mds-no=on,pschange-mc-no=on,tsx-ctrl=on,hle=off,rtm=off,mpx=off,intel-pt=off,hv-time=on,tsc-frequency=2394374000,hv-relaxed=on,hv-vapic=on,hv-spinlocks=0x1fff,hv-vpinde x=on,hv-runtime=on,hv-synic=on,hv-stimer=on,hv-stimer-direct=on,hv-reset=on,hv-frequencies=on,hv-reenlightenment=on,hv-tlbflush=on,hv-ipi=on -smp 1,sockets=1,dies=1,cores=1,threads=1

CPU 部分

memory:

guest: 4Gi

-m 4096

内存 部分

```
devices: -device
interfaces: {"driver":"e1000e","netdev":"hostua-default","id":"ua-default","mac":"02:e1:81:00:00:0c"}
- macAddress: 02:e1:81:00:00:0c
masquerade: {}
model: e1000e
name: default
-netdev {"type":"tap","fd":"19","id":"hostua-default"}
-device
{"driver":"pcie-root-port","port":16,"chassis":1,"id":"pci.1","bus":"pcie.0","multifunction":true,"addr":"0x2"}
```

网卡部分

```
devices: -blockdev
disks: {"driver":"host_device","filename":"/dev/rootdisk","aio":"native","node-name":"libvirt-4-storage","cache":{"direct":true,"no-flush":false},"auto-read-only":true,"discard":"unmap"}
- bootOrder: 2
disk: -blockdev
bus: sata {"node-name":"libvirt-4-format","read-only":false,"discard":"unmap","cache":{"direct":true,"no-flush":false},"driver":"raw","file":"libvirt-4-storage"}
name: rootdisk -device
{"driver":"ide-hd","bus":"ide.0","drive":"libvirt-4-format","id":"ua-rootdisk","bootindex":2,"write-cache":"on","werror":"stop","rerror":"stop"}
```


磁盘部分

```
machine: -machine
type: pc-q35-rhel9.2.0 pc-q35-rhel9.2.0,usb=off,smm=on,dump-guest-core=off,memory-backend=pc.ram,pflash0=libvirt-pflash0-format,pflash1=libvirt-pflash1-format
firmware:
bootloader: -chardev socket,id=char_vtpm_tpm0,path=/tmp/guest-swtpm.sock \
efi: -tpmdev emulator,chardev=char_vtpm_tpm0,id=emulator_vtpm_tpm0 \
secureBoot: true -device tpm-crb,id=tpm-crb_vtpm_tpm0,tpmdev=emulator_vtpm_tpm0 \
```

固件部分



Virtio-Win 组是做什么的



每个文件夹都是一个具体的功能。

每个功能都是 QE + Dev 同时负责。

作为 QE 主要负责:

- Bug reported
- function test
- WHQL test
- SVVP test

SVVP 认证

SVVP (Server Virtualization Validation Program) 是 Microsoft 推出的计划, 旨在提升客户在虚拟化技术上运行 Windows Server 时的支持体验。该计划允许虚拟化产品供应商验证其整个虚拟化解决方案, 使客户能够在这些经过验证的虚拟化环境中运行 Windows Server 并获得技术支持。

以 OpenShift Virtualization 为例, 我们 QE 会使用 Microsoft 的硬件测试套件对 OpenShift 中的运行的虚拟机进行全面测试, 确认其兼容性和稳定性。然后将测试结果打包并提交给 Microsoft 进行认证。通过认证后, 我们的 OpenShift Virtualization 就获得了在虚拟化环境中运行 Windows Server 的支持资格。

总结

SVVP 是针对虚拟化环境中运行 Windows Server 的系统测试认证, 确保虚拟化产品能够兼容并稳定地运行 Windows Server, 为客户提供可靠的技术支持。这个认证过程涵盖了整个虚拟化解决方案, 而不仅仅是单个驱动程序或特性。

Drivers

Device metadata

File Signing Services

IoT

Device model hierarchy

Analyze

Partner analytics

Driver flight status

Driver install and health summary

<u>Private Product ID</u>	<u>Product Name</u>	<u>Submission Status</u>	<u>Submission Created Date</u> ↓	<u>Submission Type</u>	<u>Permission</u>	<u>Source</u>	<u>Shared Product ID</u>	<u>Submission ID</u>	<u>Certification Type</u>
14151641070558861	Red Hat VirtIO NetKVM Drivers for Windows Server 2016	Complete	12/4/2023	Initial	Author	Red Hat, Inc.	401581428	1152921505697102514	HLK
14622131746130626	Red Hat VirtIO NetKVM Drivers for Windows 11	Complete	12/4/2023	Initial	Author	Red Hat, Inc.	401581540	1152921505697102051	HLK
14566800634744199	Red Hat VirtIO NetKVM Drivers for Windows 10	Complete	12/4/2023	Initial	Author	Red Hat, Inc.	401581427	1152921505697102511	HLK
14356239045148425	Red Hat OpenShift Container Platform 4.14	Complete	10/23/2023	Initial	Author	Red Hat, Inc.	401528612	1152921505696928591	HLK
14506907384338080	Red Hat OpenShift Container Platform 4.14	Complete	10/23/2023	Initial	Author	Red Hat, Inc.	401528643	1152921505696928748	HLK
14215570904046445	virtio-win-prewhql-0.1-242-x64	Complete	10/10/2023	Initial	Author	Red Hat, Inc.	401513261	1152921505696872923	Attestation
13795175919970282	Red Hat Enterprise Linux Version 8.9	Complete	9/21/2023	Initial	Author	Red Hat, Inc.	401492978	1152921505696801660	HLK
14294893626140636	Red Hat Enterprise Linux Version 8.9	Complete	9/20/2023	Initial	Author	Red Hat, Inc.	401492486	1152921505696798723	HLK
14201698163753225	Red Hat Enterprise Linux Version 9.3	Complete	9/18/2023	Initial	Author	Red Hat, Inc.	401488794	1152921505696787325	HLK
13599019641588778	Red Hat Enterprise Linux Version 9.3	Complete	9/18/2023	Initial	Author	Red Hat, Inc.	401488743	1152921505696787272	HLK



Microsoft

Hardware certification report **Approved**

Private product ID: **14506907384338080**
Shared product ID: **401528643**
Submission ID: **1152921505696928748**
Submission date: **10/23/2023**
Completion date: **11/1/2023**
Company: **Red Hat, Inc.**
Product name: **Red Hat OpenShift Container Platform 4.14**
Category: **System**
Product type: **Server Virtualization Validation Program**
Qualification level: **Certified for Microsoft Windows Hardware Compatibility Program – Server family version 21H2, x64**
Marketing name: **Red Hat OpenShift Container Platform 4.14**

Microsoft
Certified

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - SCSIAdapter - 100.94.104.24700](#)

[Red Hat, Inc. - Net - 100.94.104.24200](#)

[Red Hat, Inc. - NetTrans - 100.94.104.2](#)

[Red Hat, Inc. - NetTrans - 100.94.104.2](#)

[Red Hat, Inc. - NetTrans - 100.94.104.2](#)

[Red Hat, Inc. - NetTrans - 100.94.104.2](#)

[Red Hat, Inc. - NetTrans - 100.94.104.2](#)

Microsoft Update Catalog - Google Chrome

Update Details

Red Hat, Inc. - SCSIAdapter - 100.94.104.24700

Last Modified: 1/22/2024

Size: 45 KB

UpdateID: 3d9369fa-fc04-4fce-9fdb-3f949f7b2510

Details:

Overview

Language Selection

Package Details

Install Resources

Description: Red Hat, Inc. SCSIAdapter driver update released in January 2024

Architecture: AMD64

Classification: Drivers

Supported products: Windows 10, Vibranium and later, Servicing Drivers , Windows 10, Vibranium and later, Upgrade & Servicing Drivers

Supported languages: Arabic , Bulgarian , Chinese (Traditional) , Czech , Danish , German , Greek , English , Spanish , Finnish , French , Hebrew , Hungarian , Italian , Japanese , Korean , Dutch , Norwegian , Polish , Portuguese (Brazil) , Romanian , Russian , Croatian , Slovak , Swedish , Thai , Turkish , Ukrainian , Slovenian , Estonian , Latvian , Lithuanian , Chinese (Simplified) , Portuguese (Portugal) , Serbian (Latin) , Chinese - Hong Kong SAR

Company: Red Hat, Inc.

Driver Manufacturer: Red Hat, Inc.

Driver Class: OtherHardware

Driver Model: Red Hat VirtIO SCSI controller

Driver Provider: Red Hat, Inc.

Version: 100.94.104.24700

Version Date: 1/22/2024

More information:

<https://learn.microsoft.com/en-us/windows-hardware/drivers/dashboard/hardware-submission-support>

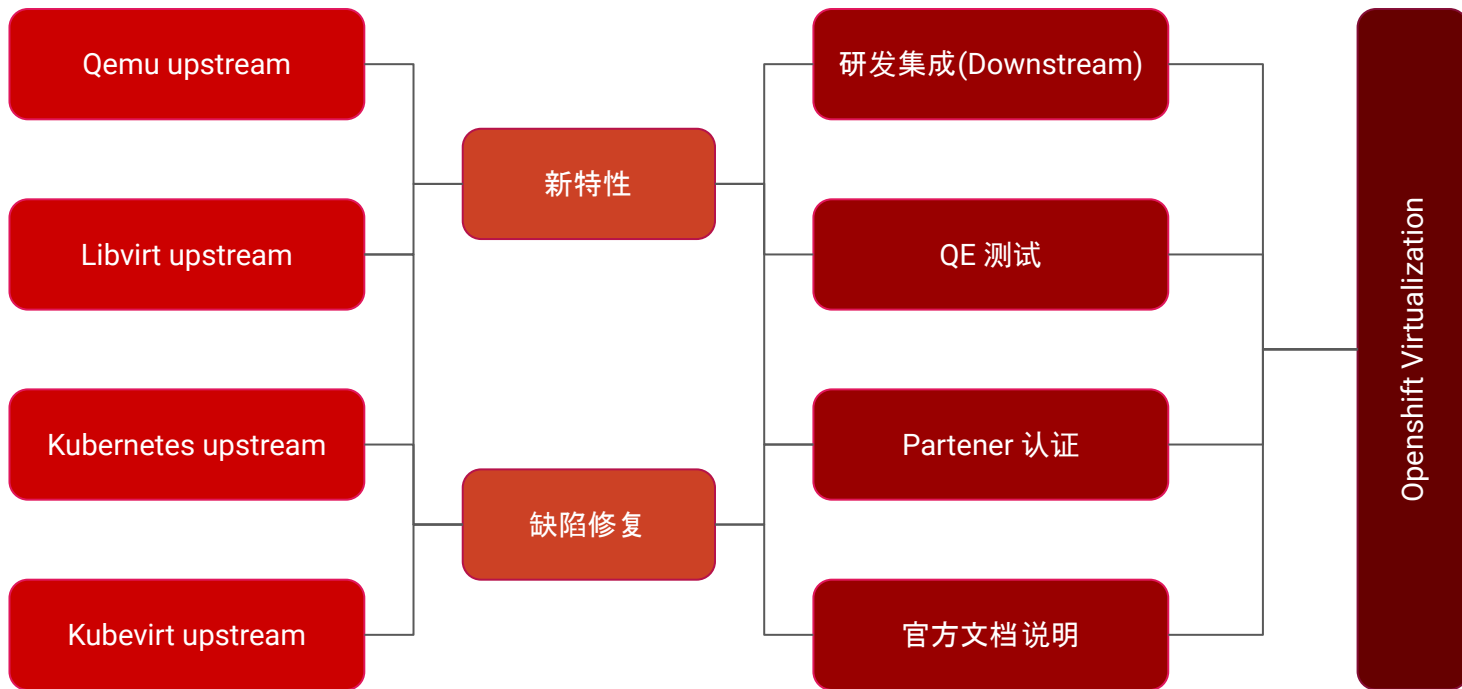
Support Url:

<https://support.microsoft.com/select/?target=hub>

Openshift Virtualization

从社区到客户

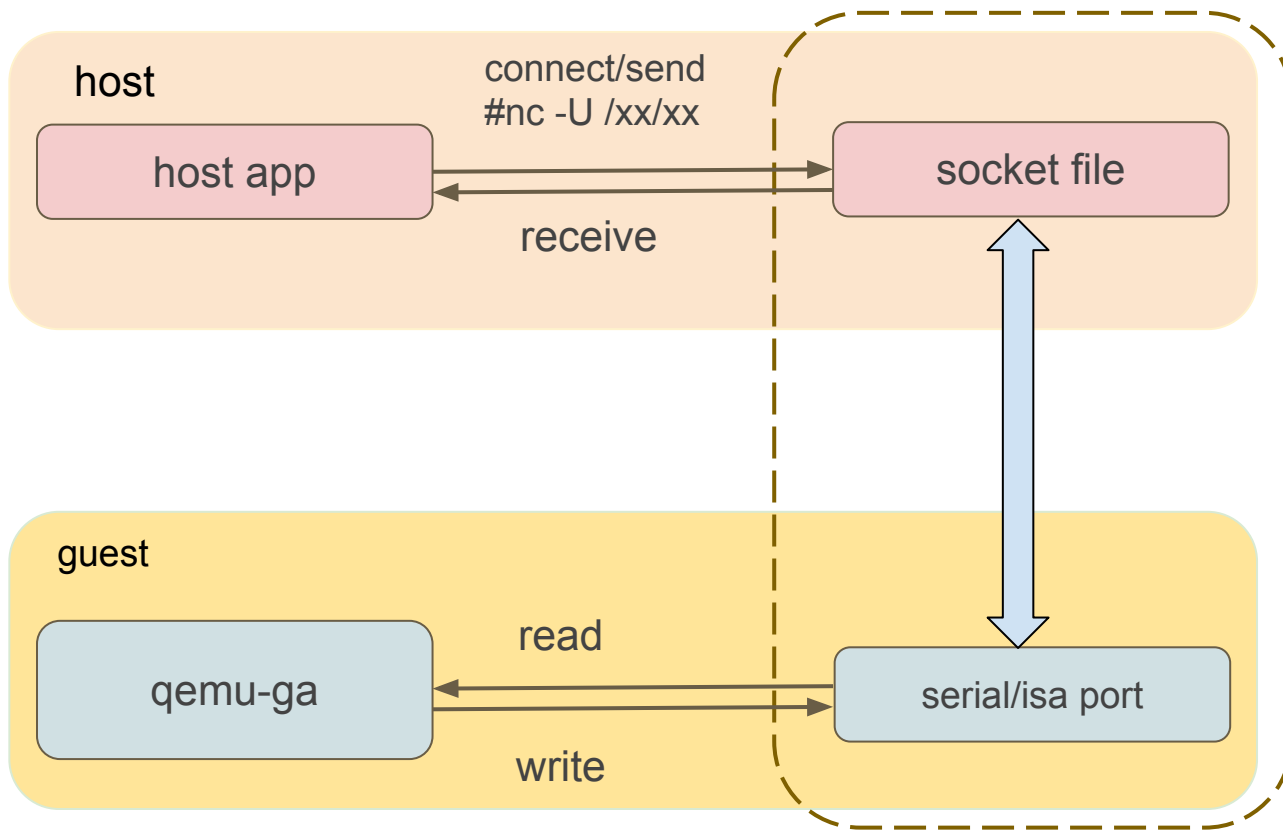
Openshift Virtualization 从社区到客户



底层虚拟机信息收集

Qemu-Guest-Agent

What is Qemu-Guest-Agent?



RHEL Version of QGA

```
# cat /etc/redhat-release
```

```
Red Hat Enterprise Linux release 9.3 (Plow)
```

```
# rpm -q qemu-guest-agent
```

```
qemu-guest-agent-8.0.0-16.el9_3.3.x86_64
```

```
# rpm -ql qemu-guest-agent
```

```
/etc/qemu-ga
```

```
/etc/qemu-ga/fsfreeze-hook
```

```
/etc/qemu-ga/fsfreeze-hook.d
```

```
/etc/sysconfig/qemu-ga
```

```
/usr/bin/qemu-ga
```

```
/usr/lib/.build-id
```

```
/usr/lib/.build-id/c5
```

```
/usr/lib/.build-id/c5/b489aac998b30b29ce3a8d15
```

```
64c8cb3aa8e4cf
```

```
/usr/lib/systemd/system/qemu-guest-agent.service
```

```
/usr/lib/udev/rules.d/99-qemu-guest-agent.rules
```

```
/usr/share/doc/qemu-guest-agent
```

```
/usr/share/doc/qemu-guest-agent/COPYING
```

```
/usr/share/doc/qemu-guest-agent/README.rst
```

```
/usr/share/man/man8/qemu-ga.8.gz
```

```
/usr/share/qemu-kvm/qemu-ga
```

```
/usr/share/qemu-kvm/qemu-ga/fsfreeze-hook.d
```

```
/usr/share/qemu-kvm/qemu-ga/fsfreeze-hook.d/
```

```
mysql-flush.sh.sample
```

```
/var/log/qemu-ga
```

```
#
```

```
# rpm -qi qemu-guest-agent
```

```
Name      : qemu-guest-agent
```

```
Epoch    : 17
```

```
Version   : 8.0.0
```

```
Release   : 16.el9_3.3
```

```
Architecture: x86_64
```

```
Install Date: Thu Feb 29 08:58:52 2024
```

```
Group     : Unspecified
```

```
Size      : 1871632
```

```
License   : GPLv2 and GPLv2+ and CC-BY
```

```
Signature : RSA/SHA256, Tue Dec 5 21:03:55 2023, Key ID 199e2f91fd431d51
```

```
Source RPM : qemu-kvm-8.0.0-16.el9_3.3.src.rpm
```

```
Build Date : Tue Dec 5 02:56:51 2023
```

```
Build Host : x86-64-02.build.eng.rdu2.redhat.com
```

```
Packager   : Red Hat, Inc. <http://bugzilla.redhat.com/bugzilla>
```

```
Vendor     : Red Hat, Inc.
```

```
URL        : http://www.qemu.org/
```

```
Summary    : QEMU guest agent
```

```
Description :
```

```
qemu-kvm is an open source virtualizer that provides hardware emulation for the KVM hypervisor.
```

```
This package provides an agent to run inside guests, which communicates with the host over a virtio-serial channel named "org.qemu.guest_agent.0"
```

```
This package does not need to be installed on the host OS.
```

```
#
```


Windows Version of QGA

Server Manager Dashboard

CD Drive (E:) Virtio-Win

Name	Date modified	Type	Size
Balloon	10/30/2021 11:26 PM	File folder	
fwcfg	10/30/2021 11:24 PM	File folder	
guest-agent	10/30/2021 11:48 PM	File folder	
NetKVM	10/30/2021 11:31 PM	File folder	
pvpanic	10/30/2021 11:35 PM	File folder	
qemufwcfg	10/30/2021 11:24 PM	File folder	
qemupciserial	10/30/2021 11:36 PM	File folder	
qxl	10/30/2021 11:48 PM	File folder	
sriov	10/30/2021 11:24 PM	File folder	
viosfs	10/30/2021 11:47 PM	File folder	
viogpudo	10/30/2021 11:24 PM	File folder	
vioinput	10/30/2021 11:38 PM	File folder	
viornio	10/30/2021 11:41 PM	File folder	
vioscsi	10/30/2021 11:43 PM	File folder	
vioserial	10/30/2021 11:45 PM	File folder	
viostor	10/30/2021 11:47 PM	File folder	

16 items 1 item selected

BPA results



File Explorer: guest-agent

Name	Date modified	Type	Size
qemu-ga-i386	10/30/2021 11:48 PM	Windows Installer ...	3,015 KB
qemu-ga-x86_64	10/30/2021 11:48 PM	Windows Installer ...	3,222 KB



File Explorer: guest-agent

Name	Date modified	Type	Size
qemu-ga-i386	10/30/2021 11:48 PM	Windows Installer ...	3,015 KB
qemu-ga-x86_64	10/30/2021 11:48 PM	Windows Installer ...	3,222 KB

QEMU guest agent

Please wait while Windows configures QEMU guest agent

Time remaining: 0 seconds

Cancel

Common commands

guest-ping: Return all commands supported by qga ;

guest-get-fsinfo: Get guest file system information ;

guest-fstrim: Discard (or "trim") blocks which are not in use by the filesystem ;

guest-suspend-disk*: Suspend guest to disk ; (currently: enable/disable) logical processors inside the guest ;

guest-suspend-ram*: Suspend guest to ram ;

guest-suspend-hybrid: Save guest state to disk and suspend to ram (This command requires the pm-utils package to be installed in the guest.) ;

guest-network-get-interfaces: Get list of guest IP addresses, MAC addresses and netmasks ;

guest-get-memory-blocks: Retrieve the list of the guest's memory blocks.

guest-set-memory-blocks: Attempt to reconfigure (currently: enable/disable) state of memory blocks inside the guest ;

guest-get-memory-block-info: Get information relating to guest memory blocks ;

guest-exec: Execute command (asynchronous operation) ;

guest-exec-status: View the result of executing the command ;

Common commands

guest-sync-delimited: Check the synchronization status of the host and qga communication, State synchronization check of the handover process between clients when the client communicates with qga;

guest-sync: like above ,but no 0xff ;

guest-ping: ping guest agent,a non-error return implies success ;

guest-get-time:Get the guest time (the return value is relative to 1970-01-01in UTC, Timein nanoseconds.) ;

guest-set-time: Set the guest time

guest-shutdown: Initiate guest-activated shutdown (this is an asynchronous shutdown request, with no guarantee of successful shutdown.) ;

guest-get-users: Get current login username ;

guest-get-host-name: Get hostname of current guest ;

guest-set-user-password:Set user's login password

guest-get-vcpus: Retrieve the list of the guest's logical processors;

guest-set-vcpus: Attempt to reconfigure (currently: enable/disable) logical processors inside the guest ;

Common commands

guest-file-flush: Write file changes buffered in userspace to disk/kernel buffers ;

guest-file-open : Open target file (return file handle);

guest-file-close : Close the files in the opened virtual machine;

guest-file-read: Read the file content in the guest according to the file handle (return the file content in base64 format);

guest-file-write : Write the contents of the file to the file in the virtual machine according to the file handle;

guest-file-seek: Seek to a position in the file, as with `fseek()`, and return the current file position afterward. Also encapsulates `ftell()`'s functionality, just Set offset=0, whence=SEEK_CUR;

guest-fsfreeze-status: Get guest fsfreeze state. error state indicates ;

guest-fsfreeze-freeze-list: Sync and freeze specified guest filesystems ;

guest-fsfreeze-freeze: Sync and freeze all freezable, local guest filesystems ;

guest-fsfreeze-thaw: Unfreeze all frozen guest filesystems ;

guest-get-disks:

guest-ssh-add-authorized-keys:

guest-ssh-remove-authorized-keys:

guest-ssh-get-authorized-keys , "arguments": {"username": "fedora"}

- [1967716](#)

Reference from:

[QEMU Guest Agent Protocol Reference](#)

Get os info via "guest-get-osinfo"

Step	Expected Result
<p>1. Start guest with virtio serial and start guest agent inside the guest</p> <p>e.g..</p> <pre>-device virtio-serial-pci,id=virtio-serial0,bus=pcie.0,addr=0x7 \ -chardev socket,path=/tmp/qga.sock,server,nowait,id=qga0 \ -device virtserialport,bus=virtio-serial0.0,chardev=qga0,name=org.qemu.guest_agent.0</pre>	<p>boot up successfully and qemu-ga works well</p>
<p>2. get os info via qga</p> <pre>{"execute": "guest-get-osinfo"}</pre>	<p>return the guest os info: such as: win2019 guest</p> <pre>{"return": {"name": "Microsoft Windows", "kernel-release": "17763", "version": "Microsoft Windows Server 2019", "variant": "server", "pretty-name": "Windows Server 2019 Datacenter", "version-id": "2019", "variant-id": "server", "kernel-version": "10.0", "machine": "x86_64", "id": "mswindows"}}</pre> <p>rhel8.1 guest:</p> <pre>{"return": {"name": "Red Hat Enterprise Linux", "kernel-release": "4.18.0-103.el8.x86_64", "version": "8.1 (Ootpa)", "pretty-name": "Red Hat Enterprise Linux 8.1 Beta (Ootpa)", "version-id": "8.1", "kernel-version": "#1 SMP Sat Jun 8 15:07:32 UTC 2019", "machine": "x86_64", "id": "rhel"}}</pre>

Execute a command with "guest-exec" and check status with "guest-exec-status"

Step	Expected Result
<p>1. Start guest with virtio serial and start guest agent inside the guest</p> <p>e.g..</p> <pre>-device virtio-serial-pci,id=virtio-serial0,bus=pcie.0,addr=0x7 \ -chardev socket,path=/tmp/qga.sock,server,nowait,id=qga0 \ -device virtserialport,bus=virtio-serial0.0,chardev=qga0,name=org.qemu.guest_agent.0</pre>	<p>boot up successfully and qemu-ga works well</p>
<p>2. Issue cmd with wrong args to get the error output.</p> <p>linux guest:</p> <pre>{ "execute": "guest-exec", "arguments": { "path": "/usr/bin/date", "arg": ["+%Y-%m-%d"], "capture-output": true }} {"return": {"pid": 3894}} {"execute": "guest-exec-status", "arguments": {"pid": 3894}}</pre> <p>windows guest:</p> <pre>{ "execute": "guest-exec", "arguments": { "path": "whoami", "capture-output": true }} {"return": {"pid": 4672}} {"execute": "guest-exec-status", "arguments": {"pid": 4672}}</pre>	<p>linux guest:</p> <pre>{"return": {"exitcode": 0, "out-data": "MjAyNC0wNi0xMgo=", "exited": true}} decode out-data: # echo MjAyNC0wNi0xMgo= base64 -d 2024-06-12</pre> <p>windows guest:</p> <pre>{"return": {"exitcode": 0, "out-data": "bnQgYXV0aG9yaXR5XHN5c3RlbQ0K", "exited": true}} decode out-data: nt authority\system</pre>

Qemu-kvm 底层的容灾备份

Qemu-kvm 功能: 虚拟机的容灾备份 - Qemu 快照

```
# qemu-img snapshot -c haha xxx.qcow2
```

'snapshot' is the name of the snapshot to create, apply or delete

```
# qemu-img snapshot -d haha xxx.qcow2
```

'-a' applies a snapshot (revert disk to saved state)

```
# qemu-img snapshot -l xxx.qcow2
```

'-c' creates a snapshot

```
# qemu-img snapshot -a haha xxx.qcow2
```

'-d' deletes a snapshot

'-l' lists all snapshots in the given image

Qemu-kvm 功能: 虚拟机的容灾备份 - Libvirt 快照 (Live)

```
# virsh snapshot-create avocado-vt-vm1
```

```
Domain snapshot 1717585245 created
```

```
# virsh snapshot-list avocado-vt-vm1
```

Name	Creation Time	State
------	---------------	-------

1717585245	2024-06-05 19:00:45 +0800	running
------------	---------------------------	---------

```
# virsh snapshot-revert avocado-vt-vm1 1717585245
```

```
Domain snapshot 1717585245 reverted
```

```
# virsh snapshot-delete avocado-vt-vm1 1717585245
```

```
Domain snapshot 1717585245 deleted
```

Snapshot (help keyword 'snapshot')

snapshot-create	Create a snapshot from XML
snapshot-create-as	Create a snapshot from a set of args
snapshot-current	Get or set the current snapshot
snapshot-delete	Delete a domain snapshot
snapshot-dumpxml	Dump XML for a domain snapshot
snapshot-edit	edit XML for a snapshot
snapshot-info	snapshot information
snapshot-list	List snapshots for a domain
snapshot-parent	Get the name of the parent of a snapshot
snapshot-revert	Revert a domain to a snapshot

Qemu-kvm 底层的 CPU 内存的热插拔

Qemu-kvm 功能: 热插拔 qemu cmdline

```
/usr/libexec/qemu-kvm \  
-name 'avocado-vt-vm1' \  
-sandbox on,elevateprivileges=deny,obsolete=deny,resourcecontrol=deny \  
-blockdev '{"node-name": "file_ovmf_code", "driver": "file", "filename": "/usr/share/OVMF/OVMF_CODE.secboot.fd", "auto-read-only": true, "discard": "unmap"}' \  
-blockdev '{"node-name": "drive_ovmf_code", "driver": "raw", "read-only": true, "file": "file_ovmf_code"}' \  
-blockdev '{"node-name": "file_ovmf_vars", "driver": "file", "filename": "/root/avocado/data/avocado-vt/avocado-vt-vm1_rhel930-64-virtio-scsi-ovmf_qcow2_filesystem_VARS.raw", "auto-read-only": true, "discard": "unmap"}' \  
-blockdev '{"node-name": "drive_ovmf_vars", "driver": "raw", "read-only": false, "file": "file_ovmf_vars"}' \  
-device '{"id": "pcie-root-port-0", "driver": "pcie-root-port", "multifunction": true, "bus": "pcie.0", "addr": "0x1", "chassis": 1}' \  
-device '{"id": "pcie-pci-bridge-0", "driver": "pcie-pci-bridge", "addr": "0x0", "bus": "pcie-root-port-0"}' \  
-nodefaults \  
-device '{"driver": "VGA", "bus": "pcie.0", "addr": "0x2"}' \  
-machine q35.pflash0=drive_ovmf_code.pflash1=drive_ovmf_vars.memory-backend=mem1 \  
-m 1024.slots=256.maxmem=40G \  
-object memory-backend-ram,id=mem1,size=1G \  
-smp 1.threads=2.cores=1.sockets=3.maxcpus=6 \  
-cpu EPYC-IBPB,x2apic=on,tsc-deadline=on,hypervisor=on,tsc-adjust=on,arch-capabilities=on,cmp-legacy=on,perfctr-core=on,clzero=on,xsaveptr=on,virt-ssbd=on,npt=on,lbrv=on,nrip-save=on,tsc-scale=on,vmcb-clean=on,pause-filter=on,pfthreshold=on,v-vmsave-vmload=on,vgif=on,svme-addrchk=on,lfence-always-serializing=on,rdctl-no=on,skip-l1dfl-vmentry=on,mds-no=on,pschange-mc-no=on,gds-no=on,monitor=off,kvm_pv_unhalt=on \  
-device '{"ioport": 1285, "driver": "pvpanic", "id": "idq8EXEd"}' \  
-device '{"id": "pcie-root-port-1", "port": 1, "driver": "pcie-root-port", "addr": "0x1.0x1", "bus": "pcie.0", "chassis": 2}' \  
-device '{"driver": "qemu-xhci", "id": "usb1", "bus": "pcie-root-port-1", "addr": "0x0"}' \  
-device '{"driver": "usb-tablet", "id": "usb-tablet1", "bus": "usb1.0", "port": "1"}' \  
-device '{"id": "pcie-root-port-2", "port": 2, "driver": "pcie-root-port", "addr": "0x1.0x2", "bus": "pcie.0", "chassis": 3}' \  
-device '{"id": "virtio_scsi_pci0", "driver": "virtio-scsi-pci", "bus": "pcie-root-port-2", "addr": "0x0"}' \  
.....
```

Qemu-kvm 功能: CPU 热插拔

```
{"execute": "query-hotpluggable-cpus"}
{"return": [{"props": {"core-id": 0, "thread-id": 1, "socket-id": 2}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 2}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 1, "socket-id": 1}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 1}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 0}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 0}, "vcpu-count": 1, "qom-path": "/machine/unattached/device[0]", "type": "EPYC-IBPB-x86_64-cpu"}]}
```

```
{"execute": "device_add", "arguments": {"driver": "EPYC-IBPB-x86_64-cpu", "core-id": 0, "thread-id": 1, "socket-id": 2, "id": "core1"}}
{"return": {}}
{"timestamp": {"seconds": 1717587239, "microseconds": 464843}, "event": "ACPI_DEVICE_OST", "data": {"info": {"device": "core1", "source": 1, "status": 0, "slot": "5", "slot-type": "CPU"}}
```

```
{"execute": "query-hotpluggable-cpus"}
{"return": [{"props": {"core-id": 0, "thread-id": 1, "socket-id": 2}, "vcpu-count": 1, "qom-path": "/machine/peripheral/core1", "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 2}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 1, "socket-id": 1}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 1}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 1, "socket-id": 0}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 0}, "vcpu-count": 1, "type": "EPYC-IBPB-x86_64-cpu"}, {"props": {"core-id": 0, "thread-id": 0, "socket-id": 0}, "vcpu-count": 1, "qom-path": "/machine/unattached/device[0]", "type": "EPYC-IBPB-x86_64-cpu"}]}
```

```
{"execute": "device_del", "arguments": {"id": "core1"}}
{"return": {}}
{"timestamp": {"seconds": 1717587312, "microseconds": 715366}, "event": "ACPI_DEVICE_OST", "data": {"info": {"device": "core1", "source": 3, "status": 132, "slot": "5", "slot-type": "CPU"}}}
{"timestamp": {"seconds": 1717587312, "microseconds": 733275}, "event": "DEVICE_DELETED", "data": {"path": "/machine/peripheral/core1/lapic"}}
{"timestamp": {"seconds": 1717587312, "microseconds": 733680}, "event": "DEVICE_DELETED", "data": {"device": "core1", "path": "/machine/peripheral/core1"}}
{"timestamp": {"seconds": 1717587312, "microseconds": 735741}, "event": "ACPI_DEVICE_OST", "data": {"info": {"source": 3, "status": 0, "slot": "5", "slot-type": "CPU"}}}
```

Qemu-kvm 功能: CPU 热插拔

```
[root@localhost ~]# cat /proc/cpuinfo | grep "^processor" -c  
1  
[root@localhost ~]# cat /proc/cpuinfo | grep "^processor" -c  
2  
[root@localhost ~]# cat /proc/cpuinfo | grep "^processor" -c  
1
```

Qemu-kvm 功能: 内存 热插拔

```
{"execute": "object-add", "arguments": {"id": "mem2", "qom-type": "memory-backend-ram", "policy":  
"default", "size": 1073741824}}  
{"return": {}}  
{"execute": "device_add", "arguments": {"id": "dimm1", "driver": "pc-dimm", "memdev": "mem2"}}  
{"return": {}}  
{"timestamp": {"seconds": 1717589771, "microseconds": 347223}, "event": "ACPI_DEVICE_OST", "data": {"info": {"device":  
"dimm1", "source": 1, "status": 0, "slot": "0", "slot-type": "DIMM"}}}}  
{"execute": "query-memory-devices"}  
{"return": [{"type": "dimm", "data": {"memdev": "/objects/mem2", "hotplugged": true, "addr": 4294967296, "hotpluggable":  
true, "size": 1073741824, "slot": 0, "node": 0, "id": "dimm1"}}]}  
{"execute": "device_del", "arguments": {"id": "dimm1"}}  
{"return": {}}  
{"timestamp": {"seconds": 1718206738, "microseconds": 329833}, "event": "ACPI_DEVICE_OST", "data": {"info": {"device":  
"dimm1", "source": 3, "status": 132, "slot": "0", "slot-type": "DIMM"}}}}  
{"timestamp": {"seconds": 1718206738, "microseconds": 333628}, "event": "DEVICE_DELETED", "data": {"device": "dimm1",  
"path": "/machine/peripheral/dimm1"}}  
{"timestamp": {"seconds": 1718206738, "microseconds": 336193}, "event": "ACPI_DEVICE_OST", "data": {"info": {"source":  
3, "status": 0, "slot": "0", "slot-type": "DIMM"}}}}  
{"execute": "object-del", "arguments": {"id": "mem2"}}  
{"return": {}}  
{"execute": "query-memory-devices"}  
{"return": []}
```

Qemu-kvm 功能: 内存 热插拔

```
[root@localhost ~]# lsmem
RANGE                                SIZE  STATE  REMOVABLE  BLOCK
0x0000000000000000-0x000000003fffffff 1G  online        yes   0-7
0x00000000100000000-0x0000000013fffffff 1G  offline       no    32-39

Memory block size:      128M
Total online memory:    1G
Total offline memory:   1G
[root@localhost ~]# free -mh
              total        used         free      shared  buff/cache   available
Mem:          682Mi        508Mi         61Mi         36Ki        200Mi        173Mi
Swap:         2.0Gi         405Mi         1.6Gi

[root@localhost ~]# for i in `ls /sys/devices/system/memory/memory*/state`; do echo online > $i ;done
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
-bash: echo: write error: Invalid argument
[root@localhost ~]# free -mh
              total        used         free      shared  buff/cache   available
Mem:          1.7Gi        715Mi         930Mi         40Ki        260Mi        991Mi
Swap:         2.0Gi         402Mi         1.6Gi

[root@localhost ~]# lsmem
RANGE                                SIZE  STATE  REMOVABLE  BLOCK
0x0000000000000000-0x000000003fffffff 1G  online        yes   0-7
0x00000000100000000-0x0000000013fffffff 1G  online        yes  32-39

Memory block size:      128M
Total online memory:    2G
Total offline memory:   0B
[root@localhost ~]#
```


Qemu-kvm 功能: 内存 热插拔

```
[root@localhost ~]# for i in `ls /sys/devices/system/memory/memory*/state`; do echo offline > $i ;done
-bash: echo: write error: Invalid argument
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Invalid argument
[root@localhost ~]# free -mh
              total        used         free      shared  buff/cache   available
Mem:           810Mi       588Mi       100Mi         40Ki       240Mi       222Mi
Swap:          2.0Gi       391Mi       1.6Gi
[root@localhost ~]# lsmem
RANGE                SIZE  STATE  REMOVABLE  BLOCK
0x0000000000000000-0x000000003fffffff  1G  online      yes    0-7
0x0000000010000000-0x00000000107fffffff 128M online      yes    32
0x00000000108000000-0x0000000013fffffff 896M offline           33-39

Memory block size:      128M
Total online memory:    1.1G
Total offline memory:   896M
[root@localhost ~]#
```

Qemu-kvm 功能: 内存 热插拔

```
[root@localhost ~]# for i in `ls /sys/devices/system/memory/memory*/state`; do echo offline > $i ;done
-bash: echo: write error: Invalid argument
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Device or resource busy
-bash: echo: write error: Invalid argument
[root@localhost ~]# free -mh
```

	total	used	free	shared	buff/cache	available
Mem:	810Mi	588Mi	100Mi	40Ki	240Mi	222Mi
Swap:	2.0Gi	391Mi	1.6Gi			

```
[root@localhost ~]# lsmem
```

RANGE	SIZE	STATE	REMOVABLE	BLOCK
0x0000000000000000-0x000000003fffffff	1G	online	yes	0-7
0x0000000010000000-0x00000000107fffffff	128M	online	yes	32
0x0000000010800000-0x00000000113fffffff	896M	offline		33-39

```
Memory block size:      128M
Total online memory:    1.1G
Total offline memory:   896M
[root@localhost ~]#
```

1. grubby
--args="memhp_default_state=online_movable"
--update-kernel
/boot/vmlinuz-4.18.0-144.el8.ppc64le
2. grubby --update-kernel
/boot/vmlinuz-4.18.0-144.el8.ppc64le
3. grubby --info
/boot/vmlinuz-4.18.0-144.el8.ppc64le

Qemu-kvm 底层的热迁移


Qemu-kvm 功能: 热迁移

Preparation (live): 准备阶段, 虚拟机运行中, 内存页开始复制。

Pre-Copy Rounds: 多个轮次复制脏页。

Downtime: 短暂停机, 复制剩余内存页。

Resume Time (live): 虚拟机在目标主机上恢复运行。




Preparation (live): 准备阶段, 虚拟机运行中, 非可分页内存复制。

Downtime: 短暂停机, 传输CPU状态和最小状态信息。

Resume Time (live): 虚拟机在目标主机上恢复运行, 内存页按需传输。

Post-Copy Preparing: 后台继续传输剩余内存页。



```
virsh migrate <VMname> --live  
qemu+ssh://<domain name>/system  
--verbose --postcopy
```

```
# virsh migrate --help | grep postcopy  
--postcopy          enable post-copy migration; switch to it using migrate-postcopy command  
--postcopy-after-precopy  automatically switch to post-copy migration after one pass of pre-copy
```

Industry Edge: OpenShift Virtualization Real Time

什么是 RT？



装有防抱死刹车系统的宝马 R1200RT

如果一个系统作业的正确性不仅取决于其逻辑的正确性，还与其执行的时间有关，此系统就被称为实时系统。

防抱死系统是个实时运算系统的简单例子，在这个系统中的实时约束是为了避免车轮锁死，刹车必须释放的时间。实时电脑系统的反应最后期限一般和事件有关，若没能在最后期限前完成，即为实时电脑系统的失效。不论系统负载如何，实时电脑系统均需满足最后期限的限制条件。



什么是 PLC？



三菱公司的Q系列大型PLC，以模块化设计来扩展

PLC(可编程逻辑控制器)是一种用于工业自动化的数字电子设备，它可以通过编程来控制 and 监控机械设备和生产过程，取代了传统的继电器控制系统。



在汽车制造过程中，PLC 控制机器人和自动化设备，确保各个生产环节(如焊接、涂装、和总装)的精确操作和实时响应，以实现高效的自动化生产。

软件定义的 PLC (Software Defined PLC, SDPLC)

OPENPLC

TO A MORE
OPEN FUTURE



EDITOR  **RUNTIME** 



OpenShift Container 的 RT 配置

```
apiVersion: v1
kind: Pod
metadata:
  name: PodName
  annotations:
    cpu-load-balancing.crio.io: "disable"
    irq-load-balancing.crio.io: "disable"
    cpu-quota.crio.io: "disable"
spec:
  runtimeClassName: performance-rt
  restartPolicy: Never
  containers:
    .....
    resources:
      requests:
        memory: "8Gi"
        cpu: "5"
      limits:
        memory: "8Gi"
        cpu: "5"
```

```
spec:
  containers:
    .....
    ports:
      - containerPort: 8080
      - containerPort: 502
      - containerPort: 22
      command: ["sleep", "infinity"]
      securityContext:
        privileged: true
    .....
  nodeSelector:
    kubernetes.io/hostname: NodeName
  tolerations:
    - effect: NoSchedule
      key: dedicated
      operator: Equal
      value: realtime
```

OpenShift Virtualization 的 RT 配置



Packages

```
# kernel-rt
# rt-tests(rhel8) or
  realtime-tests(rhel9)
# tuned
#
tuned-profiles-realtime
# tuned-profiles-nfv
```



CPU/Tuned

```
# echo "isolated_cores=2-4" >>
/etc/tuned/realtime-virtual-guest-variabl
es.conf
# echo "isolate_managed_irq=Y" >>
/etc/tuned/realtime-virtual-guest-variabl
es.conf
# tuned-adm profile
realtime-virtual-guest
# systemctl reboot
```



Hugepage/iommu


```
# grubby --args="iommu=pt
intel_iommu=on
default_hugepagesz=1G
idle=poll"
--update-kernel=`grubby
--default-kernel`
# systemctl reboot
```




Swap/irqbalance

```
# swapoff -a(Run after each
reboot))
# systemctl disable
irqbalance
# systemctl stop irqbalance
```

OpenShift Virtualization 的 RT 配置



DS18B20 sensor & led connect to Arduino boards



PLC program makes the led warning




Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 [linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)

 [facebook.com/redhatinc](https://www.facebook.com/redhatinc)

 [youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)

 twitter.com/RedHat



Reference

[https://www.google.com.hk/search?newwindow=1&sa=N&sca_esv=45a32e260ac88b18&cs=0&q=ope\[...\]sg=AI4_-kS_JYQvc5UVkg0EBrz03ecZ0PDJ7Q&biw=1920&bih=961&dpr=1](https://www.google.com.hk/search?newwindow=1&sa=N&sca_esv=45a32e260ac88b18&cs=0&q=ope[...]sg=AI4_-kS_JYQvc5UVkg0EBrz03ecZ0PDJ7Q&biw=1920&bih=961&dpr=1)

https://developers.redhat.com/blog/2015/03/24/live-migrating-gemu-kvm-virtual-machines#table_of_contents

[https://archive.fosdem.org/2016/schedule/event/live_migration/attachments/slides/1251/\[...\]live_migration/slides/1251/live_migration_from_bottom_up.pdf](https://archive.fosdem.org/2016/schedule/event/live_migration/attachments/slides/1251/[...]live_migration/slides/1251/live_migration_from_bottom_up.pdf)

<https://www.slideshare.net/slideshow/rhevm-live-storage-migration/61719661>

<https://www.slideshare.net/slideshow/openshift-virtualization-technical-overviewpdf/253204784>

[https://access.redhat.com/documentation/vi-vn/red_hat_openshift_data_foundation/4.9/h\[...\]_deployment/introduction-to-openshift-data-foundation-rhodf](https://access.redhat.com/documentation/vi-vn/red_hat_openshift_data_foundation/4.9/h[...]_deployment/introduction-to-openshift-data-foundation-rhodf)

[https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/virtual\[...\]_guest_agent-running_the_gemu_guest_agent_on_a_windows_guest](https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/virtual[...]_guest_agent-running_the_gemu_guest_agent_on_a_windows_guest)

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/9.4_release_notes/new-features#new-features-virtualization