QEMU-KVM Upgrade Test

Stable Guest ABI / In Place Upgrade

Min Deng（邓 敏)
mdeng@redhat.com
September, 2022

@twitterhandle

About me

- I’m Min Deng
- I’m QE from KVM-QE team Red Hat
- I’m responsible for Stable Guest ABI and In Place Upgrade and some other features’ tests on x86 and ppc

Agenda

- Stable Guest ABI
  - Machine type
  - PC and Q35 on x86
  - Seabios and OVMF on x86
  - Upper Layer Products and Stable Guest ABI test
  - Test workflow
- In place upgrade
  - What's in place upgrade
  - Upgrade paths
  - Test

QEMU-KVM Upgrade Test

- **Stable Guest ABI** allows virtual machines to be presented with the same ABI across QEMU upgrade.
  - Regarding it as sub feature of migration (test point of view)
- Why need the test ?
  - Avoiding breaking down virtual machines
    - Apply critical bug fix, security mitigation
    - Support new features and new capabilities of existing features

QEMU-KVM Machine Type

- Machine Type
  - Emulate different chipsets and related devices
  - Provide Stable Guest ABI
- Know machine type on different architectures
  - Check it on qemu-kvm by
    - /usr/libexec/qemu-kvm -M ?
  - You can refer to source code if you need

Machine type on different architectures

- x86_64
  - The **pc** machine type
  - The **q35** machine type
- ppc64le
  - The **pseries** machine type
- s390x
  - The **s390-ccw-virtio** machine type
- aarch64
  - The **virt** machine type

QEMU-KVM Machine Type

- PC and Q35
  - **PC**
    - QEMU corresponds to Intel® i440FX chipset (released in 1996)
      - pc RHEL 7.6.0 PC (i440FX + PIIX, 1996) (alias of pc-i440fx-rhel7.6.0)
  - **Q35**
    - QEMU corresponds to Intel® 82Q35 chipset (released in 2007)
      - Supported modern features
      - q35 RHEL-8.6.0 PC (Q35 + ICH9, 2009) (alias of pc-q35-rhel8.6.0)

QEMU-KVM Machine Type Q35 chipset Overview

- Two primary components:
  - Graphic Memory Controller Hub
  - IO Controller Hub (ICH9/ICH9 DO)
- New features:
  - PCIe
  - AHCI storage controller
  - vIOMMU emulation
  - “Secure” Secure Boot
  - …

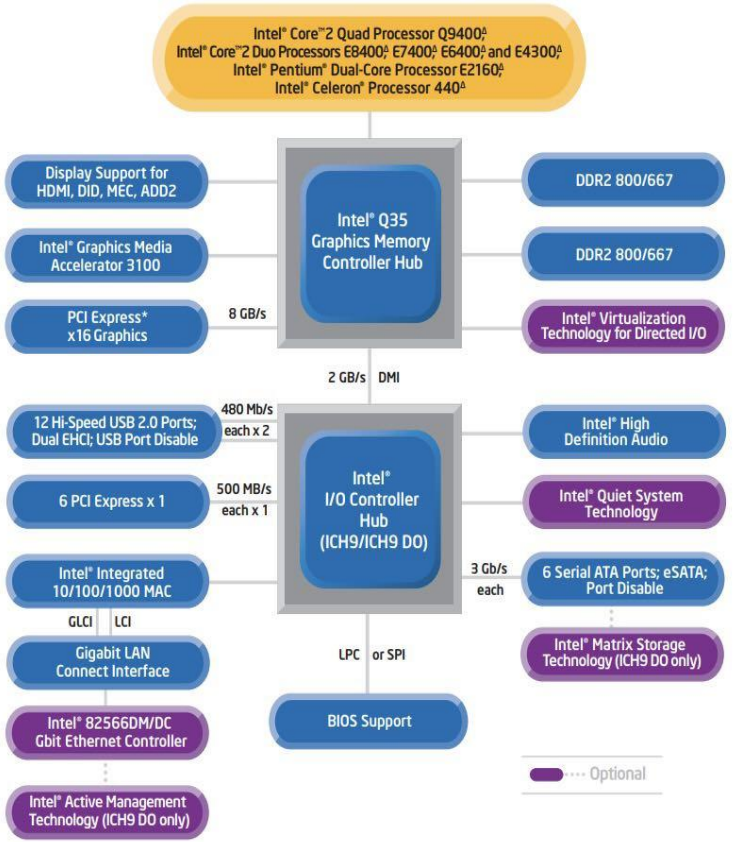

Block Diagram for Intel® Q35 Express Chipset

https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/q35-chipset-brief.pdf

QEMU-KVM Machine Type
Seabios and OVMF(Open Virtual Machine Firmware)

- Seabios
  - SeaBIOS runs inside an emulator, it's the default BIOS for the QEMU-KVM
- OVMF
  - UEFI(Unified Extensible Firmware Interface) for x86 VMs is called OVMF
- Test Matrix of Stable Guest ABI on x86
  - PC and Seabios
  - Q35 and Seabios
  - Q35 and OVMF

Upper Layer Products And Stable Guest ABI

- **Red Hat OpenStack Platform**
  - Red Hat OpenStack Platform-16.1
  - Red Hat OpenStack Platform-16.2
  - …
- **Red Hat OpenShift**
  - Container-native virtualization 4.8
  - Container-native virtualization 4.9
  - Container-native virtualization 4.10
  - …
- **Red Hat Virtualization**
  - RedHat Virtualization 4.4.8
  - RedHat Virtualization 4.4.9
  - RedHat Virtualization 4.4.10
  - …

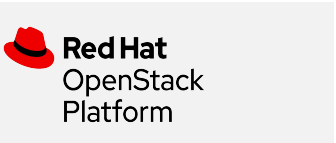

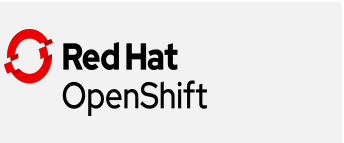

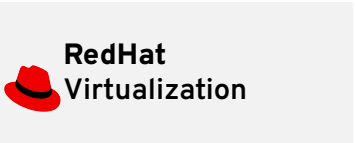

Upper Layer Products And Stable Guest ABI

- Red Hat OpenShift virtualization support for mixed applications running on virtual machines (“VMs”) and containers. Previously known as container-native virtualization (“**CNV**”)
- OpenShift virtualization is a feature of the OpenShift platform

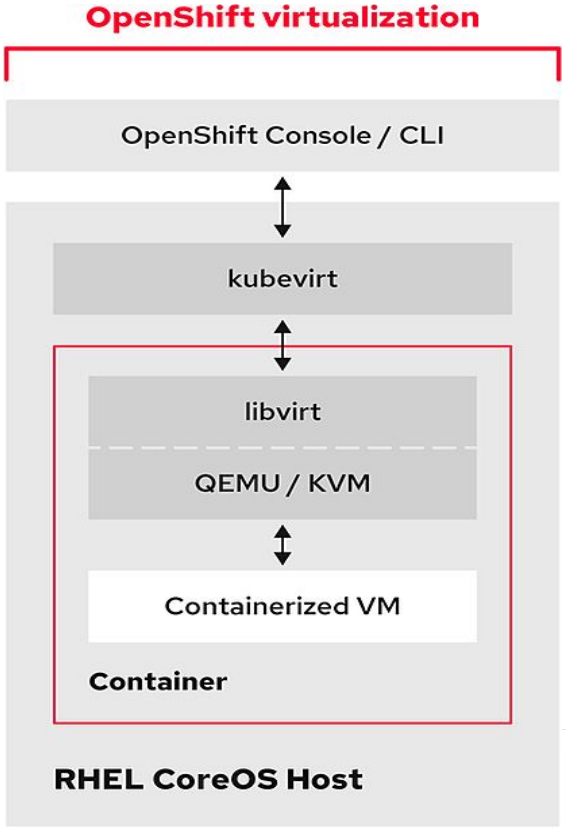

Four aspects

- Product line
  - RHEL.7 to RHEL.8
  - RHEL.8 to RHEL.9
- Cover supported features
  - New supported features
  - New capabilities of existing features
- Versioned machine type on different architectures
  - x86_64
  - ppc64le
  - S390x
- Hardware
  - x86_64
    - Intel, AMD
  - ppc64le
    - Power 8, Power 9
  - s390x
    - IBM z

Test Principles

- **Test Principles**
  - Ping-Pong migration
    - Live migration
    - Post copy
  - Test intersection of machine types
  - Consider priority of Seabios and OVMF for VM on different product lines. (x86_64)

Test Workflow

- x86_64
- ppc64le
- s390x

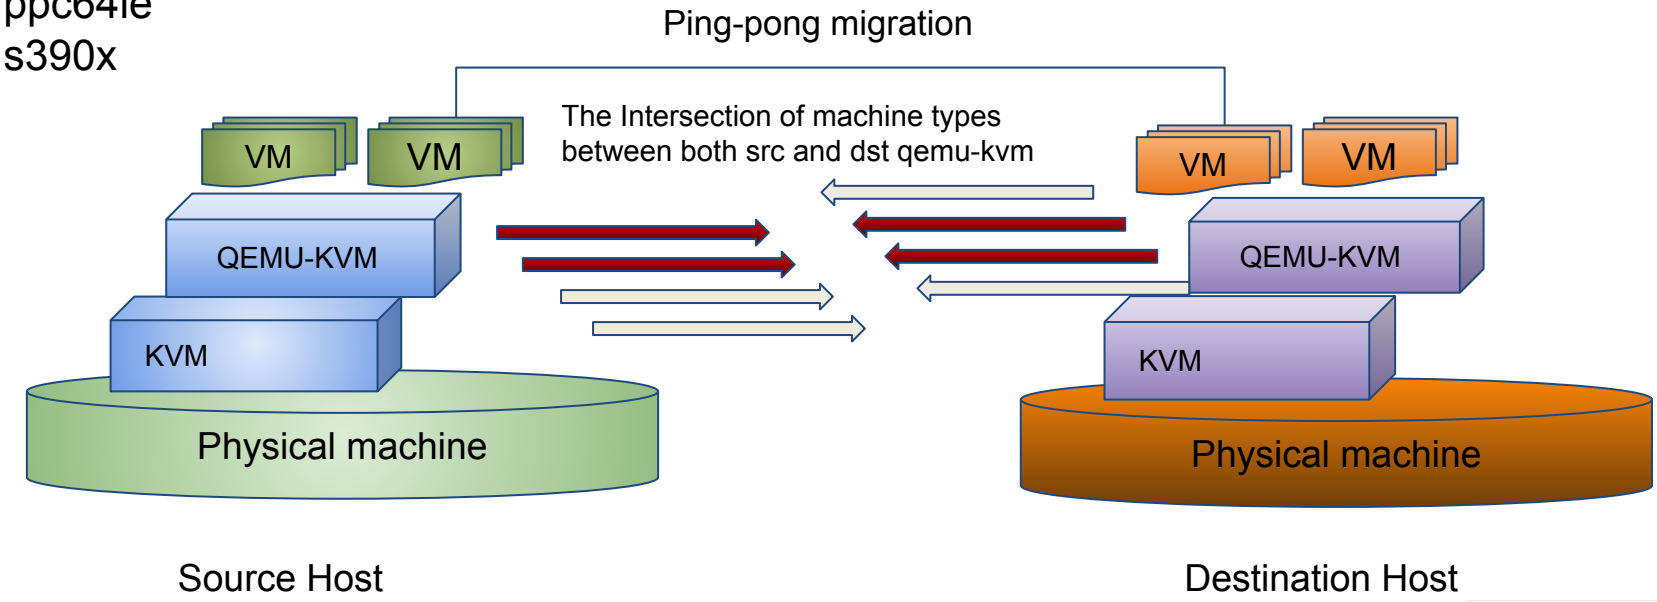

In Place Upgrade

- **What's In Place Upgrade ?**
  - In place upgrade(IPU) is a way of upgrading a system to a new major release of Red Hat Enterprise Linux by replacing the existing operating system.
  - The in place upgrade tool is leapp utility
- **QEMU-KVM related test on x86_64, ppc64le and s390x**
  - IPU on the VM
  - IPU on the host

Advantages

In place upgrade vs re-deployment

01	Preserve configuration	• Old configuration will be removed and need to set up new configuration again
02	Retain subscription management	• Machines have to be re-subscribed
03	Save time and cost	• Additional time and cost
04	Low bar of seniority required	• Require expertise to ensure the setup

Upgrade paths

In place upgrade from RHEL 7 to RHEL 8	
HOST	RHEL 7.(minor release) -> 8.(even-numbered minor release)
VIrtual Machine	
In place upgrade from RHEL 8 to RHEL 9	
HOST	RHEL 8.(minor release) -> RHEL 9.(even-numbered minor release)
VIrtual Machine	

Implementation

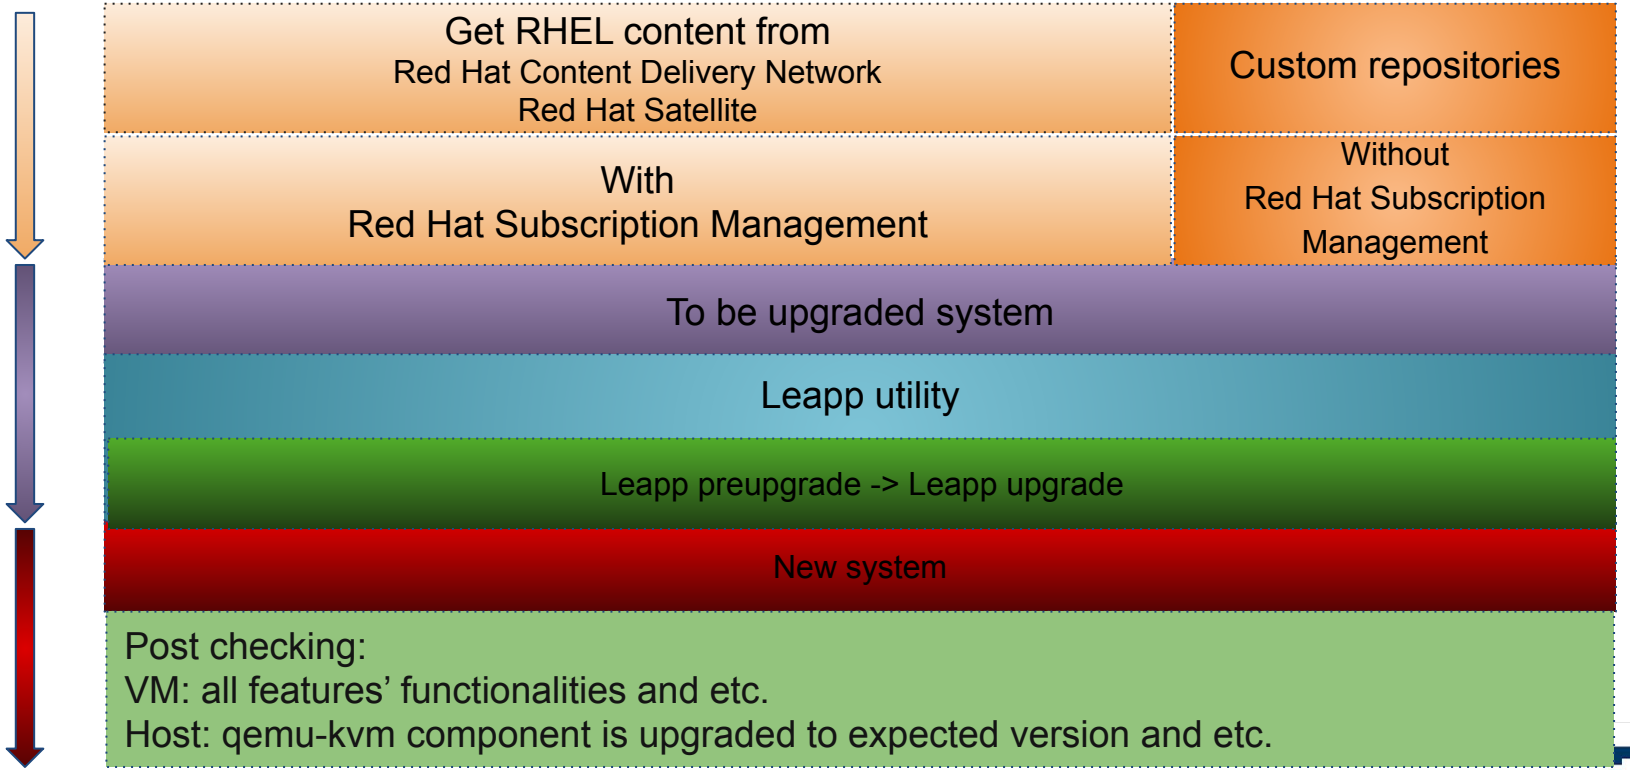

In Place Upgrade With/Without RHSM

- **With RHSM**
  - Red Hat Subscription Management (RHSM)
    - RHSM is the service which manages your Red Hat subscriptions and entitlements
- **Without RHSM**
  - Custom repos should be provided at the beginning of In Place Upgrade

In Place Upgrade with RHSM

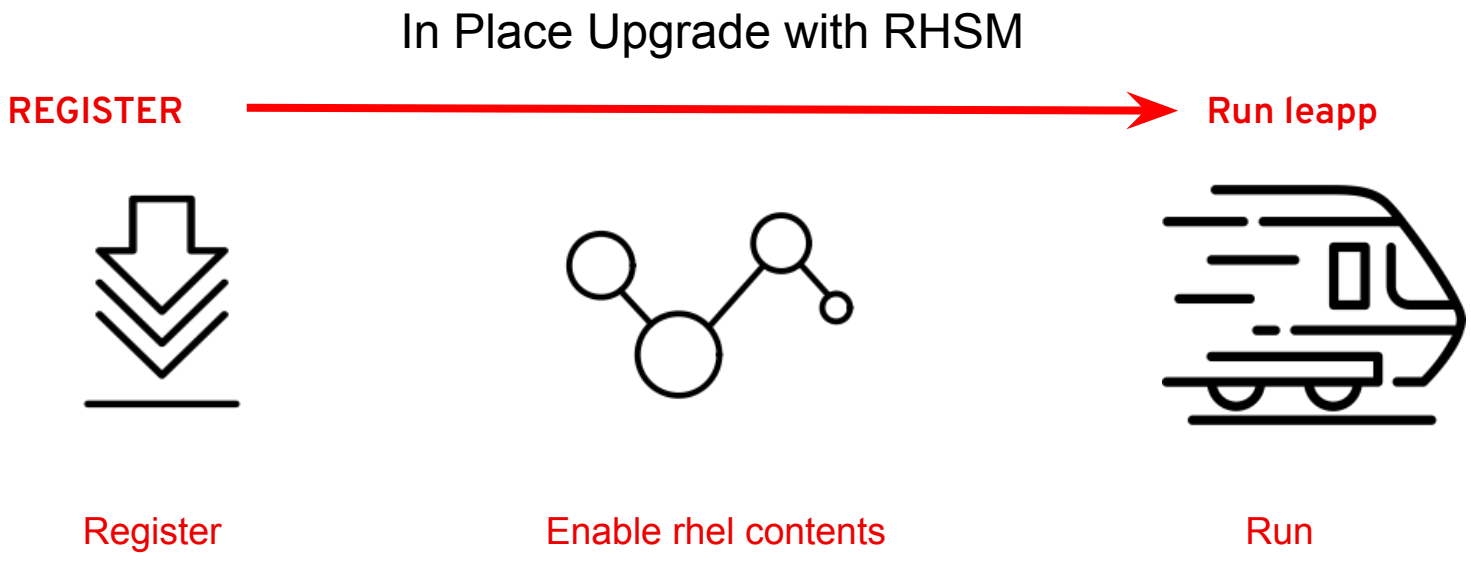

In Place Upgrade with RHSM

Register your system

- #subscription-manager config --rhsm.baseurl=https://xxx.redhat.com
  - configure your server first
- #subscription-manager register --username in_place_upgrade --password in_place_upgrade --serverurl "subscription.xxxx.redhat.com"
  - register your system by the user attached the SKU already

In Place Upgrade with RHSM

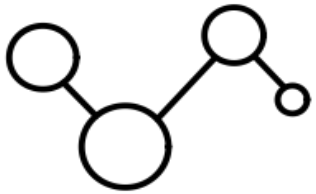

Enable RHEL content

- #subscription-manager list --available
  - According to user's account information and then you can get a pool id
- #subscription-manager attach --pool poolid
  - Attach your old system to above pool where you can get the product(rhel) content for upgrading your system later.
- #subscription-manager list --installed
  - Check if you have the Red Hat Enterprise Linux Server subscription attached
- #subscription-manager repos --enable rhel-7-server-extras-rpms
- #subscription-manager repos --enable rhel-7-server-rpms
  - Enable the base repository and enable the extras repository where leapp and its dependencies are available
  - It's RHEL 7 repos here, and you need to adjust repos according to your current upgrade path
- #yum update
  - Update old system to the corresponding minor version
- #reboot
  - Reboot your old system if required

In Place Upgrade with RHSM

Run

- #yum install leapp-upgrade
  - Install leapp tool
- #leapp preupgrade
  - To assess upgradability of your system, start the pre-upgrade process by the leapp preupgrade command
- #leapp upgrade (eg. --target 8.6/9.0)
  - Leapp takes over the role to upgrade your system
- … less one hour, just need to wait !
  - New system will be ready soon … :)

In Place Upgrade with RHSM

Results

- ----------Finish Upgrade ----------
- Verify that the current OS version is Red Hat Enterprise Linux X:
  - #cat /etc/redhat-release
  - #uname -r
- Verify that the correct product is installed
  - #subscription-manager list --installed
  - #subscription-manager release

Reference

[Reference]

https://github.com/qemu/qemu/blob/master/docs/pcie.txt

https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/q35-chipset-brief.pdf

https://wiki.qemu.org/Features/Q35

Q&A

KVM FORUM