

# SUSE® SLES® 12 KVM

Taking Advantage of the Latest Features of KVM Virtualization

**Bruce Rogers**

Software Engineer

SUSE / brogers@suse.com



# SLES 12 QEMU / KVM Highlights

- General Status of QEMU/KVM:
  - KVM (v3.12 kernel) and QEMU v2.0
  - Available on Intel 64 / AMD 64, IBM System z and POWER 8 architectures!
- Libvirt preferred management layer ...
  - QEMU command line supported for access to more features
- virt-v2v tool for converting Xen VM to KVM
- VMDP v2.2 aligns with SLES 12 release
  - Windows Guest Agent
  - Performance improvements

# QEMU / KVM Documentation

Where do I go to learn more?

- Online Documentation: <https://www.suse.com/documentation/sles-12>
  - Product Release Notes (for technical limits and technical preview features)
  - Virtualization Guide
  - Xen to KVM Migration Guide
- White paper
- Best practices guide
- QEMU package documentation
  - From upstream: `/usr/share/doc/packages/qemu/qemu-doc.html`
  - SUSE specific: `/usr/share/doc/packages/qemu-kvm/kvm-supported.txt`
  - QEMU man page

# Some Changes to Be Aware Of...

- Command line default changes
  - Default ram size: 128MB (was 512MB)
  - Default emulated nic: e1000 (was rtl8139)
  - Default UI: GTK (was SDL)
- Packaging changes
  - kvm and virt-utils packages are gone
  - QEMU split into multiple packages (qemu-\*)
  - qemu-kvm package and binary are now considered legacy and optional
- PCI passthrough using VFIO is more strict than legacy: some legacy configs which “work”, not allowed under VFIO

# New Usability Features

- Spice graphics acceleration
- VFIO
- Hot-add vcpus
- Guest Agent for SLES 12 and Windows guests
- cache=none now possible on storage with non-512 byte buffer size / alignment requirements
- GTK interface - includes menu of common operations
- virtio-rng

# New Usability Features

- vTPM for passthrough of host TPM
- Kernel steal time accounting
- Panic device (pvpanic)
- Live migration auto-convergence and statistics
- qemu-img dump of storage metadata
- nocow attribute when creating storage images
- Built with support for Xen's usage of QEMU

# New Performance Related Features

- Spice graphics acceleration
- Guest install-time “best” vcpu matching
- Multi-queue tx networking
- ivshmem (Inter-VM shared memory)
- vhost-net zero copy
- virtio-blk “bio” based I/O for better performance with high speed storage devices
- CPU affinity support for MSI interrupts
- KVM improvements: x2apic, PV-EOI, instruction emulations

# New Technology Previews

- USB 3.0
- PCI bridge
- Q35 machine, including PCI Express
- virtio-ccw on s390



# Existing Feature Improvements

- 4 TB guest support (was 2 TB in SLES 11 SP3)
- Live migration from SLES 11 SP3 to SLES 12
- QCOW2 and QED images now migratable
- All vcpu types supported (except -cpu host)
- No reboot needed after installing kvm hypervisor
- KVM modules autoloaded by hardware detection
- vhost-net module autoloaded by /dev/vhost-net open
- More command line options and features now supported: megasas scsi adapter, more char devices, Hyper-V emulation, vcpu improvements, hotplug char devices

# QEMU v2.1 and Nested Virtualization Updates In the Queue for SLES 12

- NUMA topology passthrough finally added to QEMU
  - NUMA awareness is important for good performance on NUMA hosts
  - Poor NUMA affinity can result in as much as ~10x guest memory performance hit depending on memory access hops
  - Previously binding guest to single NUMA node was only real solution
  - v2.1 adds ability to specify guest NUMA bindings and policies
- Hotplug memory
- Improved USB 3.0 support
- Q35 machine PCI Express hotplug / unplug
- More complete nested virtualization support



**QEMU / KVM Keeps  
Getting Better So That  
You Can As Well!**

Thank you.





## **Unpublished Work of SUSE LLC. All Rights Reserved.**

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE LLC. Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

## **General Disclaimer**

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. SUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for SUSE products remains at the sole discretion of SUSE. Further, SUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All SUSE marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

