

Incremental volume snapshots/backups in KVM plugin of Apache CloudStack

Daniel Augusto Veronezi Salvador
<gutoveronezi@apache.org>



Agenda

- ❖ Introduction
- ❖ Snapshots (backups?) on Apache CloudStack
- ❖ KVM volume snapshots History in ACS
- ❖ Incremental volume snapshots for KVM
- ❖ Unlimited or limited backup chains?
- ❖ Limitations
- ❖ Future works
- ❖ Hands-on
- ❖ Questions

Introduction

- ❖ **Backup:** an extra copy of information on a computer that is stored separately*
- ❖ **Snapshot:** capture of a system at a point in time
- ❖ CloudStack native features
 - Snapshot: for **volumes** and VMs
 - Backup: for **volumes**

* <https://dictionary.cambridge.org/dictionary/english/backup>

Snapshots (backups?) on Apache CloudStack

❖ VM snapshots

- Per VM
- Can be memory + disk or disk-only
- Cannot be scheduled
- Cannot be backed-up to the secondary storage
- Volume can be extracted

❖ Volume snapshots


- Per volume
- Can be scheduled
- Can be **backed-up** to the secondary storage (that is done by default)
- Can be used to create new volumes and templates

❖ We will focus on **volume snapshots**

KVM volume snapshots History in ACS v1.0

Original volume snapshot process:


- ❖ Slow
- ❖ Costly
- ❖ VM downtime



KVM volume snapshots History in ACS v2.0 (current)


Current volume snapshot process ([#5124](#)):

- ❖ Faster than the old one, but its still slow
- ❖ Costly
- ❖ No downtime



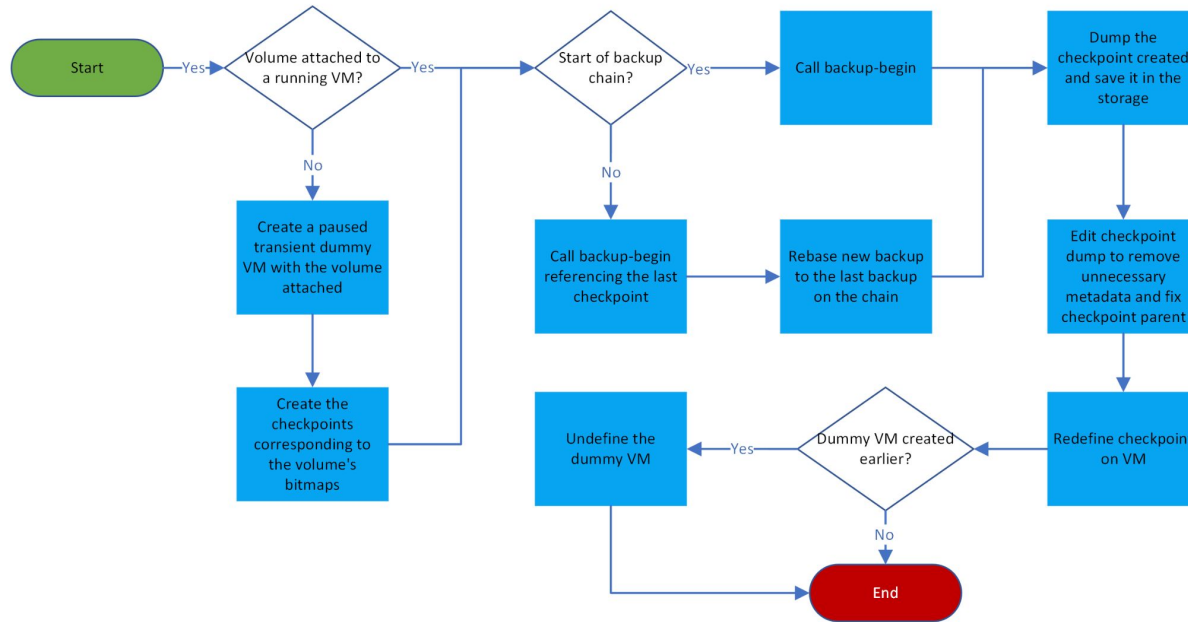
Incremental volume snapshots for KVM

- ❖ Only for volumes using QCOW2
- ❖ Backing chains
- ❖ Bitmaps
- ❖ Checkpoints
- ❖ At least one full volume snapshot
- ❖ “Manually” maintain the snapshot backing chain



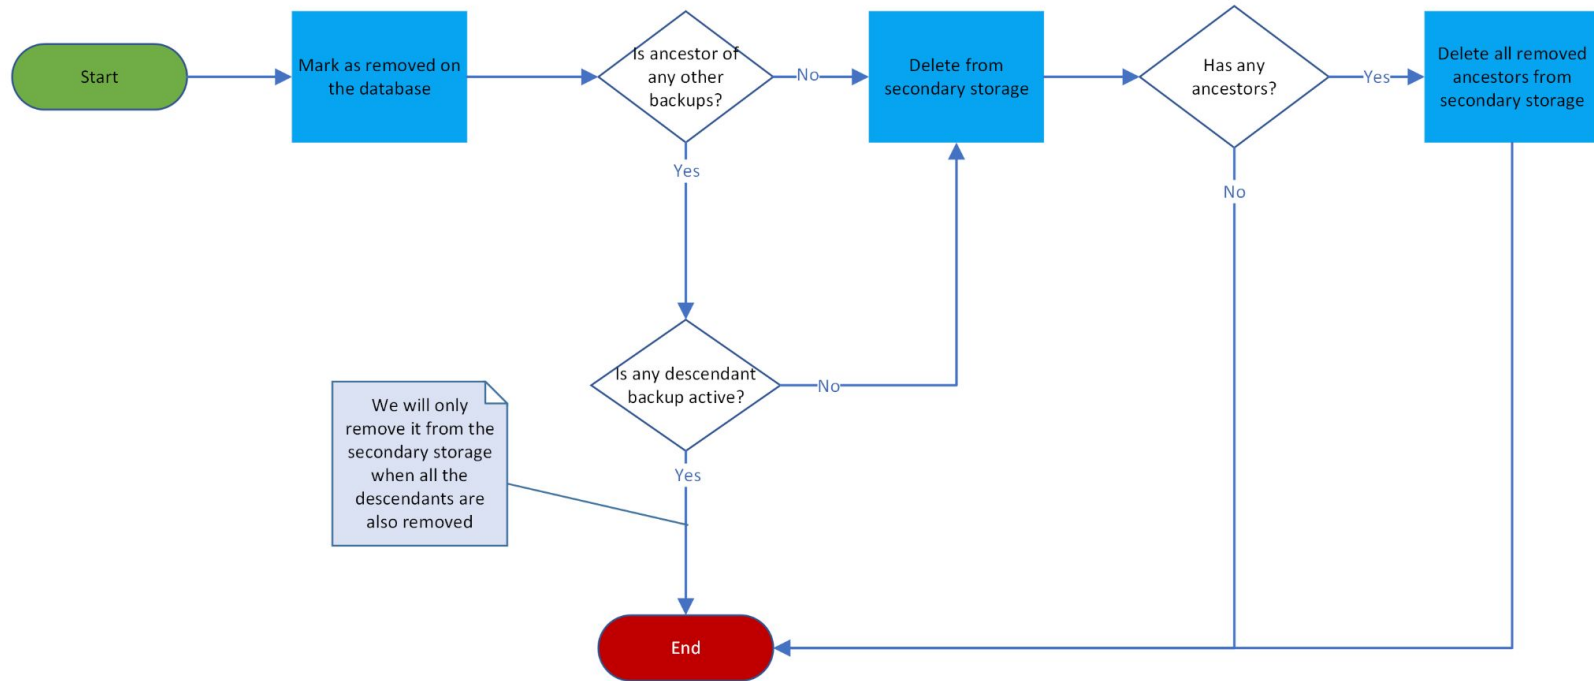
Incremental volume snapshots for KVM

Creation workflow



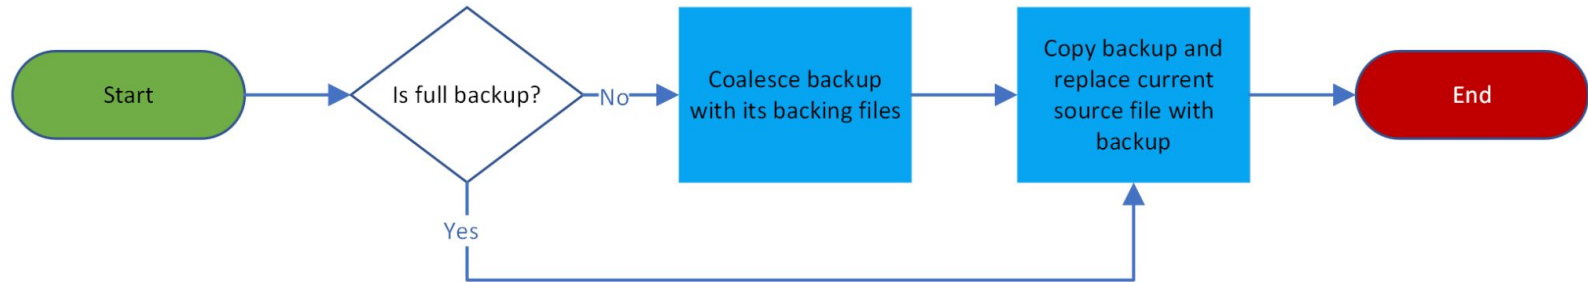
Incremental volume snapshots for KVM

Deletion workflow



Incremental volume snapshots for KVM

Reversion workflow



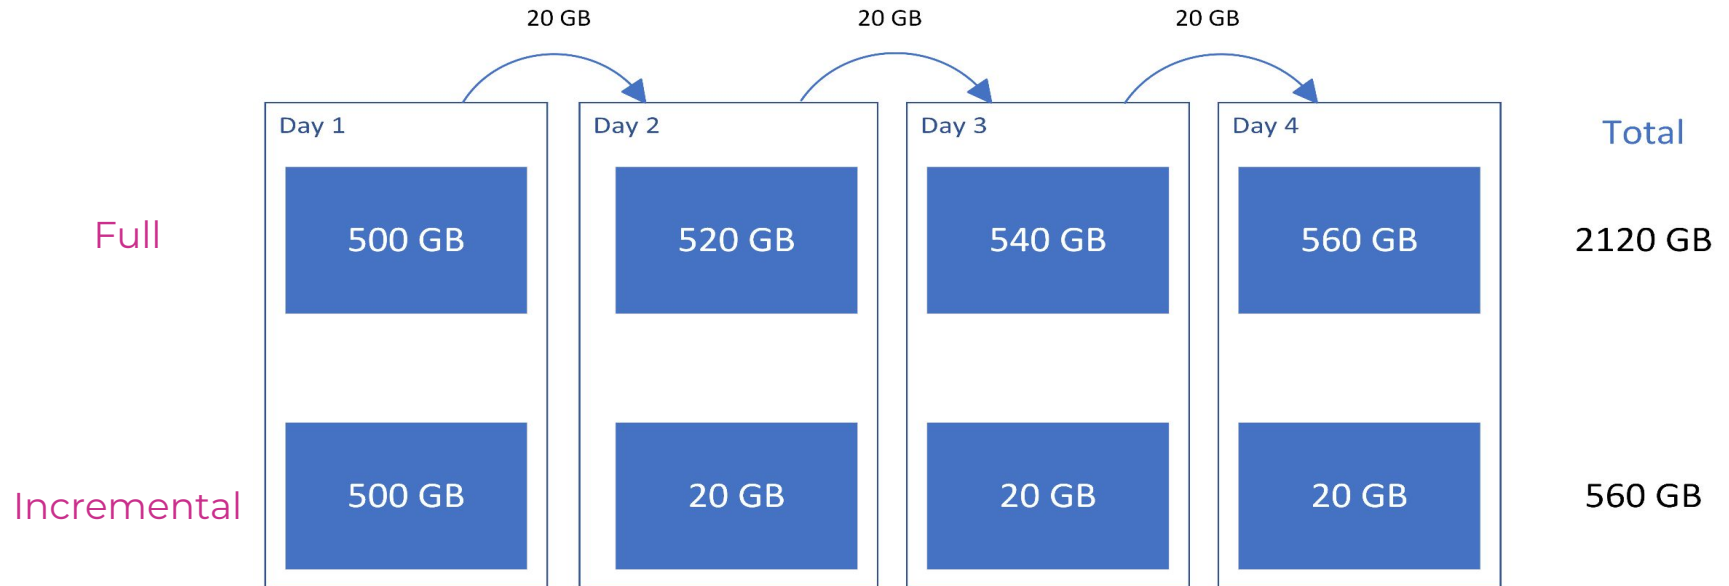
Incremental volume snapshots for KVM

Other processes affected

- ❖ Cold Volume Migration
- ❖ Volume Attach
- ❖ Volume Detach
- ❖ VM Start
- ❖ VM Migration
- ❖ Template/Volume Creation from Snapshot
- ❖ Snapshot Download
- ❖ Snapshot copy between zones

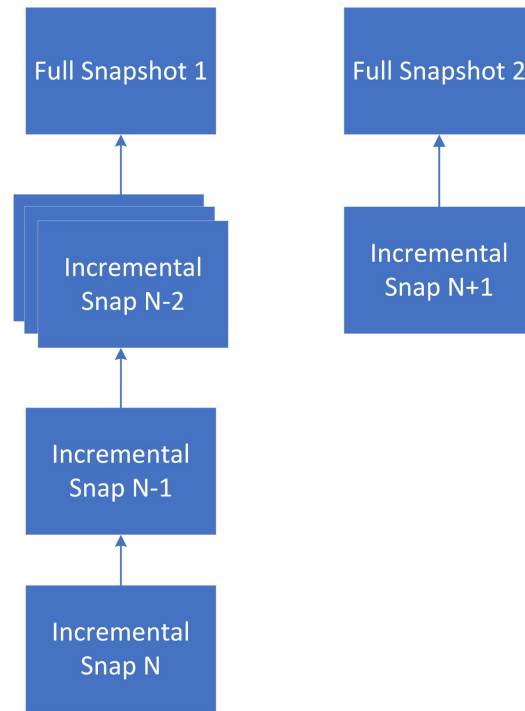
Incremental volume snapshots for KVM

Why use them?



Unlimited or limited backup chains?

- ❖ “Unlimited” chains:
 - Only the first snapshot is a full snapshot
 - After the first snapshot, all snapshots are really fast
 - Saves time for taking snapshots
 - Saves network/storage resources
 - What if the first snapshot is lost?
 - Slow deletion process for old snapshots
 - It would require snapshots to be merged in the secondary storage
- ❖ Limited chains:
 - A full snapshot after X number of snapshots
 - Lower risk of losing all the data
 - Faster snapshot deletion
 - Consumes more network/storage resources



Limitations

- ❖ Libvirt 7.6+ and Qemu 6.1+
- ❖ New snapshot chain after
 - Hot volume migration
 - Snapshot reversion
- ❖ Only files based storage (local, NFS, shared mount point)
 - RBD (Ceph), CLVM and iSCSI not supported yet

Future works

- ❖ Incremental Snapshots for RBD storage (Ceph)
- ❖ Incremental Snapshots for VMware
- ❖ Atomic disk-only VM Snapshots for KVM
 - To be pushed upstream ([#9254](#))
- ❖ Native VM backup (WIP)

Let's look at a
demonstration!

Questions?

gutoveronezi@apache.org