

libguestfs

tools for viewing and modifying virtual machine disk images

Richard W.M. Jones

Software Engineer
Red Hat

`rjones@redhat.com`

November 2012

1 Introduction

The libguestfs project is nearly 4 years old, so for my rare opportunity to talk at the KVM Forum I thought I'd cover where we've got to in those 4 years, and also talk about some of the ways we are using the new features in KVM.

2 Overview of libguestfs

libguestfs is a C library with a simple, long-term stable API that contains about 400 different calls. It also has bindings in many popular high-level languages, including Perl, Python, Ruby, OCaml, Java, PHP, and more. The library lets you write simple programs to create, view and modify disk images, like the one in Appendix A.

There are many tools which use libguestfs:

guestfish Interactive and scriptable shell.

guestmount Mount filesystems from any guest or disk image on the host.

virt-alignment-scan Check alignment of partitions in guests.

virt-cat Display a file from a guest.

virt-copy-in Copy files and directories into a guest.

virt-copy-out Copy files and directories out of a guest.

virt-df Display disk usage of a guest.

virt-edit Edit a file in a guest.

virt-filesystems Display the partitions, filesystems, logical volumes etc. in a guest.

virt-inspector Inspect a guest and produce a report detailing the operating system, version, applications installed and more.

virt-ls List files and directories in a guest.

virt-make-fs Make a new filesystem.

virt-rescue Rescue mode for guests.

virt-resize Resize a guest.

virt-sparsify Make a disk sparse.

virt-sysprep Reset a guest to “factory configuration”.

virt-tar-in Copy files from a tarball into a guest.

virt-tar-out Copy files out of a guest into a tarball.

virt-win-reg Display and edit the Windows Registry in a guest.

and libguestfs is increasingly being used by major projects:

BoxGrinder & Oz Projects that can create guests from scratch.

OpenStack Uses libguestfs to inject files into guests.

virt-manager Uses libguestfs to display icons and applications in guests.

virt-v2v & virt-p2v Migrate between physical machines and different hypervisors.

3 Using guestfish to view and inject files

An easy way to get started with libguestfs is to use our shell scripting tool, *guestfish*¹ to open up a disk image and look inside.

You can open up almost any type of disk image, such as a raw file, qcow2 file, or a logical volume or partition that contains a virtual machine. You *don't* need to be root, unless you need root to access the device.

Simply do:

```
$ guestfish -a disk.img -i
```

```
Welcome to guestfish, the libguestfs filesystem interactive shell for
editing virtual machine filesystems.
```

```
Type: 'help' for help on commands
      'man' to read the manual
      'quit' to quit the shell
```

```
Operating system: Red Hat Enterprise Linux Server release 6.3 (Santiago)
/dev/vg_rhel6x32/lv_root mounted on /
/dev/sda1 mounted on /boot
```

You can now use guestfish commands to list files and directories and create files:

```
><fs> ls /
bin
boot
cgroup
[etc]
```

```
><fs> write-append /etc/rc.d/rc.local "service sshd start"
```

```
><fs> cat /etc/rc.d/rc.local
#!/bin/sh
touch /var/lock/subsys/local
service sshd start
```

Some useful commands include:

¹<http://libguestfs.org/guestfish.1.html>

`cat` Display small text files.

`edit` Edit a file.

`less` Display longer files.

`ll` List (long) directory.

`ls` List directory.

`mkdir` Make a directory.

`rm` Remove a file.

`touch` Touch a file.

`upload` Upload a local file to the disk.

`write` Create a file with content.

As there are hundreds of commands, it's a good idea to read the *guestfish(1)* and *guestfs(3)*² man pages. These are also available online at the website³.

4 Introducing virt-rescue

`virt-rescue`⁴ is a good way to rescue virtual machines that don't boot, or just generally make ad hoc changes to virtual machines. It's like a rescue CD for virtual machines.

`virt-rescue` is a little different from `guestfish` in that you get an ordinary shell and ordinary tools. However unlike `guestfish`, `virt-rescue` cannot be used from shell scripts, so it's not useful if you want to make repeatable changes to lots of your guests.

You must not use `virt-rescue` on running VMs.

²<http://libguestfs.org/guestfs.3.html>

³<http://libguestfs.org>

⁴<http://libguestfs.org/virt-rescue.1.html>

If you had a libvirt guest called “Fedora” then:

```
# virt-rescue -d Fedora
[lots of boot messages]
```

```
Welcome to virt-rescue, the libguestfs rescue shell.
```

Note: The contents of / are the rescue appliance.
You have to mount the guest’s partitions under /sysroot
before you can examine them.

```
><rescue> lvs
  LV      VG          Attr   LSize Origin Snap%   Move Log Copy%  Convert
  lv_root vg_f13x64 -wi-a- 7.56g
  lv_swap vg_f13x64 -wi-a- 1.94g
><rescue> mount /dev/vg_f13x64/lv_root /sysroot/
[ 107.912813] EXT4-fs (dm-0): mounted filesystem with ordered data mode.
Opts: (null)
><rescue> ls -l /sysroot/etc/fstab
-rw-r--r--. 1 root root 781 Sep 16  2010 /sysroot/etc/fstab
><rescue> vi /sysroot/etc/fstab
```

There is a lot more information about virt-rescue in the man page *virt-rescue(1)*.

5 New features in libguestfs 1.20

At the end of this year (2012) we hope to make a major new release of libguestfs. Since the last version was released in May 2012, this represents over 6 months of development effort, and it includes and uses some major new features from KVM.

The new version will (optionally) use libvirt to manage the libguestfs appliance. From the point of view of libguestfs this hides the complexity of dealing with the qemu command line. It also lets us leverage libvirt for:

security The appliance will now be encapsulated with sVirt (SELinux or AppArmor) to prevent a malicious guest from being able to take over the qemu process and escape to attack the host.

remote access libvirt will provide remote access to guests, using the ordinary libvirt remote URIs.

We are also using new features from KVM:

virtio-scsi The primary new feature is Paolo Bonzini's virtio-scsi driver, which is far superior in both performance and features compared to virtio-blk.

up to 255 disks With virtio-scsi, we can now access up to 255 disks at a time, and possibly many more in future.

hot-plugging of disks With virtio-scsi and libvirt together we are able to hotplug (add and remove) disks, giving libguestfs users a lot more performance and flexibility.

sparsification Coming soon, we'll be able to trim filesystems and cause the host backing file to become sparse.

qcow2 v3 We use the new version of qcow2 wherever possible because of its performance benefits.

A Example: how to inspect a virtual machine disk

```
#!/usr/bin/perl -w

use strict;
use Sys::Guestfs;

if (@ARGV < 1) {
    die "usage: inspect_vm disk.img"
}

my $disk = $ARGV[0];

my $g = new Sys::Guestfs ();
$g->add_drive_opts ($disk, readonly => 1);
$g->launch ();

# Ask libguestfs to inspect for operating systems.
my @roots = $g->inspect_os ();
if (@roots == 0) {
    die "inspect_vm: no operating systems found";
}

for my $root (@roots) {
    # Print basic information about the operating system.
    printf " Product name: %s\n",
        $g->inspect_get_product_name ($root);
    printf " Version:      %d.%d\n",
        $g->inspect_get_major_version ($root),
        $g->inspect_get_minor_version ($root);

    # Mount up the disks, like guestfish -i.
    my %mps = $g->inspect_get_mountpoints ($root);
    my @mps = sort { length $a <=> length $b } (keys %mps);
    for my $mp (@mps) {
        eval { $g->mount_ro ($mps{$mp}, $mp) };
    }
}
```

```
        if ($?) {
            print "$? (ignored)\n"
        }
    }

    # If /etc/issue.net file exists, print up to 3 lines.
    my $filename = "/etc/issue.net";
    if ($g->is_file ($filename)) {
        printf "--- %s ---\n", $filename;
        my @lines = $g->head_n (3, $filename);
        print "$_\n" foreach @lines;
    }

    # Unmount everything.
    $g->umount_all ()
}
```