oVirt

# The future integration points for oVirt storage

***Birds of a feather session***

Sean Cohen, Ayal Baron
*Red Hat*

KVM Forum, October 2013

# Agenda

oVirt

- What's new in oVirt 3.3 Storage?
- Storage directions
  - Software Defined Storage
  - Storage Offloading
  - IO Performance
  - Replication
- The future integration points for oVirt Storage
  - Openstack
  - DRDB
  - Extending oVIrt API
- 3.4 & Beyond

# What's new in oVirt 3.3 Storage?

oVirt

**Business Continuity:**

- Backup and Restore API for Independent Software Vendors *(3.3.2)*
- Manage Storage Connections (Multipath & DR)*(3.3.1)*

**Disk Management:**

- Enable online virtual drive resize
- Virtio-SCSI support
- Disks Block Alignment scan
- Disk Hooks (for disk hot-plug/unplug)

# Storage directions

oVirt

**Software Defined Storage (SDS)**

- GlusterFS native storage domain in oVirt 3.3
- Converged Storage Hypervisor vision (self hosted)

- **Storage Offloading:**
  - **LibStorageMgmt** integration
    - Ability to plugin external storage array into oVirt/VDSM virtualization stack, in a vendor neutral way
    - Can provide Array Offload capabilities, such as Snapshots/Clones, xCopy

# Storage directions

oVirt

## LibStorageMgmt

## Storage Offloading

- Open source, vendor agnostic library which provides an API for managing external storage arrays
- Current array support (varying levels of functionality)
  - Linux software target
  - SMI-I compliant arrays (NetApp, EMC etc)

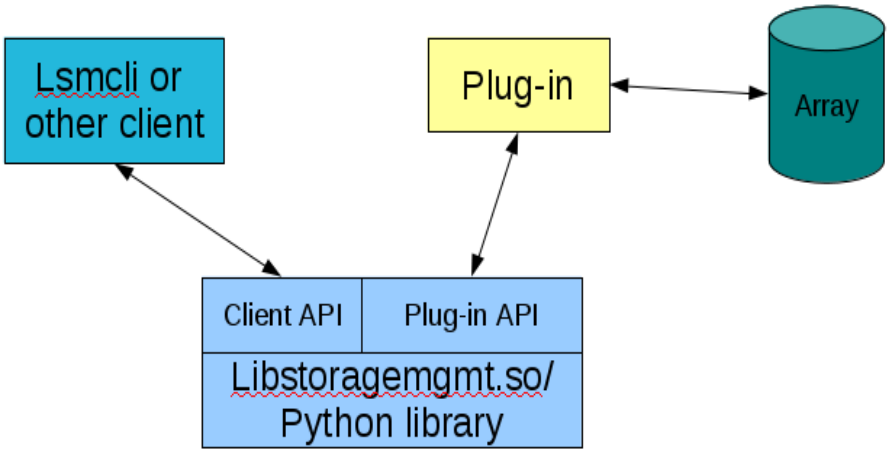

# Storage directions

oVirt

**Storage Offloading:**

- **OpenStack CInder** integration
  - Storage offload by design
  - LibSM Driver
  - Existing Eco System

**IO Performance:**

- On Hypervisors running multiple VMs all writes end up Random
  - Bad for:
    - SSD
    - SMR

# Storage Replication directions

oVirt

- **Array Based**
  - 3.3 DR support via Manage Storage Connections
- **Distributed**
  - Gluster
  - Ceph
  - DRBD
- **Hypervisor based**
  - QEMU-KVM
- **Cinder?**

# Future integration points

oVirt

- **OpenStack Storage stack integration**
  - Glance - consumed in oVirt 3.3
  - Cinder
    - Consume
    - Provide
- **Extending oVirt API**
  - oVirt third-party UI plug-in framework:
    - NetApp VCS 1.0 integration available for oVirt 3.2
  - oVirt Backup & Restore API next phases:
    - Leverage qemu-qa Microsoft Windows Volume Shadow Copy Service (VSS)
    - Leverage qemu block layer Change Block Tracking (CBT) to cover incremental backups.

# 3.4 & Beyond

oVirt

- Get rid of Storage Pool Manager
- Proper iSCSI-multipathing
- Single disk snapshots
- Import existing storage domain
- Read-only disks
- Support VM Fleecing (qemu-kvm/Libvirt)
- Snapshots Live Merge

oVirt

# THANK YOU !

http://www.ovirt.org

http://lists.ovirt.org/mailman/listinfo
vdsm-devel@lists.fedorahosted.org
engine-devel@ovirt.org
#ovirt irc.oftc.net

abaron@redhat.com
scohen@redhat.com