



What is new in Libvirt?

Michal Privozník
mprivozn@redhat.com

LinuxCon Japan 2013

What for?

security Guest is isolated from host

reliability Guest is host independent

recovery Guest can be saved at any point

testing One can run many different environments

shareability Guest can be shared among multiple users






- ▶ Stable: protection from hypervisor API changes
- ▶ Portable: Linux, Windows, OS-X, . . .
- ▶ Secure: TLS, SASL, (lib-)SSH, PolicyKit
- ▶ Simple: rapid application development
- ▶ Open: LGPLv2+



Basic structure



- ▶ Supported HV: QEMU, KVM, ESX, Xen (XenD, XenStored, Xen Hypervisor, Xenlight), LXC, VirtualBox, **Parallels**, UML, Hyper-V, OpenVZ, Phyp, VMWare, test



What else?

- ▶ Storage: LVM, disk, SCSI, iSCSI, NFS
- ▶ Network: bridge, bonding, vlan, VEPA, OpenVSwitch
- ▶ QoS, filtering (iptables + ebttables)
- ▶ CGroups: CPU, memory, disk I/O limits
- ▶ PCI/USB device passthrough
- ▶ Guest agent
- ▶ Host management



Language bindings:

- ▶ Core: C
- ▶ Perl, Python, Java, Ocaml, Ruby, C#, PHP

Mappings to different models:

- ▶ CIM/DMTF: libvirt-cim
- ▶ AMQP/QMF: libvirt-qmf
- ▶ SNMP: libvirt-snmp
- ▶ GObject: libvirt-glib



```
zippy@bart ~ $ virsh -c qemu:///system
Welcome to virsh, the virtualization interactive terminal.
```

```
Type: 'help' for help with commands
      'quit' to quit
```

```
virsh # list --all
```

Id	Name	State
1	fl7	running
-	fl6	shut off
-	fl6_nfs	shut off
-	win7	shut off

```
virsh # start fl6
Domain fl6 started
```

```
virsh # █
```

CLI for managing libvirt



- ▶ Nearly every API is exposed
- ▶ Ported to many platforms: Linux, OS-X, Windows, Solaris, ...
- ▶ Capable of talking to some HVs directly
- ▶ Remote access
- ▶ Stable output ⇒ easily usable in scripts



libvirt-1.0.0

tag v1.0.0

Tagger: Daniel Veillard <veillard@redhat.com>

TaggerDate: Fri Nov 2 12:09:46 2012 +0800

Release of libvirt-1.0.0

Yay !!!

7th birthday of the project see the first commit:

d77e1a9642fe1efe9aa5f737a640354c27d04e02 initial revision

-----BEGIN PGP SIGNATURE-----

Version: GnuPG v1.4.12 (GNU/Linux)

iEYEABECAAYFAlCTR88ACgkQRga4pd6VvB9ImgCgl0U348QSwRu3j79Nerc

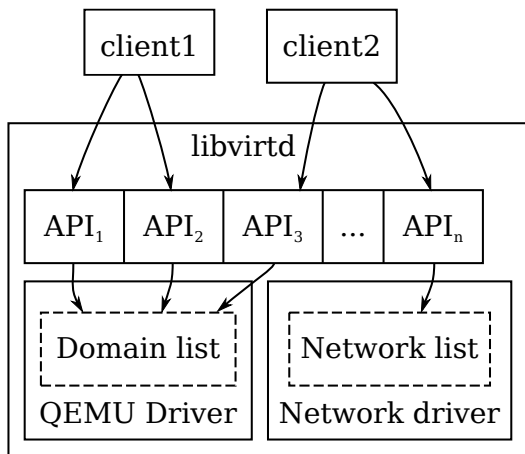
FOEAoI4wAVVB3rt2EOzqrE2a3JMz5pZ3

=d/6i

-----END PGP SIGNATURE-----



The Big Libvirt Lock



The driver's list is bottleneck. Dropping it speeds up domain starting process by 200%.



- disk** Content of disk is saved at given point of time
- memory** Tracks state of RAM and other resources
- checkpoint** Combination of disk + memory

Small live example




The "sanlock" plugin for libvirt allows integration with the sanlock daemon for locking.

Explicit leases Management application must add leads in the guest XML to represent exclusion policy is desired. Leases are not directly associated with any device.

Automatic leases libvirt creates leases for each disk device, based on the filepath of the disk backing store. Assumes that the app is using stable file paths, across all hosts.




QEMU Guest Agent



- ▶ Execute operations from inside the guest (poweroff, FS freeze, thaw, trim, ...)
- ▶ Side channel to standard qemu monitor



Multiqueue Network



- ▶ Packet sending/receiving process scale with number of vCPUs
- ▶ Create multiple TX/RX queues, each can be handled with a different vCPU
- ▶ Changes involve kernel, qemu & libvirt

<http://www.linux-kvm.org/wiki/images/e/e3/Ver1.jpg>



Sandbox on the top of LXC/KVM with libvirt

```
zippy@bart ~ $ virt-sandbox -p -c lxc:/// /bin/date  
Mon Apr 30 21:21:46 CEST 2012  
zippy@bart ~ $ █
```

- ▶ Host filesystem passthrough bind mounts or P9FS
- ▶ small overheads: LXC 200ms, KVM 3000ms
- ▶ Boot kernel+initrd (KVM) “init” binary (LXC)

Usage scenarios: run apache per virtual host inside sandbox, mock RPM build



Produce domain XML for libvirt

- ▶ Reimplementation of virt-install
- ▶ Based on libosinfo¹ and libvirt-glib²
- ▶ To cooperate with libvirt-builder (not started yet)

¹<http://libosinfo.org/>

²<ftp://libvirt.org/libvirt/glib/>



Libvirt participate in GSoC:

- ▶ More intelligent virsh auto completion
- ▶ Libvirt RPC protocol Wireshark dissector
- ▶ Introduce API to query IP addresses for given domain



Where to go?

- ▶ Libvirt <http://libvirt.org>
- ▶ libvirt-snmp
<http://wiki.libvirt.org/page/Libvirt-snmp>
- ▶ libvirt-sandbox <http://berrange.com/tags/sandbox/>

GNOME Shell - not standalone mgmt applications, but virtualization builtin desktop (gnome-boxes)



The End

Questions?

