

Virtualization Workshop: Virtual machines and snapshots

Saving and restoring state with virtual machines

What are we trying to do ?

Make it possible to save the state of a Virtual Machine

- By saving the state, we mean saving the contents of the disk, memory and CPU
- Saving the state allows multiple things:
 - testing possibly destructive operations (installing software, upgrading the operating system, ...) without damaging / modifying the Virtual Machine
 - restoring the Virtual Machine to a previously known “good state” (possibly following one or more destructive operation)

We call this “saving of the state” *Snapshots*

All hypervisor software suites offer Snapshotting as a feature (VirtualBox, VMWare, KVM/QEMU, XEN, Hyper-V, ...)
Snapshotting is one of the unique benefits of using virtualization in the first place.

How do snapshots work ?

In a nutshell, when you create a snapshot:

The hypervisor. . . :

- pauses operation of the Virtual Machine
- creates a special disk image ^a (the snapshot file)
- saves the state of the CPU, and the RAM
- ensures that the original disk image for this virtual machine will not be modified by write operations
- resumes operation of the Virtual Machine

^aThe snapshot file doesn't necessarily appear as a separate file. It can be part of the disk image file used by the Virtual Machine, in a special section.

Hypervisor and disk image files

Remember the hypervisor controls the Virtual Machine's resources, including disk access.



Figure 1:

Every time the Virtual Machine wants to write to the disk, the hypervisor intercepts this, looks up the block number being written, and maps this to the proper location in the Disk Image File.

When using snapshots

Once a snapshot has been created, the Hypervisor stops writing changes to the original disk image file. Instead, it will *write* changes to the snapshot file.




Figure 2:

Note the direction of the arrows.

Snapshots and copy-on-write

The snapshot file only contains *changes* (differences) to the disk image file: the disk image file itself is not modified.

The Hypervisor still uses the disk image file to READ, but any changes go to the snapshot.

This is called Copy On Write.

Copy on Write

Let's see what happens when we read and write disk blocks to a disk image file. . .

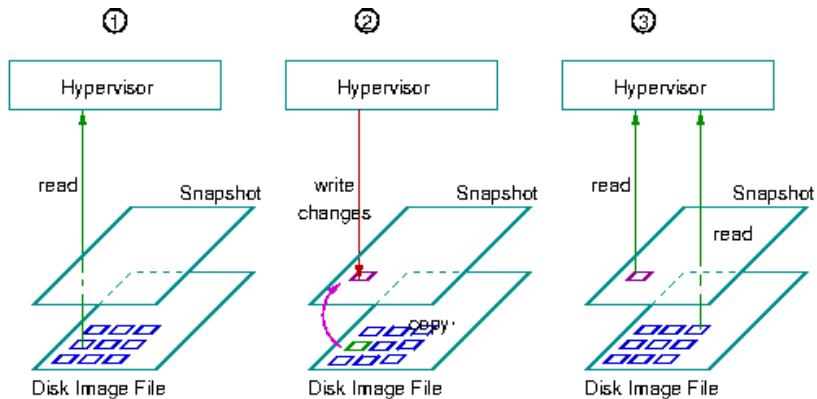


Figure 3:

Let's imagine that our change didn't go very well:

- failed upgrade
- software install for testing
- or simply because we made a mistake and deleted some files

We have the possibility of rolling back to a previously known state (snapshot). This will lose the current state of the machine, and “travel back in time” to the previous state.

Another benefit of snapshots: Cloning

A side benefit of using snapshots, is that we can create an entirely new disk image file without having to copy all of the image to a new file.

Cloning is slightly different, because it will allow us to use the “source” machine and its disk images, while we run “clone” machines based on the “source” machine.

We will see this later.