Automated Instance Creation

Brian Candler
Network Startup Resource Center
brian@nsrc.org

Partitioning: choices, choices!

MSDOS partitioning, single partition, Linux filesystem

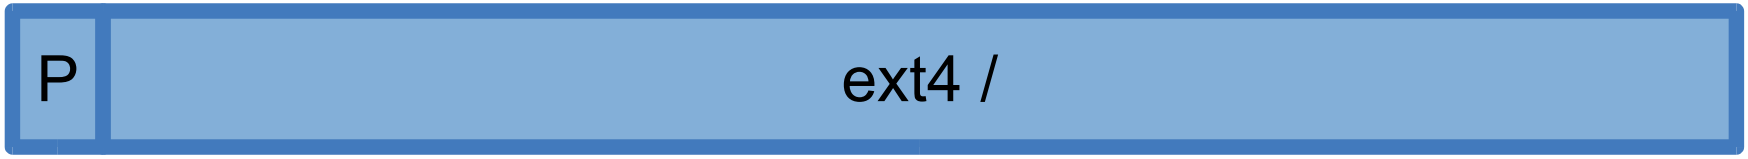

MSDOS partitioning, multiple partitions

No partitions (whole disk filesystem)

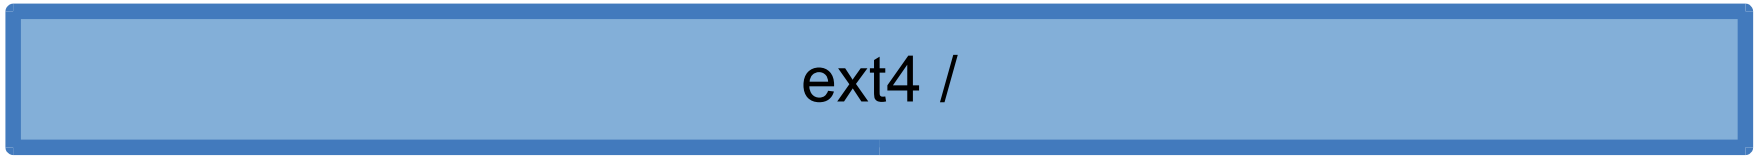

Master boot record is here (required to boot from the instance disk, and contains partition table)

Why does partitioning matter?

- To automatically create instances you need to have scripts to make the final adjustments
  - e.g. change hostname, change IP address, set password...
- Hence the guest filesystem(s) must be mounted temporarily
- Therefore scripts must know how it's partitioned

With a partition table / MBR

- BIOS boots from the Master Boot Record in the guest's disk image
  - standard, just like a normal PC
- Fairly easy to grow the *last* partition
  - increase size of disk
  - delete and re-add the last partition
  - grow the filesystem (resize2fs)

Without a partition table / MBR

- Very easy to grow filesystem
  - increase size of disk, grow the filesystem (resize2fs)
  - but no partitions = no MBR = can't boot from disk!
- KVM can be told to boot a kernel and initrd stored on the *host* filesystem
  - kernel_path, initrd_path, kernel_args
  - no bootstrap loader (e.g. GRUB) required
  - only works for Linux VMs

Instance export/import

- The OS definition scripts include 'export' and 'import' functions
  - gnt-backup {export|import}
  - see "man gnt-backup"
- Useful for backups and for moving instances between different ganeti clusters
- But only works if the disk is partitioned and formatted in the way the scripts expect

Instance export/import (2)

- When exporting the instance, we want to use a tool like "dump" which understands the filesystem and skips over unused blocks
- Export the correct partition(s)
- We don't want to export swap partition!
  - it would be a waste of space
  - it would be a security risk (contains RAM contents)

Safe resizing

- If you have export and import working correctly (with dump), then you can export the instance and re-import it to a different-sized filesystem, or to a different cluster
- This is the safest way to resize a VM

Partitioning: non-Linux guests

e.g. Windows instance

- This type of filesystem can only be exported as a raw disk file
  - ganeti-instance-image: set "NOMOUNT=yes" to export/import whole disk as a qcow2 image

Note about raw dumps

- Less efficient than ext3/4 dump/restore as it will include blocks from deleted files
- Can only be restored to a disk image of *exactly the same size*

Questions?

ganeti-instance-image

- Most basic usage is installing from ISO with option "CDINSTALL=yes"
  - same effect as "--no-install"
  - creation of partitions is the responsibility of the ISO installer
- But it can also do scripted installs from prepared image dumps

ganeti-instance-image

- Script creates 2 or 3 partitions (boot+root, or boot+swap+root)
  - ext3 or ext4
  - See /usr/share/ganeti/os/image/create

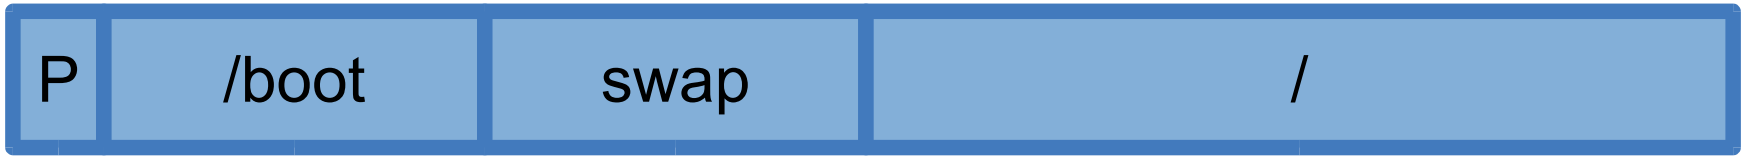

Consequences

- If you want to export/import an image created with instance-image, either it must be partitioned this way OR you must make a raw disk image (NOMOUNT=yes)
- If you are preparing an OS image to clone using instance-image, you need to make two partitions (boot and root) and dump them
- The README file explains this
  - /usr/share/doc/ganeti-instance-image/README.gz

ganeti-instance-debootstrap

- OS definition to automatically install Debian or Debian-derived OS (e.g. Ubuntu)
- Downloads all the .deb packages and unpacks them
- Creates a cache so that subsequent installs are very fast

Debootstrap partitioning

- Default is to install a partition table and *one* partition
- Or you can install without any partition table
  - PARTITION_STYLE="none"
  - makes resizing filesystem very easy
  - but then you *must* boot from a kernel on the host

Debootstrap booting

- You can boot from a kernel on the host
  - reasonable if you have many nearly-indentical VMs
- If you want to boot from a kernel in the guest FS you need to install grub in the guest
- An example hook script is provided to do this, but it doesn't work properly :-(
- Good practice in how to recover a VM with broken boot loader :-)

Now the good news!

- The instance creation scripts are simple, easy to read/understand and modify
- Look in directories under /usr/share/ganeti/os/
- Settings in /etc/default/ganeti-instance-<xxx>
- Documentation: "man ganeti-os-interface"
  - docs.ganeti.org/ganeti/master/man/ganeti-os-interface.html

Finally, two more options

Importing existing image

- You can take a VM image created somewhere else and copy the disk
  - VM "appliances" are now supplied this way
- You will need to convert the disk image
  - "qemu-img convert" will do this
- Create instance with exactly the right sized disk
- Beware writing to drbd-replicated volumes
  - Safer to create -t plain, then convert to -t drbd

snf-image

- Standalone component of the synnefo cloud solution
- Provides ready-made images you can clone
  - or create your own (snf-image-creator)
- Very robust
- Works with Windows and BSD images too!

snf-image architecture

- Uses raw disk dumps
  - boot loader already installed, no messing around
  - partition how you like (except no LVM)
- Post-install "helper" enlarges the last partition and filesystem to fit the chosen disk size
  - also sets password, installs ssh keys etc
  - helper runs inside a temporary VM for security
  - Ganeti is moving to this model too

Summary of installation options

- instance-image: install manually from ISO
- instance-image: unpack filesystem dump or tarball that you prepared earlier
- instance-debootstrap (Debian/Ubuntu only)
- import an existing VM disk image
- snf-image
  - probably best option for self-service installs

Exercises (depending on time)

- Create VM using debootstrap
- Create VM using VDMK disk image
- Create VM using snf-image