



# Virtual machine snapshots

(Overview and illustration of different types of snapshots with libvirt/QEMU)

Red Hat

Kashyap Chamarthi – {kashyap@redhat.com}

5 Nov 2012

# Agenda

- 1 Virtual disk image file basics
- 2 Snapshots overview
- 3 Managing multiple snapshot files
- 4 Demonstration
- 5 Upcoming...



## Section 1


# Virtual disk image file basics

## Common representation of virtual disk images

- **RAW:**
  - Storage *preallocated* ahead of time
  - Best performance for **I/O** intensive guests
- **QCOW2 (Qemu Copy-On-Write):**
  - Sparse by nature – storage allocated as needed
  - Supports base images & overlays
  - Very useful for – 'thin-provisioning' guests; snapshots

## QCOW2: backing files and overlays

- QCOW2 allows a disk image to *refer* to another image ('backing file/base image')
- Overlays record only the **difference** from its (read-only) backing file





## Section 2

# **Snapshots overview**

## Snapshot flavours

- Internal Snapshots
  - Internal **disk snapshot** – {Live/Offline}
  - Internal **system checkpoint** – {Live}
- External Snapshots
  - External **disk snapshot** – {Live/Offline}
  - External **system checkpoint** – {Live}
- VM State
  - Saves memory, device state (*not* disk-state) to a file

## Internal snapshots

- A *single* QCOW2 file holds everything {original & its delta}
- Very handy when moving around machines
- The original *must* be a QCOW2 file; guest is *paused* for more than a couple of seconds while taking snapshot




## External snapshots

- Uses **overlays** – Everytime a snapshot is taken, the current disk becomes a 'backing file' & a *new* QCOW2 overlay is created to track the 'delta'
- Snapshots are created instantaneously
- Excellent for **live** backups

## External snapshots(2)

- Multiple files to track; 'revert/delete' operations become a little tricky (more on this next..)





## Section 3

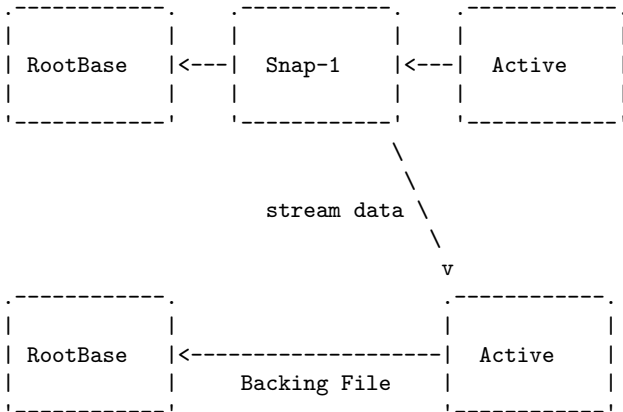
# Managing multiple snapshot files

## With external snapshots..

- **Problem:** – Lots of snapshot disk image files accumulate over time
- **Solution:** – Merging snapshot disk image files:
  - **blockpull** – (Live) merges data from **base** into **top**
  - **blockcommit** – (Live) merges data from **top** into **base**

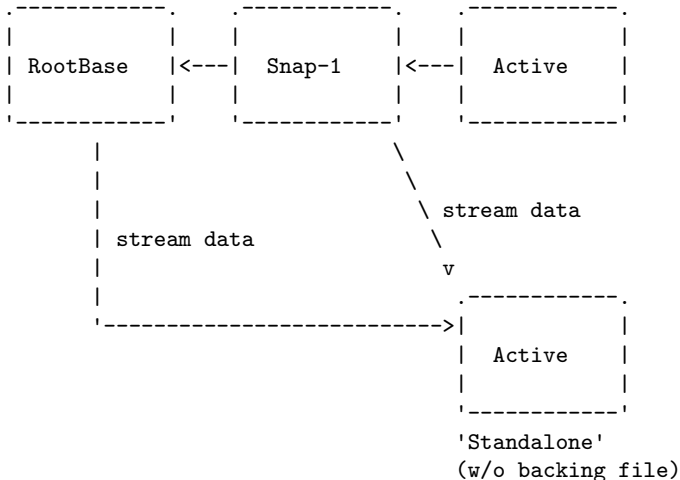
# blockpull: shorten image chain w/ 'partial' pull

- Merge data from **Snap-1** into **Active**
  - **RootBase** will now be the backing file of **Active**; (**Snap-1** remains valid & can be reused)



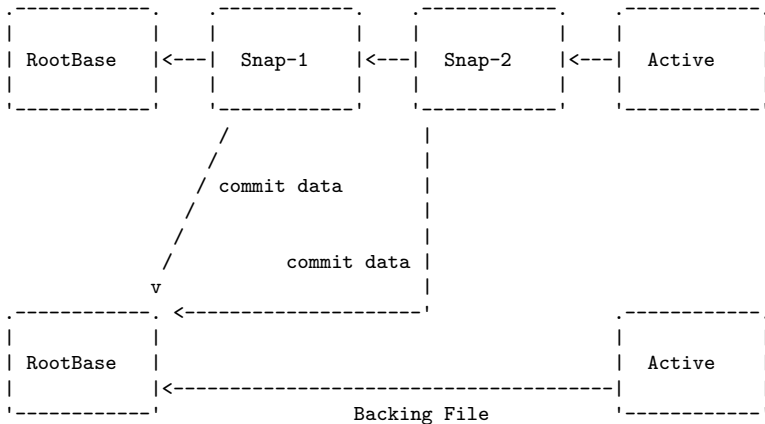
# blockpull: shorten image chain w/ 'complete' pull

- Merge data from **RootBase** and **Snap-1** into **Active** – which will now be *standalone* w/o any backing file
  - **RootBase** & **Snap-1** remain valid & can be reused



## blockcommit: commits 'top' into 'base' – shortens image chain length

- Merge data from **Snap-1** & **Snap-2** into **RootBase**
  - **RootBase** will now be the backing file of **Active**
  - **Snap-1** & **Snap-2** will be invalidated





# Section 4

## **Demonstration**



# URL to demonstration info

- `http://kashyapc.fedorapeople.org/virt/lc-2012/snapshots-illustration.txt`
- `http://kashyapc.fedorapeople.org/virt/lc-2012/live-backup-with-external-disk-snapshots-and-blockpull.txt`



# Section 5

## **Upcoming...**

## What's next

- Quite a few interesting things:
  - External [system checkpoint/offline] snapshots in Libvirt
  - Snapshot revert/delete improvements for external snapshots
  - Live/Offline **blockcommit** enhancements
  - QEMU Live **blockcopy** (for storage migration)
  - Plenty other bug fixes & updates...

# References



More notes on snapshots, blockcommit, blockpull

<http://kashyapc.fedorapeople.org/virt/lc-2012/>



Libvirt website

<http://www.libvirt.org/>



QEMU website

<http://www.qemu.org/>



KVM website

<http://www.linux-kvm.org/>



libguestfs

<http://libguestfs.org/>



Blog

<https://kashyapc.wordpress.com>



# The end.

Thanks for listening.