

KVM Forum Vancouver, 2011



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: Who contributes

- Founder: libvirt: Red Hat, libvirt-cim: IBM
- Corporate: IBM, Fujitsu, Siemens, Novell, Canonical, Sun, Citrix, HP, Dell, ...and more...
- Community: Fedora, Debian, Gentoo, Suse, ...many unaffiliated developers...



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, (soon Hyper-V)



libvirt: What platforms?

- Linux: RPC client+server, all hypervisor drivers
- OS-X: RPC client
- Windows: RPC client, ESX+VirtualBox (Hyper-V)
- Solaris: RPC client, Xen driver

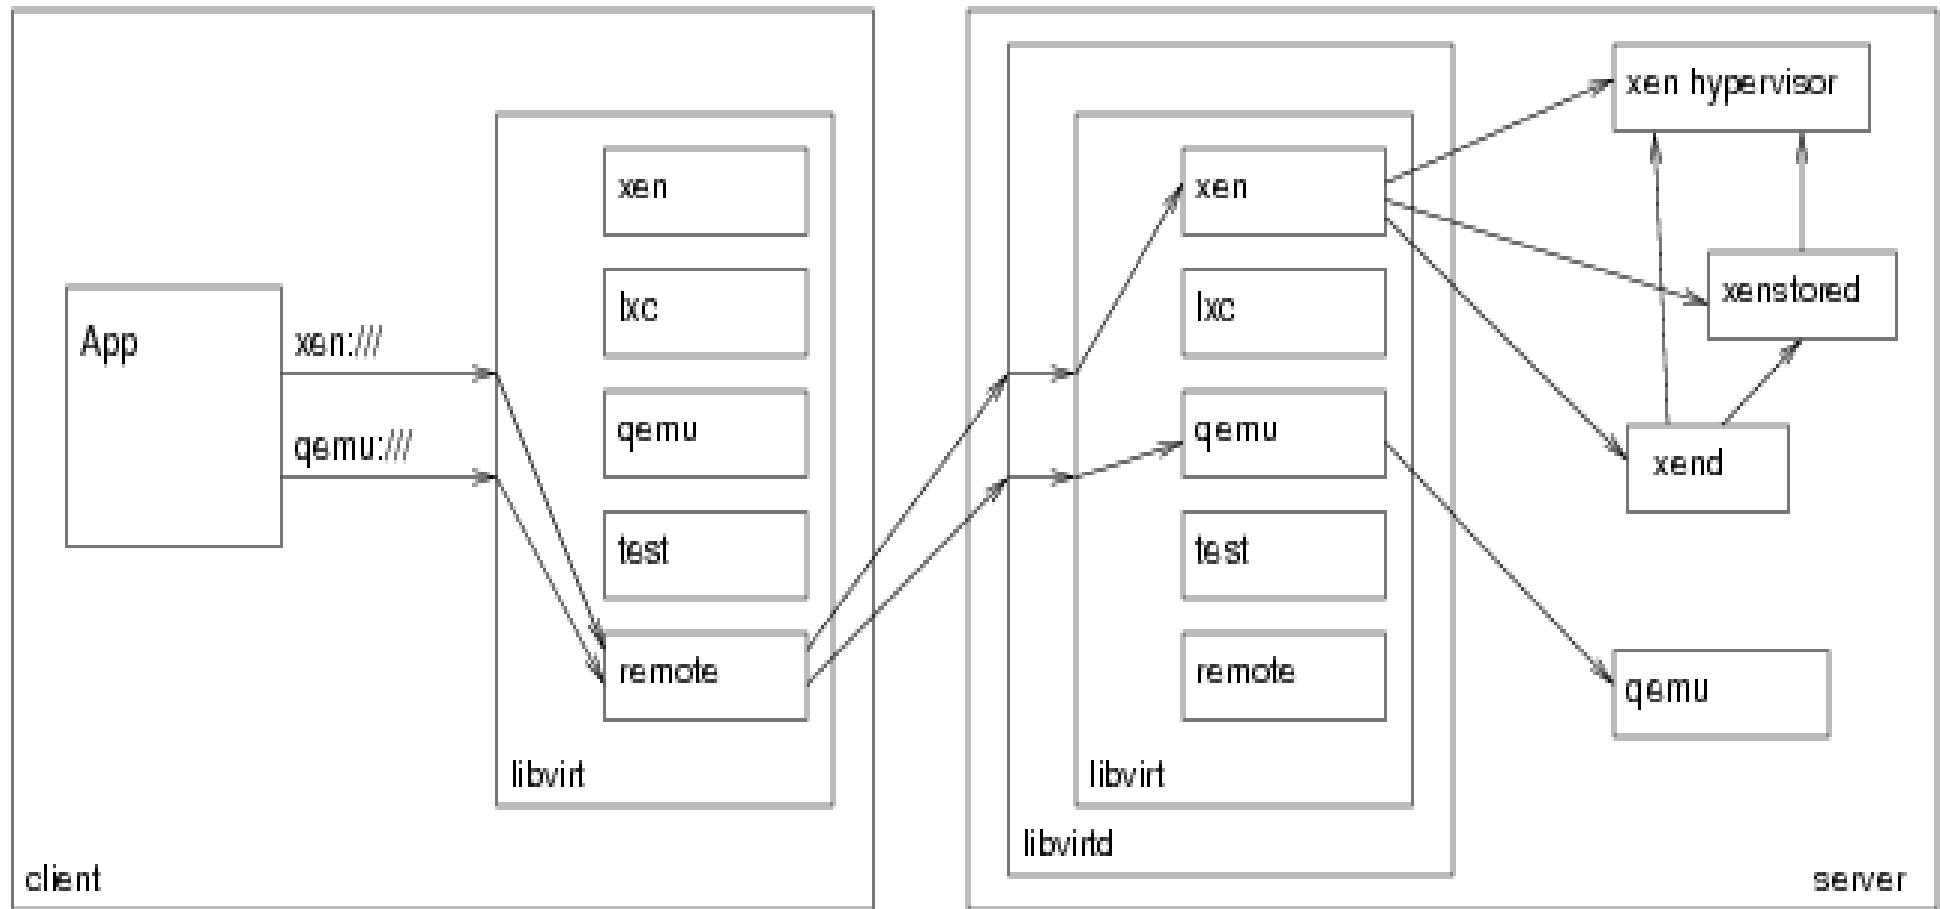


libvirt: What else?

- Storage: LVM, Disk, SCSI, iSCSI, NFS
- Network: bridge, bonding, vlan, VEPA
- Firewall: filtering ebtables, ip(6)tables
- Devices: HAL, udev
- Security: SELinux, AppArmor



libvirt: How it works



libvirt: How to use it

- Core: C library
- Bindings: Perl, Python, OCaml, Java, Ruby, C#, Php
- Mapping to CIM/DMTF: libvirt-cim
- Mapping to AMQP/QMF: libvirt-qmf
- Mapping to SNMP: libvirt-snmp
- Mapping to GObject: libvirt-glib



libvirt: Where it is used

- virt-manager: desktop manager app
- virt-inst: install, clone, deploy appliance
- cobbler: network installation service
- rhn: red hat network management
- RHEV-M/VDSM: data center management
- gnome-shell: desktop integration
- virt-dmesg: guest kernel dmesg



libvirt: KVM general 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 or AppArmor
- Capabilities: block setuid usage
- CGroups: device ACLs
- Audit: logging of operations
- Certification: common criteria
- Containers: PID, FS, Net namespace (TBD)



libvirt: for KVM developers

- Custom ARGV/ENV:

```
<qemu:arg value='-somenewarg' />
```

```
<qemu:env name='F00' value='BAR' />
```

- Monitor passthrough:

```
virsh qemu-monitor-command -hmp "info reg..."
```

- Attach to external process:

```
virsh qemu-attach $PID
```



libvirt: Where next

- RBAC: fine grained access control
- DTrace: low overhead dynamic probing
- Scalability: cope with 1,000's VMs/host
- iSCSI: LUN creation
- libvirt-gconfig, libosinfo, libvirt-install



libguestfs: Why it exists

- Stable: long term API/ABI guarantee
- Standard: access for any disk format
- Simple: rapid application development
- Open: LGPLv2+ license



libguestfs: What features

- Disks: block, raw, qcow2, vmdk, etc
- Filesystems: ext2, 3, 4, iso9660, NTFS
- Actions: read/write/edit/upload/download
- Misc: LVM, dm-crypt, fstab
- OS: Any (if Linux supports the FS)



libguestfs: How to use it

- Core: C API
- Bindings: Perl, Python, OCaml, Ruby, Java
- Shell: guestfish
- OS: FUSE filesystem
- Desktop: guest filesystem browser



libguestfs: How to use it

- virt-rescue: fix broken guest OS
- virt-df: guest filesystem usage
- virt-cat/virt-ls: read file / list directory
- virt-resize: change virtual disk size
- virt-tar: download archive of files
- virt-win-reg: edit Windows registry
- virt-v2v: convert guest to new hypervisor



libvirt/libguestfs: How to get help

- Websites <http://libvirt.org> / <http://libguestfs.org>
<http://wiki.libvirt.org>
- Core API documentation
- libvirt Application Developer Guide
- libvirt-users / libguestfs mailing lists
- #virt on irc.oftc.net (libvirt, virt-manager, etc)
- #libguestfs on irc.freenode.net





<http://libvirt.org/>