

London OpenStack Meetup, July 2012



libvirt & KVM with OpenStack Nova

Daniel P. Berrangé <berrange@redhat.com>

# libvirt: Why it exists

- Stable: isolation from HV API changes
- Standard: portable across HV
- Simple: rapid application development
- Portable: Linux, Solaris, Windows & OS-X
- Secure: TLS, SASL, SSH, PolicyKit
- Open: LGPLv2+ license



# libvirt: What virtualization?

- QEMU – KVM, QEMU
- VMWare – ESX, GSX Server, vCenter
- Xen – XenD, XenStored, Xen Hypervisor
- LXC – native Linux containers
- Test – 'mock' hypervisor
- Remote – RPC access to APIs
- ...VirtualBox, OpenVZ, Phyp, UML, XenAPI, Hyper-V, Parallels



# KVM: What is it ?

- Hardware virtualization (VT/SVM)
- Linux kernel as a hypervisor
- QEMU userspace device emulation
- Paravirtualized device drivers (VirtIO)



# libvirt: What KVM benefits...?

- CPUID: verified migration compatibility
- Migration: secure tunnelling
- Guest ABI: stable PCI addrs & hardware
- CLI: comply w/ QEMU best practice
- Monitor: comply w/ QEMU best practice
- Save/restore: compression, O\_DIRECT



# libvirt: ...KVM resource benefits?

- CGroups: CPU, memory, disk I/O limits
- Network: I/O limits, packet filtering
- Locking: disk lease management
- Encryption: qcow2 encryption key mgmt
- PCI: device assignment safety checks



# libvirt: ...KVM security benefits?

- DAC: unprivileged UID:GID
- MAC: sVirt with SELinux / AppArmor
- Capabilities: block setuid usage
- CGroups: device ACLs
- Audit: logging of operations
- Certification: common criteria



# Nova: What is it ?

- Cloud compute service / controller
- Native API + EC2 compatible API
- Multiple hypervisor drivers
- Reliable message bus
- Integrates Quantum, Cinder, Keystone





# Nova: libvirt driver

- Uses local hypervisor APIs
- KVM, Xen (Paravirt), LXC & UML
- Python API binding
- Flexible disk configuration
- Network bridging / OpenVSwitch
- Guest network traffic filtering



# Nova: libvirt booting

- LXC: /sbin/init
- UML: /usr/bin/linux
- Xen/KVM: kernel+initrd / HD / CDROM



# Nova: libvirt disks

- LXC
  - loop mount raw/qcow2 image as /
- KVM/Xen/UML
  - Primary disk (raw/qcow2)
  - Config disk (raw/qcow2)
  - Swap disk (raw/qcow2)
  - 'n' \* ephemeral disks (raw/qcow2)
  - Persistent volumes for any disk (raw)



# Nova: libvirt in Folsom

- Qual: No more libvirt.xml.template
- Perf: Optimized timers
- Perf: Configurable CPU model
- Flex: LVM for ephemeral disks
- Maint: Diagnostic stats
- Maint: Avoid /tmp for snapshots
- Flex: Xen HVM (expected)



# Nova: libvirt in Grizzly

- Flex: libosinfo
  - Identify ISO image OS type
  - Optimize OS hardware config (Win32!)
- Sec: QCow2 encryption
- Maint: Guest watchdog
- Perf: NUMA placement
- Perf: resource isolation





<http://libvirt.org> ~ <http://linux-kvm.org> ~ <http://openstack.org>