Virtualization with libVirt and KVM

Kitchener-Waterloo Linux Users Group
November 2014

Khalid Baheyeldin
http://2bits.com

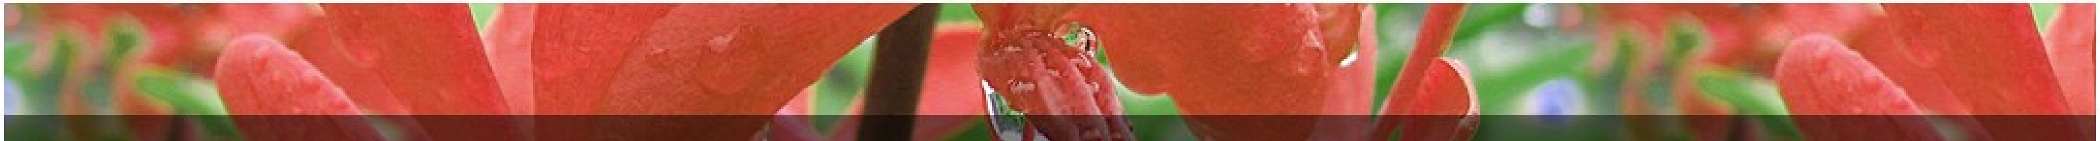

Agenda

- Goals
- What is Virtualization?
- Various virtualization technologies
- Why libvirt?
- Installation
- Creating VMs
- Managing VMs
- Questions/discussions

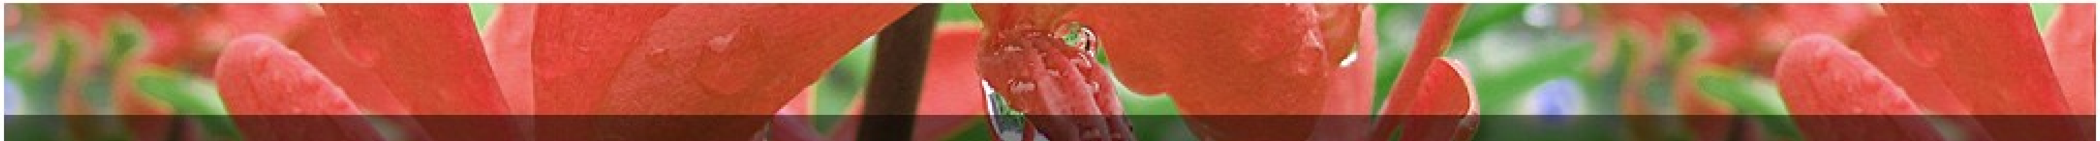

Goals

- Testing infrastructure
- Easy to setup, replicate and administer
- Different releases of the operating systems
  - Ubuntu Server 12.04, 14.04, ...
- Various combinations of software packages in different versions
  - MySQL, PHP, Apache, Drupal, ApacheSolr, ...

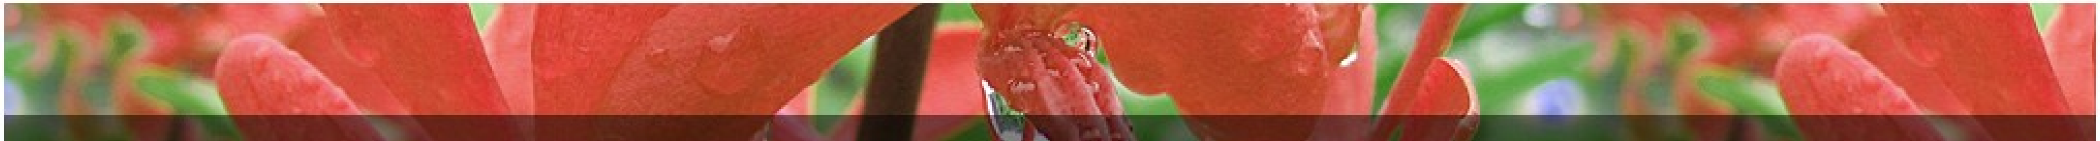

Virtualization

- Multiple instances of different operating systems on the same hardware
- IBM doing it on mainframes since the 1960s
- Started coming to the Intel architecture at the turn of the century (VMWare)
- Can run Windows guest on Linux host or vice versa
- Free variants started to come (Qemu, VirtualBox, Xen, then KVM)

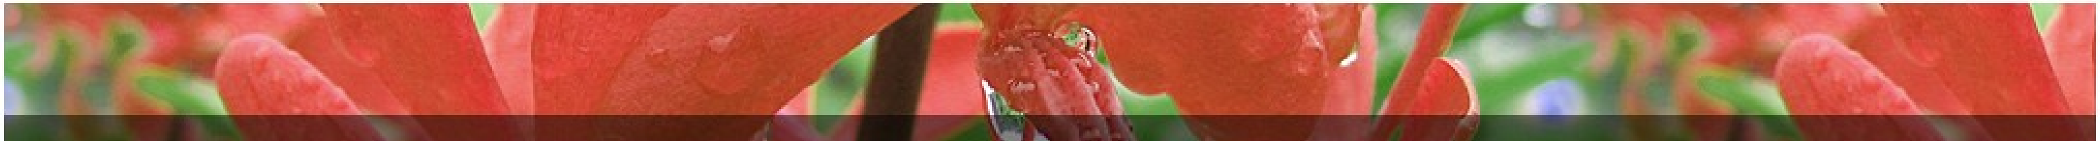

Hardware Support

- Technology to support virtualization at the CPU level
- Intel VT-x or AMD-V
- How to find out?
- `egrep "svm|vmx" /proc/cpuinfo > /dev/null ; echo $?`
  - 1 = NO virtualization support.
  - 0 = virtualization support.

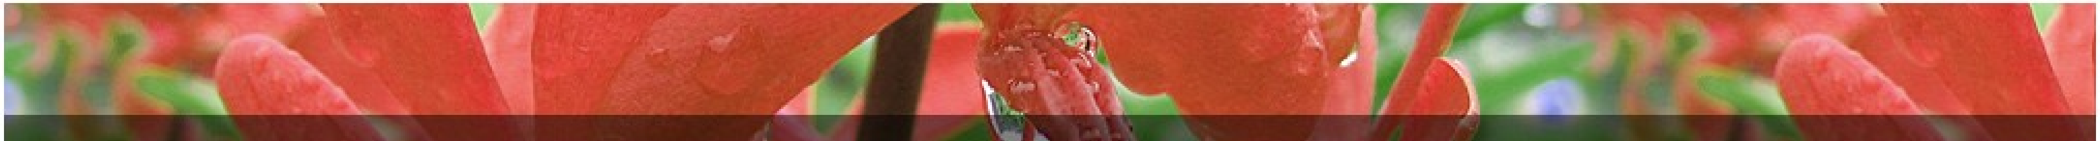

Qemu

- Fabrice Bellard
- More of an emulation than virtualization
- Useful for virtualizing FreeDOS and such lightweight operating systems
- Acceleration add-on
- Can be really slow

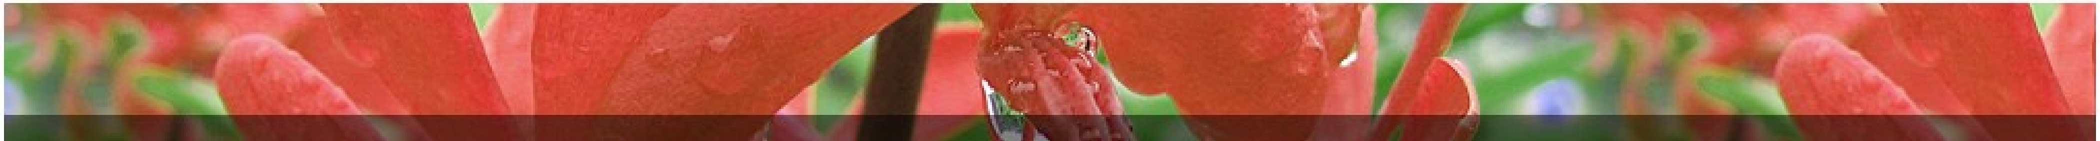

VirtualBox

- innotek (Germany), then Sun Microsystems, now Oracle
- True Virtualization
- Great features
- Nice GUI interface on the desktop
- Also has a command line interface
    - but tends to be too verbose for my taste

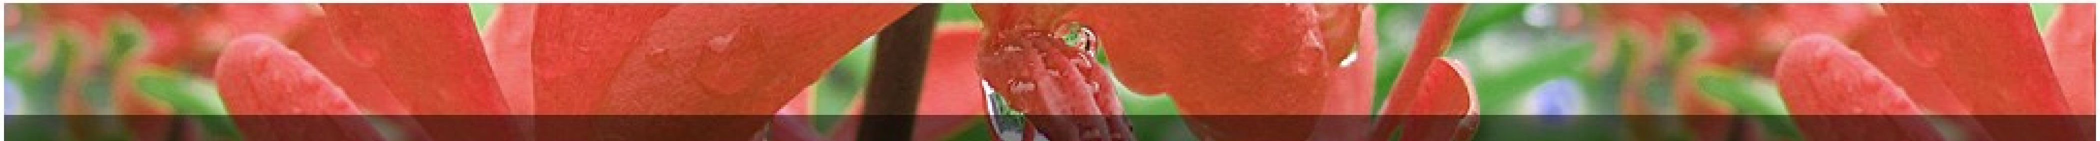

VirtualBox

- Quick Demo

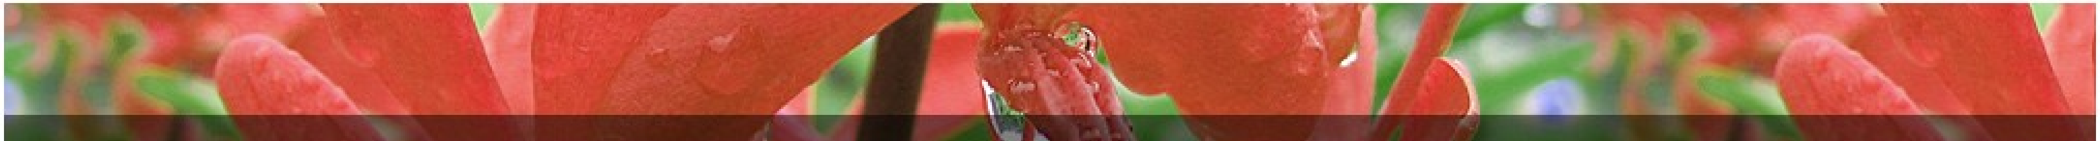

Virtuozzo/OpenVZ

- Parallels Inc.
- Operating System Level Virtualization
- File system: vzfs, maps to host file system directories
- Burst RAM (= swap)
- Slow I/O
- Used by many hosting companies
- Often slow for some Drupal sites

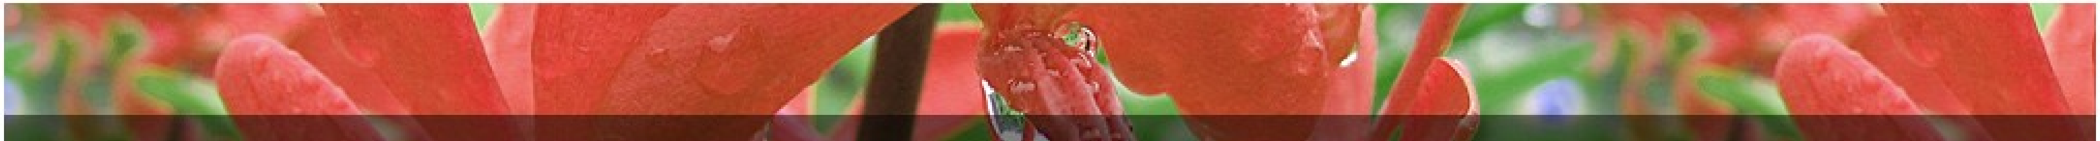

Xen

- University of Cambridge
- GPL V2
- Hypervisor with microkernel design
- Very good performance
- Popular with hosting companies
  - Linode, Amazon EC2, ...etc.)
- Been using it for years, but only as a guest

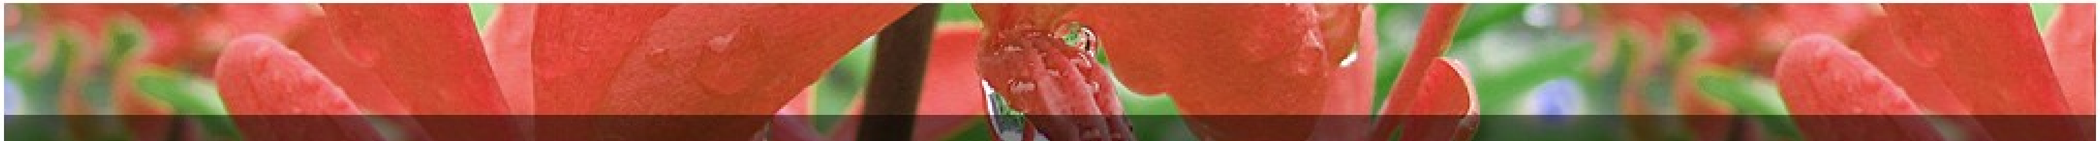

KVM

- Kernel based Virual Machine
- Full virtualization for Linux on x86
- Uses hardware virtualization
  - Kernel module for Intel, another for AMD
- Each VM has private virtualized hardware (network card, graphics adapter, ...)
- http://www.linux-kvm.org/

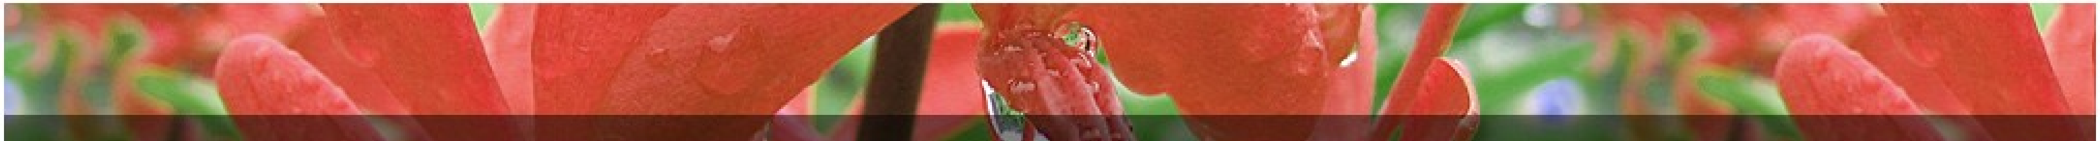

LibVirt

- Abstracts the many variants into one API
- Same commands should work on many underlying technologies (Qemu, VirtualBox, KVM, ...)
- Allows datacentres with many virutal machines to minimize the amount of work that goes into scripting their deployment and management
- http://libvirt.org/

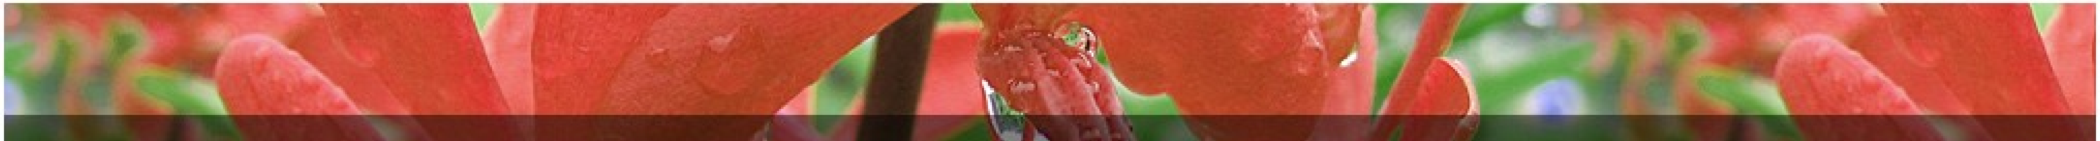

LibVirt

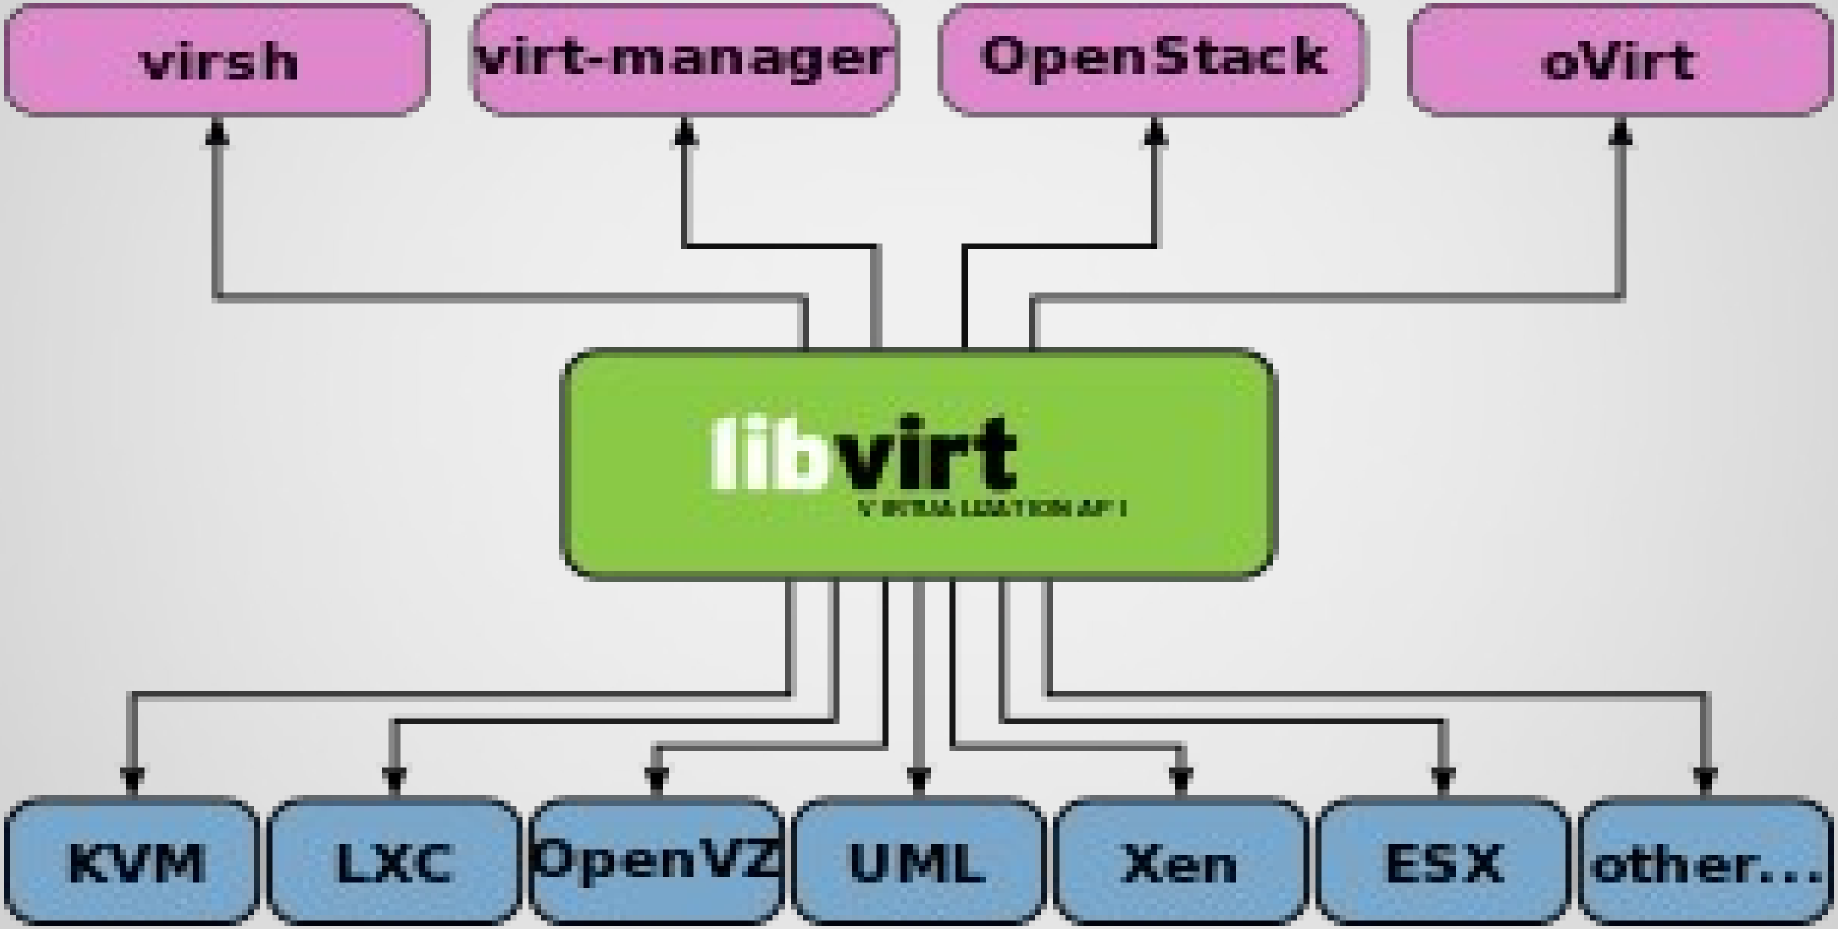
virsh
virt-manager
OpenStack
oVirt
libvirt
VIRTUALIZATION API
KVM
LXC
OpenVZ
UML
Xen
ESX
other...

Installing libvirt

- `sudo aptitude install \`
- ` qemu-kvm libvirt-bin virtinst virt-viewer virt-top`
- Add your user to the libvirt group (/etc/group)

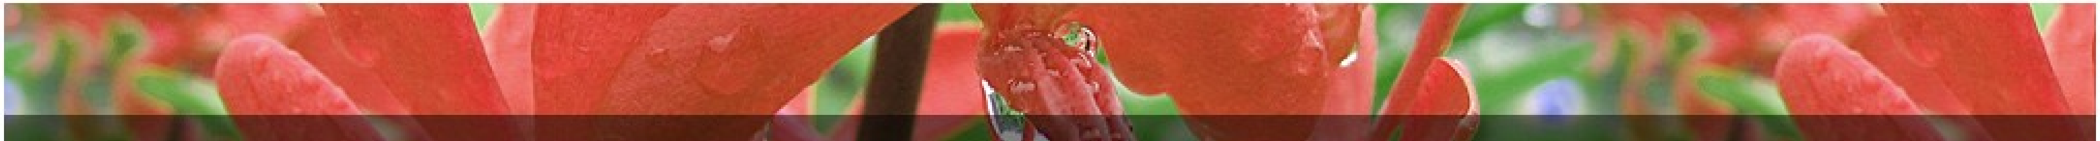

Creating a Network Bridge

- Use wlan0 if you are on a laptop, or eth0 if you on a server, in /etc/network/interfaces
- `auto br0`
- `iface br0 inet dhcp`
- `  bridge_ports eth0`
- $ sudo service networking restart
- Bad idea! Stopped the host from connecting to the network on boot.

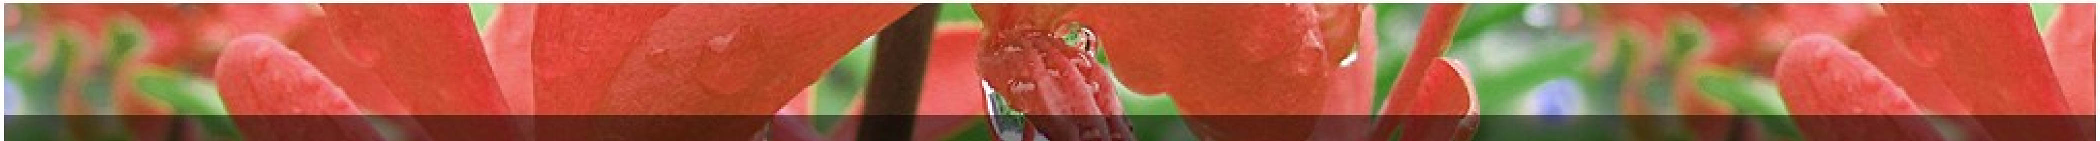

Creating a NAT Network

- $ virsh net-edit default
- `<range start='192.168.122.2' end='192.168.122.254' />`
- `<host mac='fe:54:00:00:00:01' name='p1.example.com' ip='192.168.122.11' />`
- `<host mac='fe:54:00:00:00:02' name='t2.example.com' ip='192.168.122.12' />`
- $ virsh net-destroy default
- $ virsh net-start default

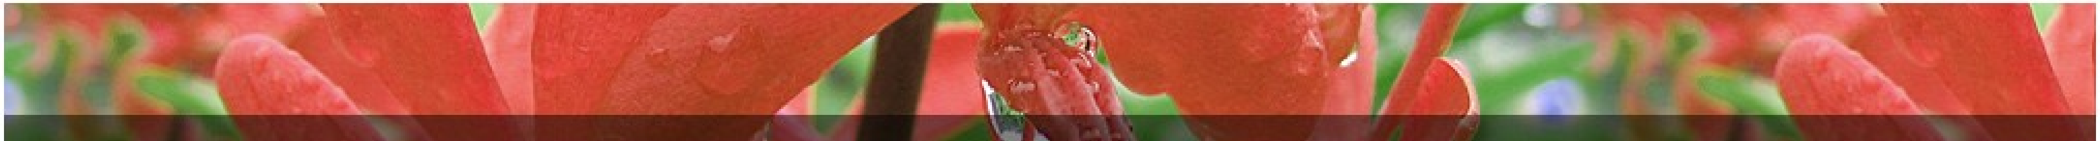

Creating a VM (domain)

- Download the ISO from a mirror
- Use this command (with sudo)
- `sudo virt-install \`
- ` --name t8 --ram 512 --disk \`
  `path=~/virt/disks/t8.img,bus=virtio,size=4 \`
- ` --cdrom ~/virt/isos/trusty64/ubuntu-14.04.1-`
  `server-amd64.iso \`
- ` --boot cdrom \`
- ` --graphics vnc`

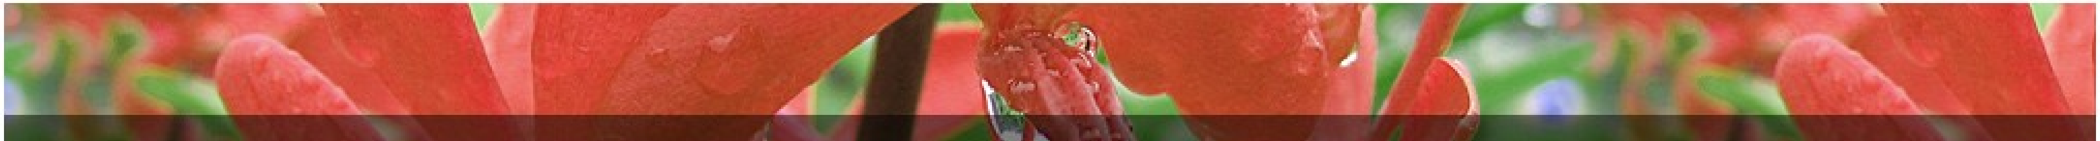

Creating a VM (domain)

- If you are using a bridged network, then use:
-
  ```
    --network bridge=br0,mac=fe:54:00:00:00:11
  ```
- Connect to the console (from your desktop to the server):
-
  ```
  $ virt-viewer -c qemu+ssh://host/system t8
  ```
- Go through normal steps for installing Linux

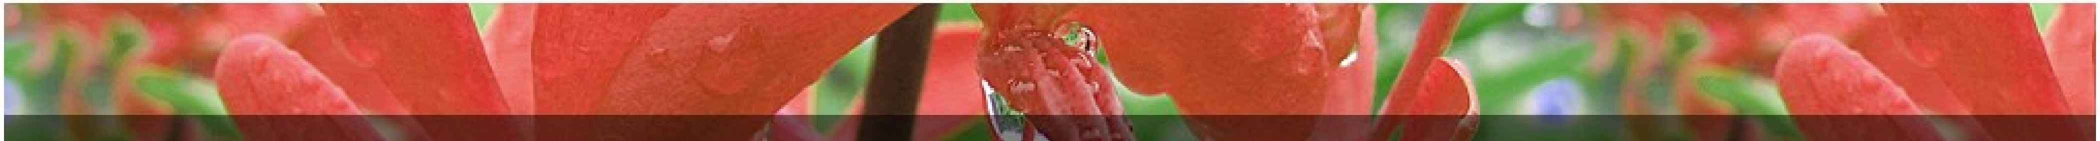

Edit your VM

- Several steps needed before you boot again ...

```
– $ virsh edit t8
```

- 1. Make it boot from disk, not CD-ROM

```
– <boot dev='cdrom'/>
– <boot dev='hd'/>
```

- 2. Change MAC address (if you want to)
- 3. Remove the CD-ROM

```
– <disk type='block' device='cdrom'> ...
  </disk>
```

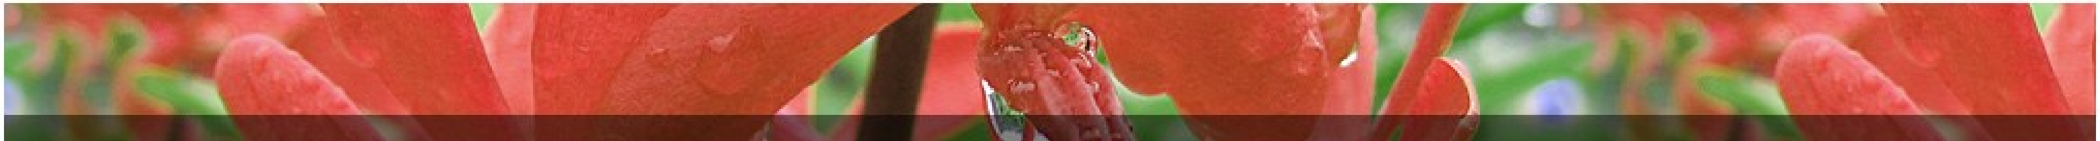

Start your VM

- On the server

  -
    ```
    $ virsh start t8
    ```
- On your desktop

  -
    ```
    virt-viewer -c qemu+ssh://host/system t8
    ```

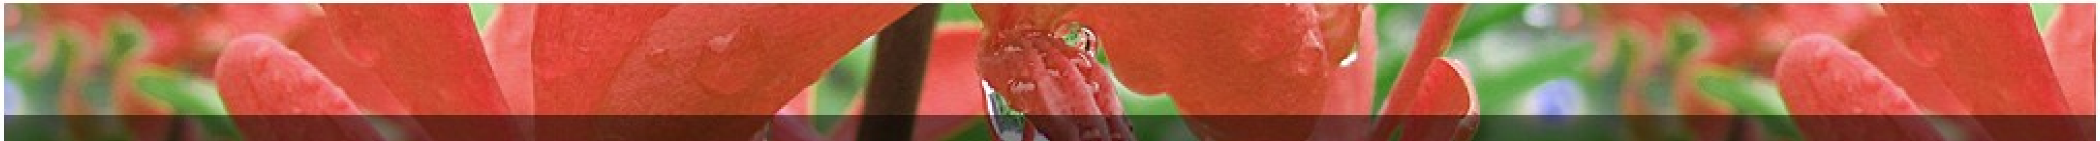

Virsh

- Interactive shell for libvirt, many subcommands
- Starting and stopping a VM
  - `virsh start vm_name`
  - `virsh shutdown vm_name`
- Delete a VM (undefine
  - `virsh undefine vm_name`
- Listing running VM
  - `virsh list --all`
- Much more ...

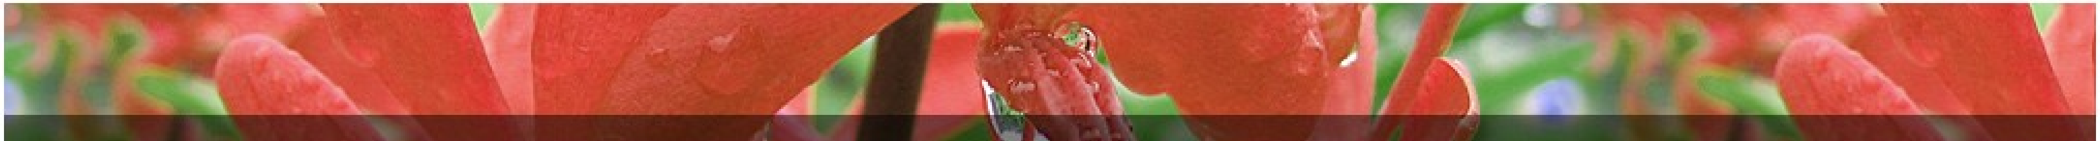

Compressing disk images

- Default images are “qcow2”, not “qcow”
- `$ qemu-img convert -f qcow2 -c -O qcow t8.img t8-new.img`
- `$ mv t8-new.img t8.img`
- `$ virsh edit t8`
- `<driver name='qemu' type='qcow2'/>`
- `<driver name='qemu' type='qcow'/>`

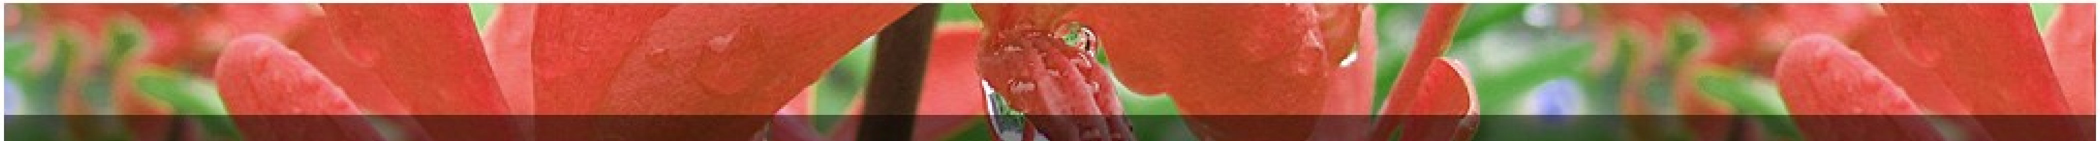

Tips and Tricks

- Save a pristine disk image (starting point)
  - Libvirt allows you to swap disk images, unlike VirtualBox, which stamps them with UUIDs, and complains if you swap them
- Snapshots are slow
- “Wait for I/O” can be a bottleneck (on regular single disk servers)
  - e.g. running aptitude full-upgrade on multiple instances simultaneously

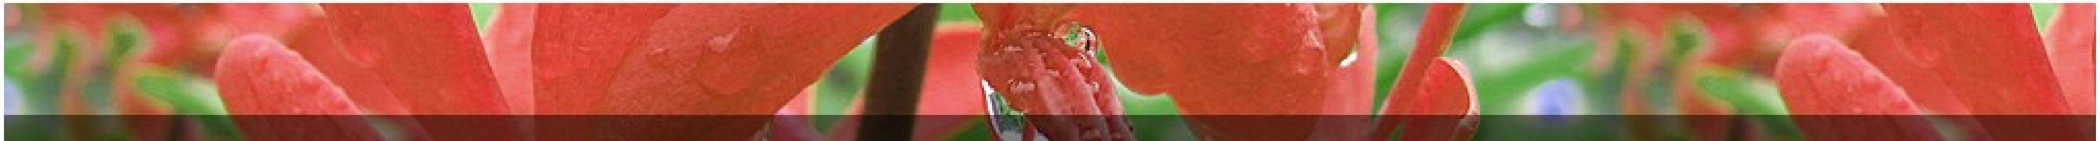

Bonus

Still with me?

virt-manager

- The easy way (GUI)
- Can manage one or more remote servers, not only your local desktop

```
$ sudo aptitude install virt-manager
```

- Add a connection to the server, and off you go
- You can monitor CPU, disk and network usage
- You can create VMs

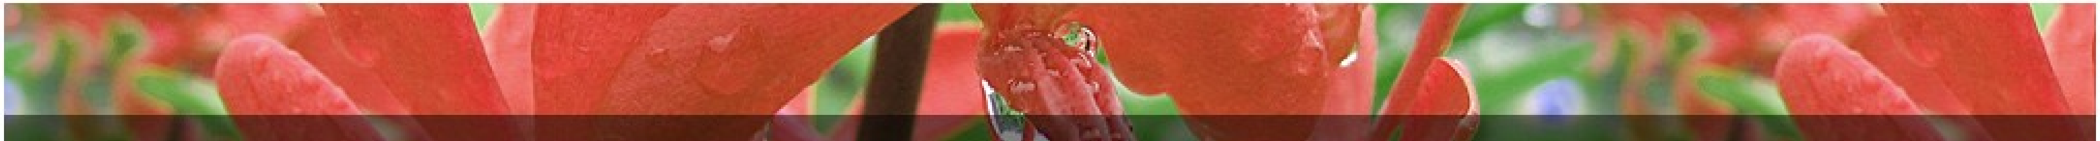

Questions?

Any questions? Comments?