

oVirt and OpenStack Storage (present and future)

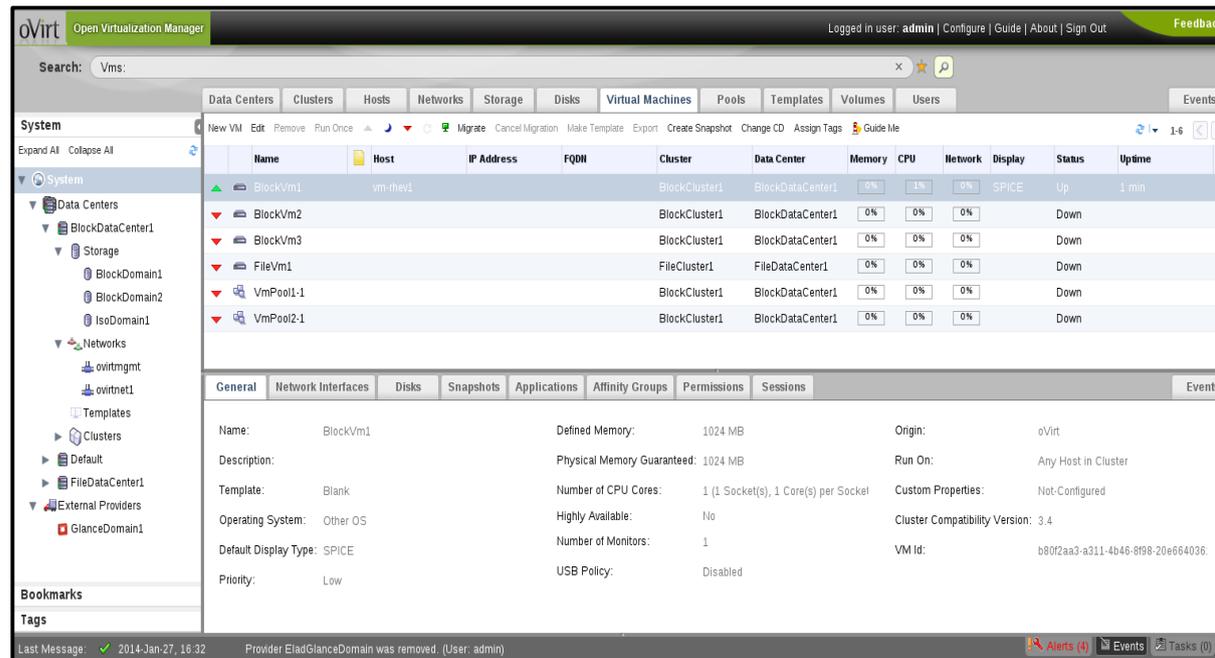
Federico Simoncelli

Principal Software Engineer, Red Hat

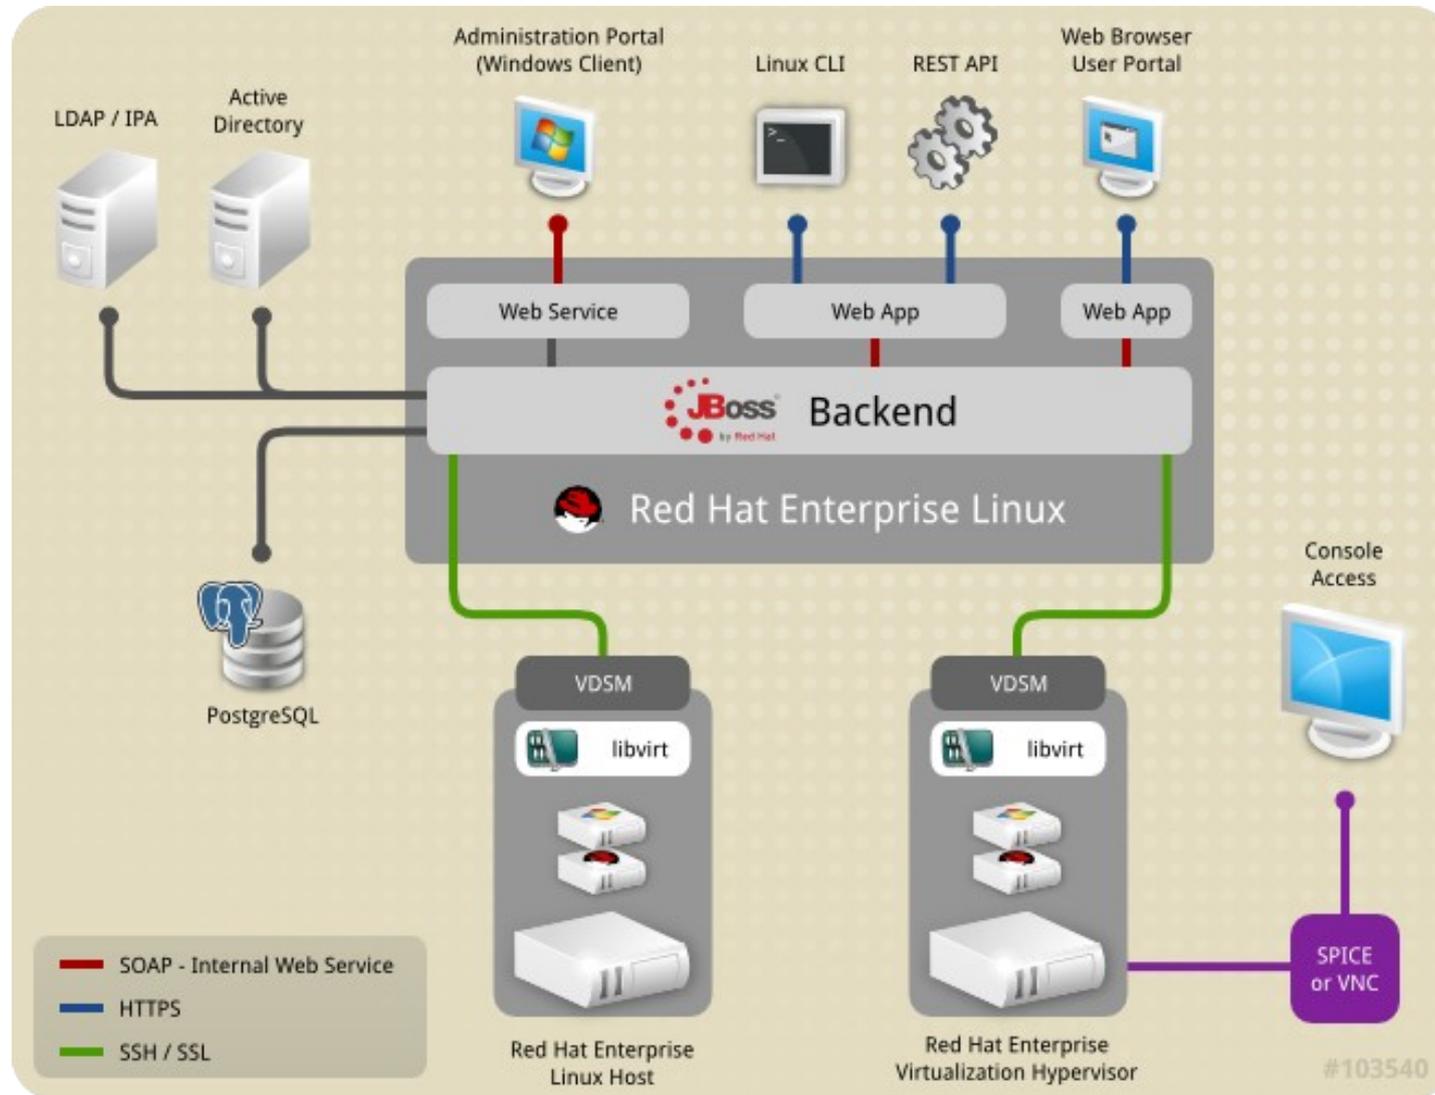
January 2014

- Introduction
 - oVirt and OpenStack Overview
- Present
 - oVirt and Glance Integration
 - Importing and Exporting Glance Images
 - Current Constraints and Limitations
- Future
 - Glance Future Integration
 - Keystone Authentication in oVirt
 - oVirt and Cinder Integration

- oVirt is a virtualization management application
- manages hardware nodes, storage and network resources, in order to deploy and monitor virtual machines running in your data center
- Free open source software released under the terms of the Apache License

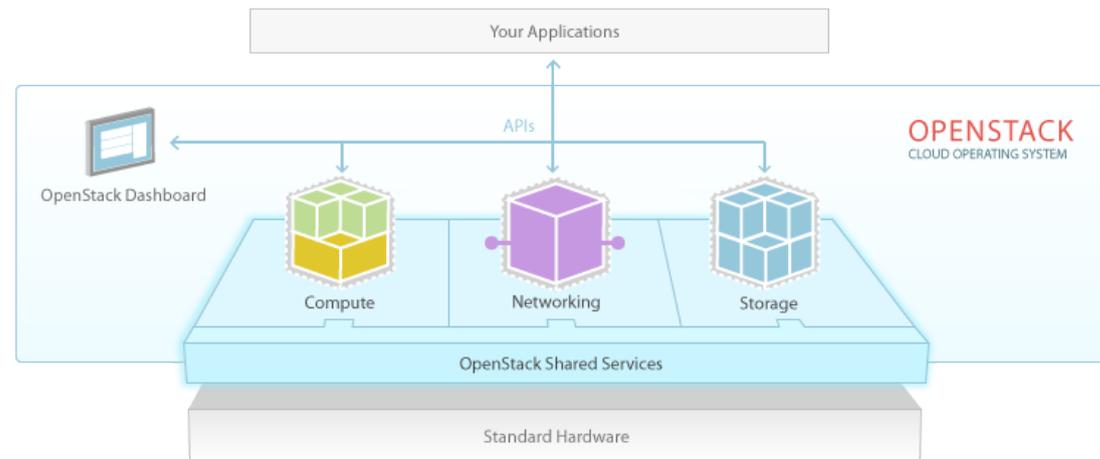


The oVirt Virtualization Architecture



OpenStack Overview

- Cloud computing project to provide an Infrastructure as a Service (IaaS)
- Controls large pools of compute, storage, and networking resources



- Free open source software released under the terms of the Apache License
- Project is managed by the OpenStack Foundation, a non-profit corporate entity established in September 2012

- Provides services for discovering, registering, and retrieving virtual machine images
- RESTful API that allows querying of VM image metadata and retrieval of the actual image
- Stored images can be used as a template
- It can also be used to store and catalog an unlimited number of backups
- Images available through Glance can be stored in a variety of locations (simple filesystems, OpenStack Swift, etc.)

Disk Formats

- **raw** (an unstructured disk image format)
- **qcow2** (disk format supported by the QEMU)
- **iso** (format for the data contents of an optical disc)
- **vhd, vmdk, vdi** (VMWare, Xen, Microsoft, VirtualBox, and others)
- **aki, ari, ami** (Amazon images)

Container Formats

- **bare** (no container or metadata envelope for the image)
- **ovf** (ovf container format)
- **aki, ari, ami** (Amazon containers)

Glance Image Base Metadata



- **checksum:** md5 checksum of the image
- **container_format:** bare, ovf, ...
- **created_at, updated_at, deleted_at:** creation, last update and deletion date and time
- **deleted:** whether the image has been deleted or not (True/False)
- **disk_format:** raw, qcow2, ...
- **id, name:** uuid of the image and name
- **is_public:** whether the image is public or not (True/False)
- **min_disk, min_ram:** generic recommendation for using the image
- **owner:** uuid of the owner of the image
- **protected:** whether the image can be deleted or not (True/False)
- **size:** real size of the image (not the virtual image size seen by the guest)
- **status:** active, saving, queued, killed, pending_delete, deleted

- In Glance API Version 1 it is possible to add custom properties to the images (not officially standardized)
- In Glance API Version 2 some additional custom properties have been standardized:
 - **architecture:** the CPU architecture that must be supported by the hypervisor (e.g. x86_64, arm, ppc64, ...)
 - **hypervisor_type:** xen, qemu, kvm, lxc, uml, ...
 - **os_type, os_distro, os_version:** OS information
- More information at:

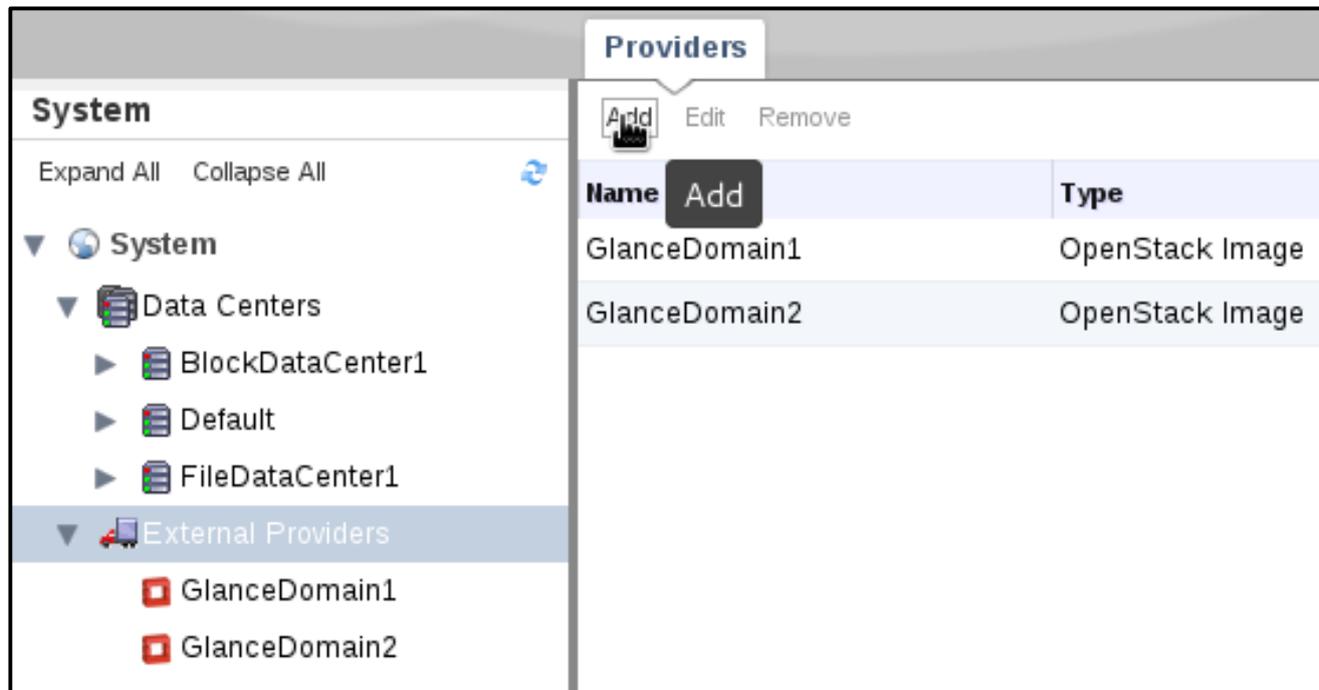
<http://docs.openstack.org/trunk/openstack-compute/admin/content/image-metadata.html>

- Import and Export single images (templates) from and to Glance for interoperability with OpenStack
- Unlimited number of backups for images and VMs
- Store ISO images in Glance and use them natively in the oVirt VMs
- Import and Export full VMs and templates supporting:
 - Complete VM definition
 - Multiple disks
 - Snapshots

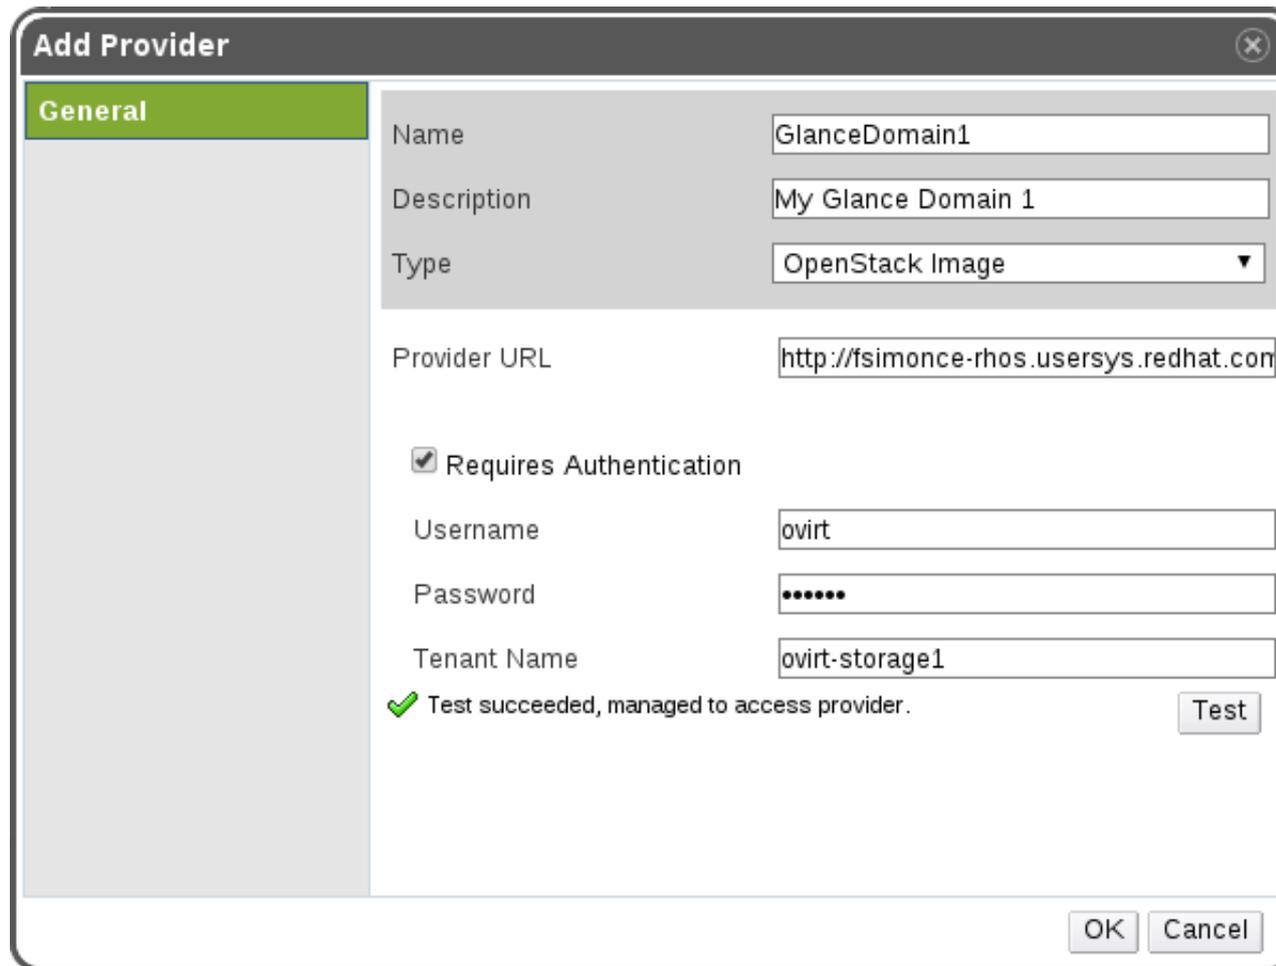
- Import and Export single images (templates) from and to Glance for interoperability with OpenStack
- Unlimited number of backups for images and VMs
- Store ISO images in Glance and use them transparently for oVirt VMs
- Import and Export full VMs and templates supporting:
 - Complete VM definition
 - Multiple disks
 - Snapshots

External Providers and Glance 1/2

- External Providers support in oVirt 3.3; oVirt is now capable to interface with some additional supported services (e.g. OpenStack Image/Network, Foreman)
- New “Providers” Tab in WebAdmin to add, edit and remove the External Providers



- Glance is the “OpenStack Image” External Provider



Add Provider

General

Name: GlanceDomain1

Description: My Glance Domain 1

Type: OpenStack Image

Provider URL: http://fsimonce-rhos.usersys.redhat.com

Requires Authentication

Username: ovirt

Password:

Tenant Name: ovirt-storage1

Test succeeded, managed to access provider.

- Glance Storage Domains are also listed in the Storage Tab
- They don't belong to a specific Data Center (Unattached)

Domain Name	Domain Type	Storage Type	Format	Cross Data-Center Status
BlockDomain1	Data	iSCSI	V3	Active
BlockDomain2	Data (Master)	iSCSI	V3	Active
ExportDomain1	Export	NFS	V1	Active
FileDomain1	Data (Master)	NFS	V3	Unknown
GlanceDomain1	Image	OpenStack Glance	V1	Unattached
GlanceDomain2	Image	OpenStack Glance	V1	Unattached
IsoDomain1	ISO	NFS	V1	Active

File Name	Type	Actual Size
Blank QCOW2 Image 20Gb (1477cd5)	Disk	< 1 GB
BlockDiskThin1 (7fdbd25)	Disk	< 1 GB

- To accomplish some operations (e.g. Glance images listing) oVirt Engine requires a direct connection to OpenStack services
- OpenStack clients are the native Python bindings for the OpenStack APIs (used to implement the command-line interfaces)
- `openstack-java-sdk` is a third-party OpenStack client API library for Java (licensed under the Apache 2 license)
- Since February 2013 (`os-java-sdk 3.0.0`) oVirt engineers have started contributing code, highlights:
 - Introduce “connectors” to reduce the number of dependencies (JBoss vs Glassfish)
 - Stabilize Keystone, Glance and Neutron (was Quantum) APIs
 - Improve Keystone authentication with Token Providers

Glance Images Discovery

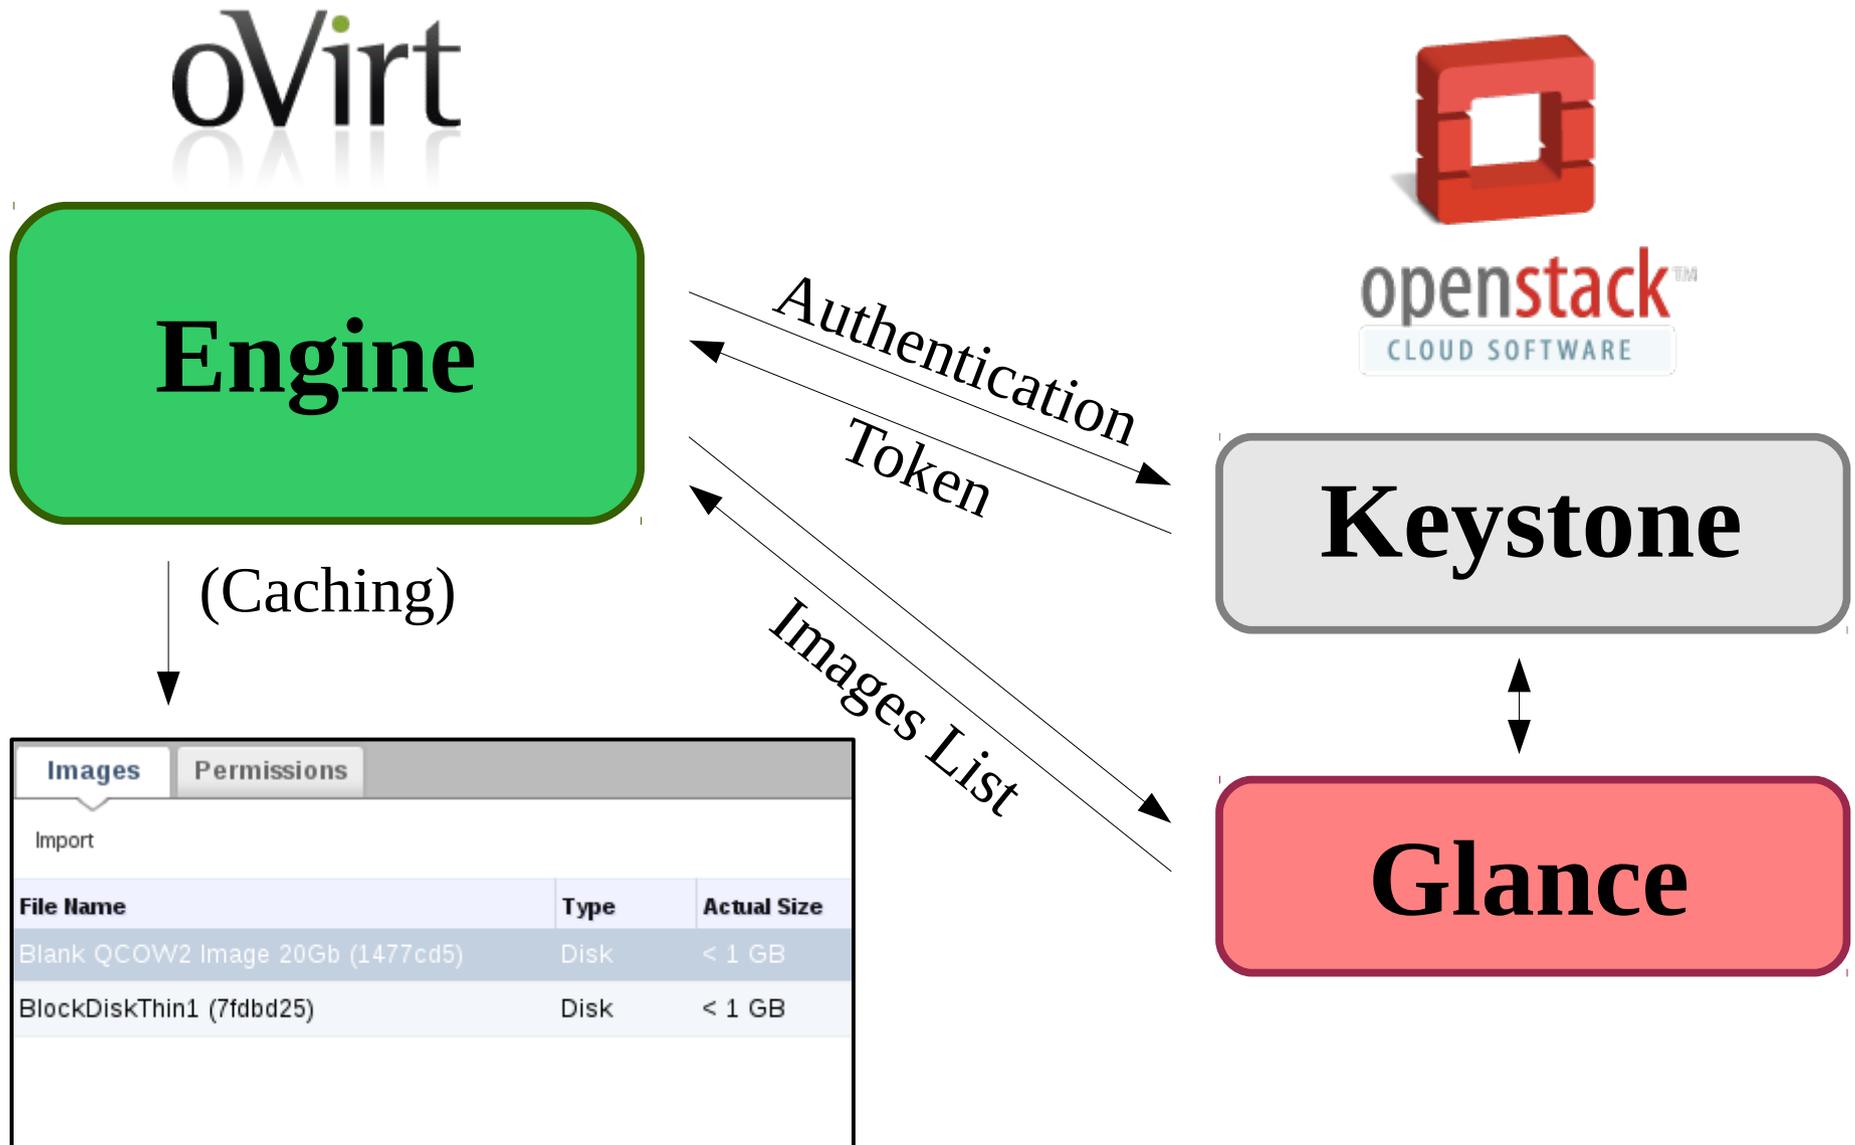
- Active Glance Images (in the known and supported formats) are listed in the Images sub-tab when a Glance domain is selected

The screenshot shows the oVirt interface with the 'Storage' tab selected. The left sidebar shows a tree view of the system hierarchy, including Data Centers and External Providers. The main area displays a table of storage domains. Below this, the 'Images' sub-tab is active, showing a list of images associated with the selected Glance domain.

Domain Name	Domain Type	Storage Type	Format	Cross Data-Center Status
▲ BlockDomain1	Data	iSCSI	V3	Active
▲ BlockDomain2	Data (Master)	iSCSI	V3	Active
▲ ExportDomain1	Export	NFS	V1	Active
▼ FileDomain1	Data (Master)	NFS	V3	Unknown
■ GlanceDomain1	Image	OpenStack Glance	V1	Unattached
■ GlanceDomain2	Image	OpenStack Glance	V1	Unattached
▲ IsoDomain1	ISO	NFS	V1	Active

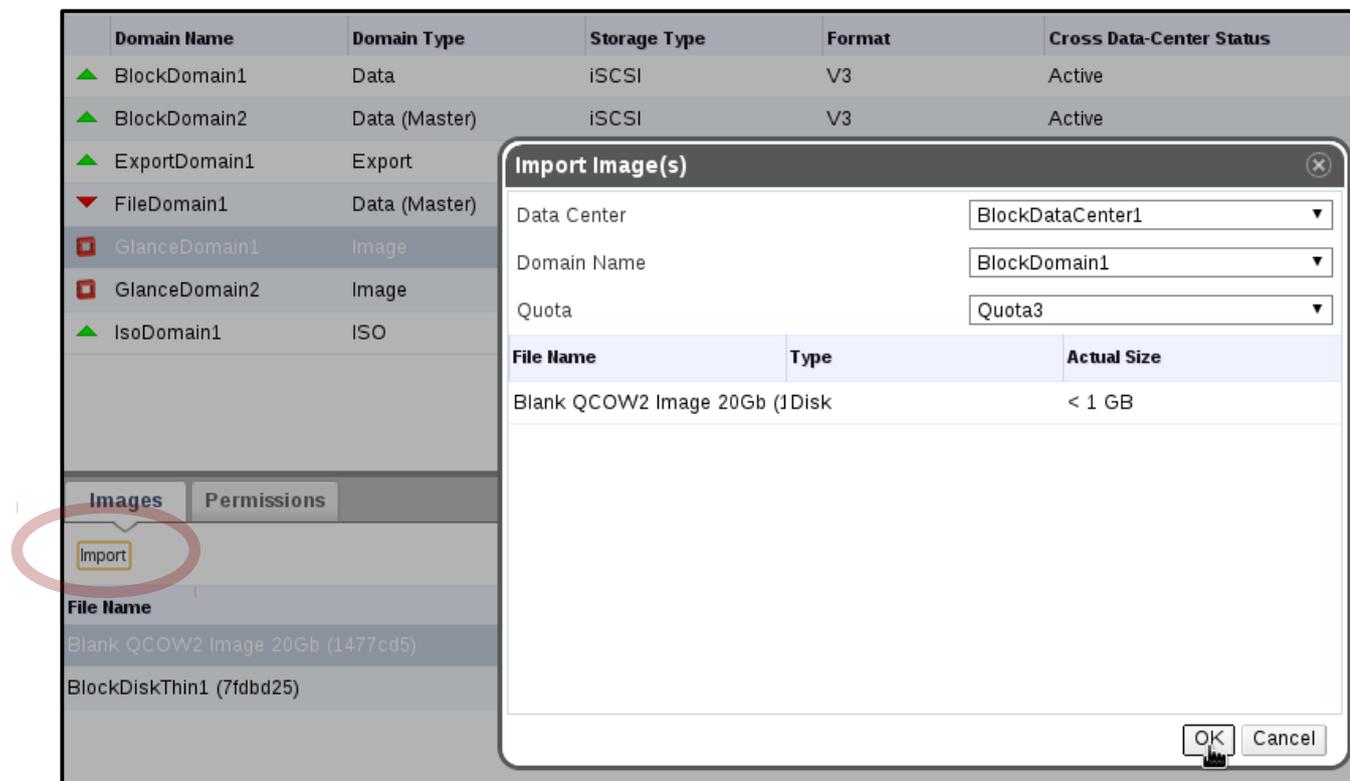
File Name	Type	Actual Size
Blank QCOW2 Image 20Gb (1477cd5)	Disk	< 1 GB
BlockDiskThin1 (7fdbd25)	Disk	< 1 GB

Glance Images Discovery Diagram

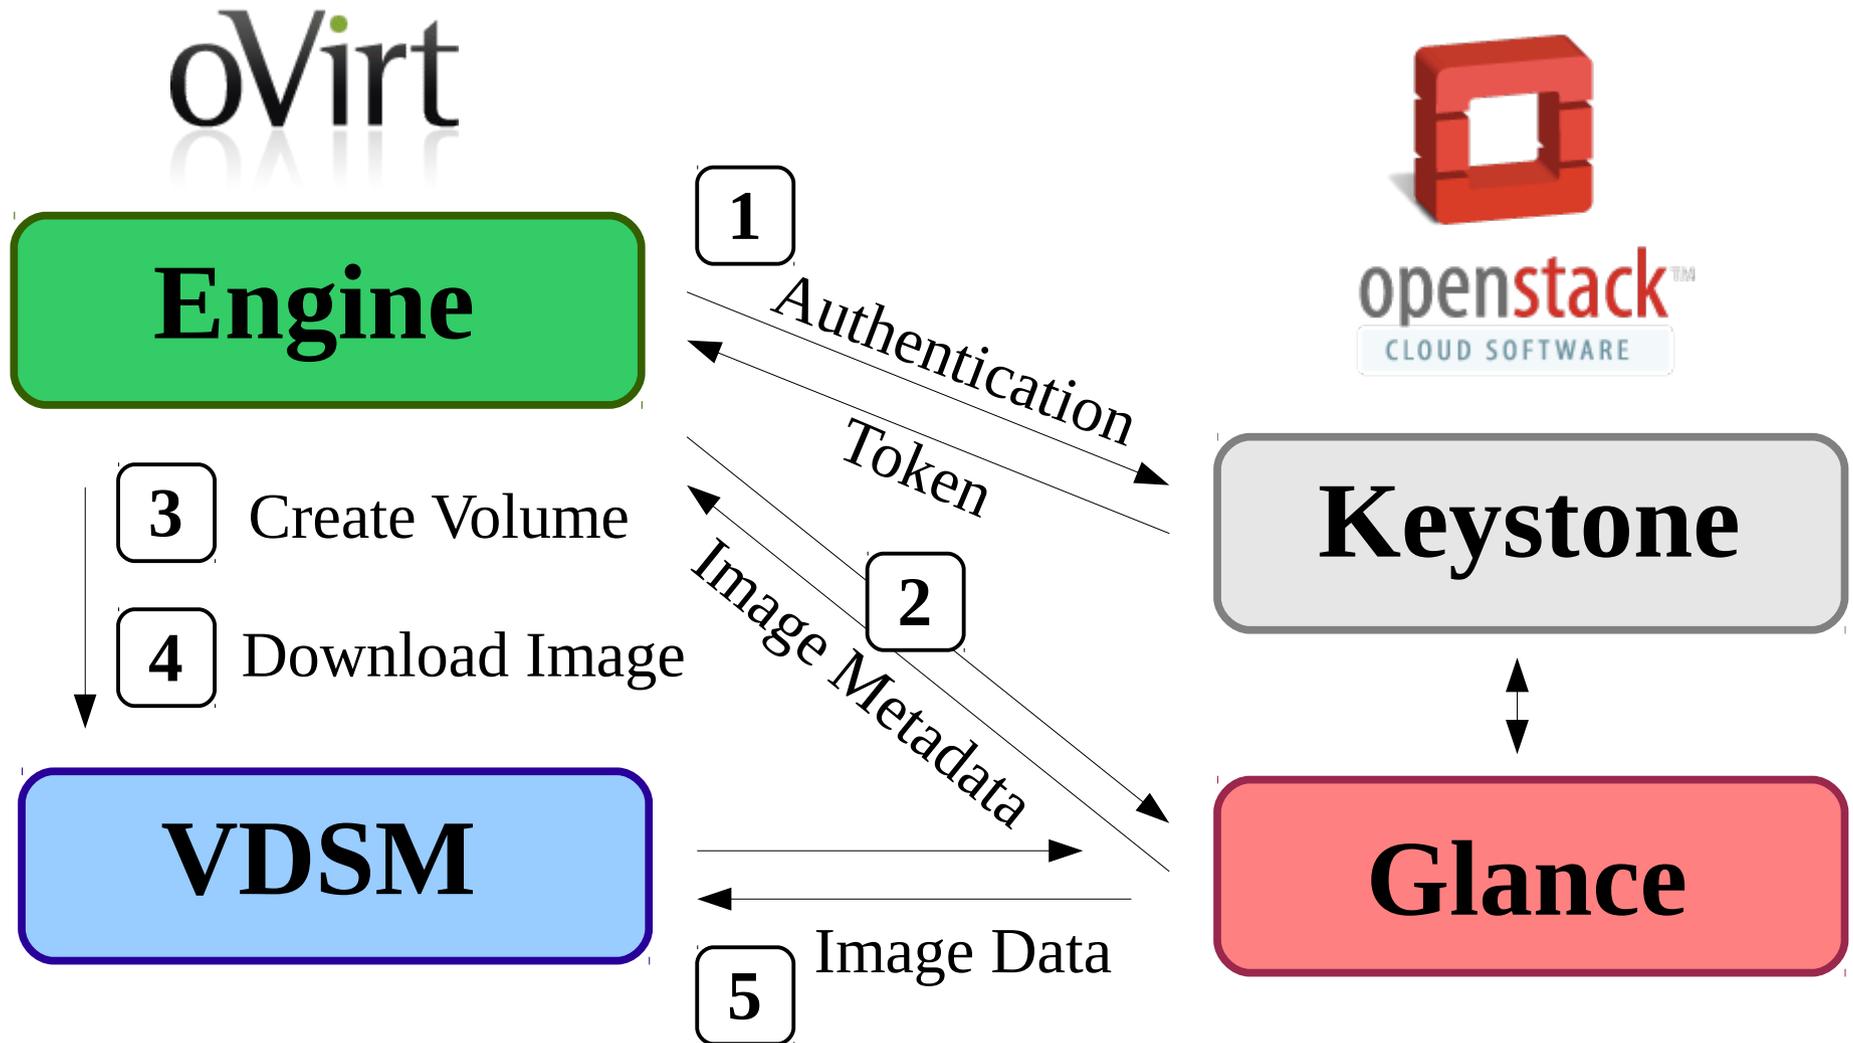


Importing Glance Images

- To import from Glance select the relevant Images and click on the “Import” button in the “Images” sub-tab
- Select the destination parameters in the pop-up dialog as the destination Data Center, Domain and Quota

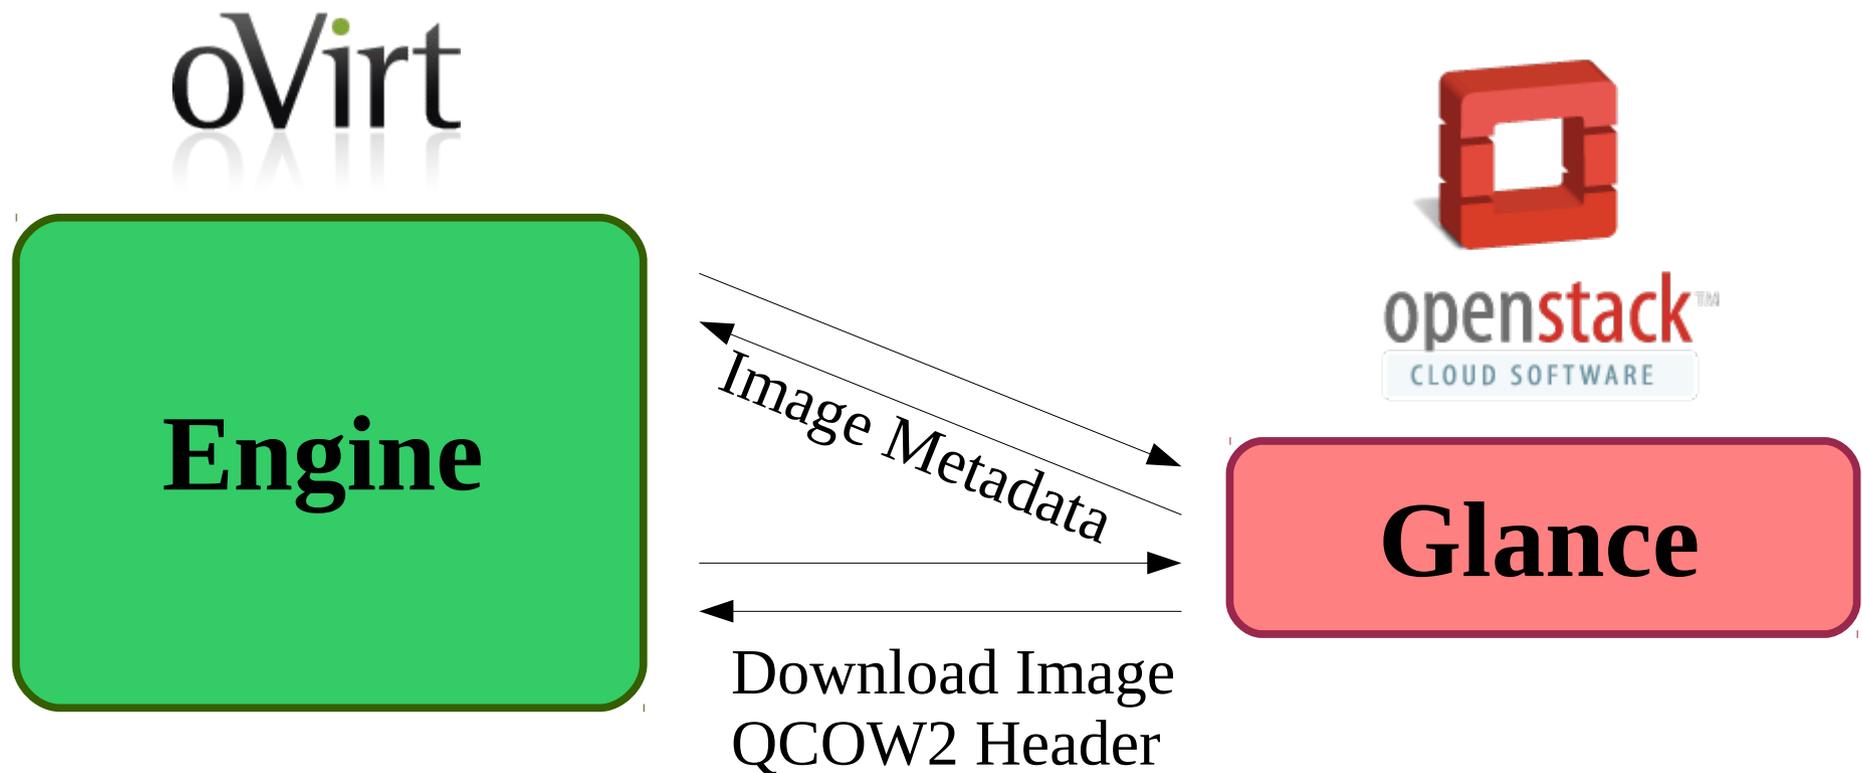


Importing Glance Images Diagram 1/2



Importing Glance Images Diagram 2/2

- Glance Metadata doesn't provide the Image Virtual Size
- When the Image is in the qcow2 format, the virtual size is discovered by downloading the image qcow2 header (72 bytes)



Imported Glance Image in oVirt

- The imported Image can be found in the “Disks” tab
- The relevant Metadata is maintained, some other parameters are set to the default (e.g. interface VirtIO) and they can be modified later
- An oVirt VM can now be instantiated with the Disk and eventually it can be transformed into a Template (automated in oVirt 3.4)

The screenshot shows the oVirt interface. At the top, a table lists disks with columns for name, UUID, VM, domain, and size. The 'GlanceDisk-1477cd5' row is highlighted with a red oval. Below the table, a navigation bar shows 'General', 'Virtual Machines', 'Storage', and 'Permissions' tabs. The 'General' tab is active, displaying details for the selected disk.

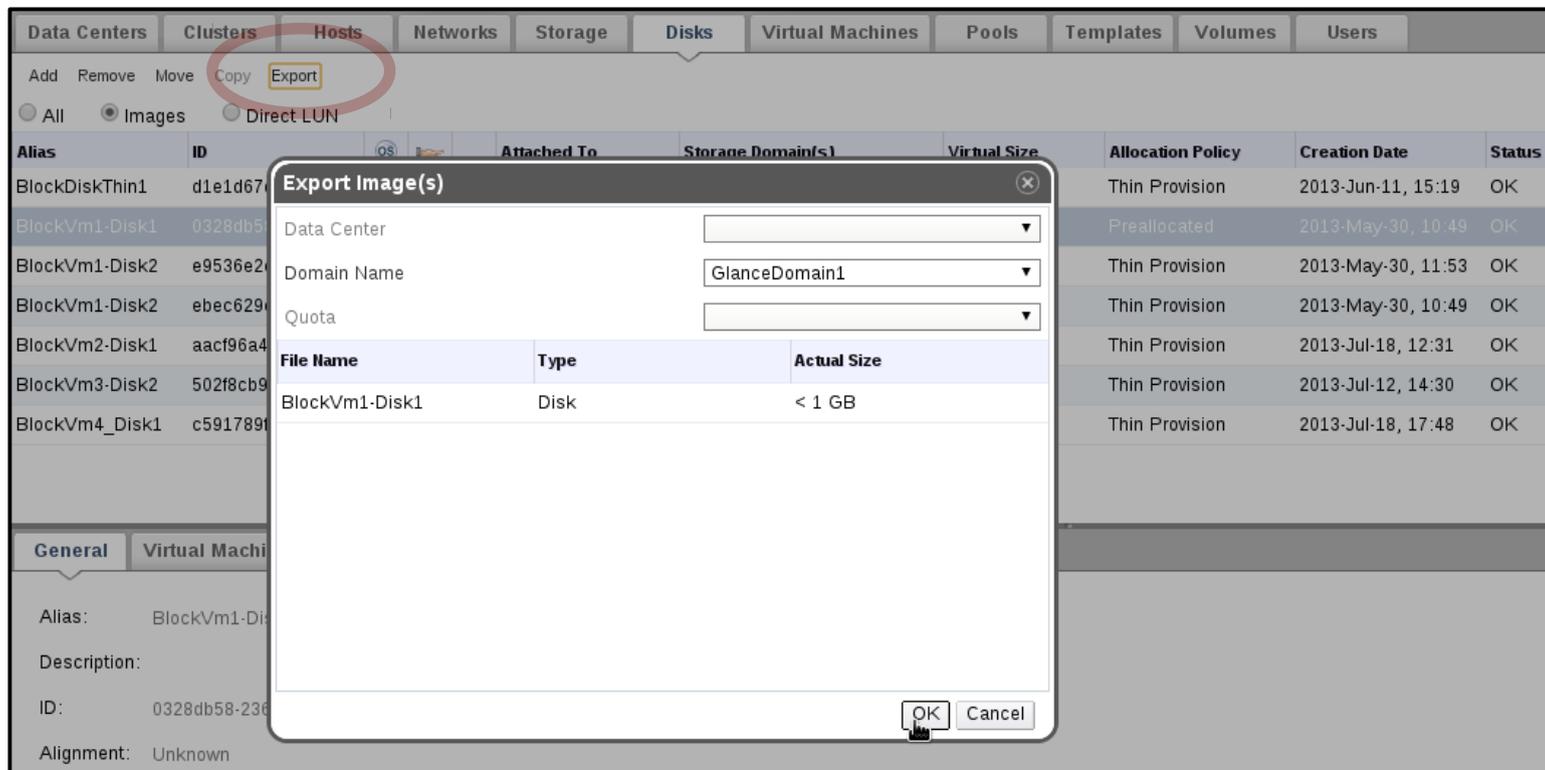
Name	UUID	VM	Domain	Size
BlockVm3-Disk2	502f8cb9-760a-4f3b-af37-6	BlockVm3	BlockDomain1	20 GB
BlockVm4-Disk1	c591789f-afd4-479e-88b6-:	BlockVm4	BlockDomain1	2 GB
GlanceDisk-1477cd5	e0e3ca0f-d998-44fa-8f70-8		BlockDomain1	20 GB

General | Virtual Machines | Storage | Permissions

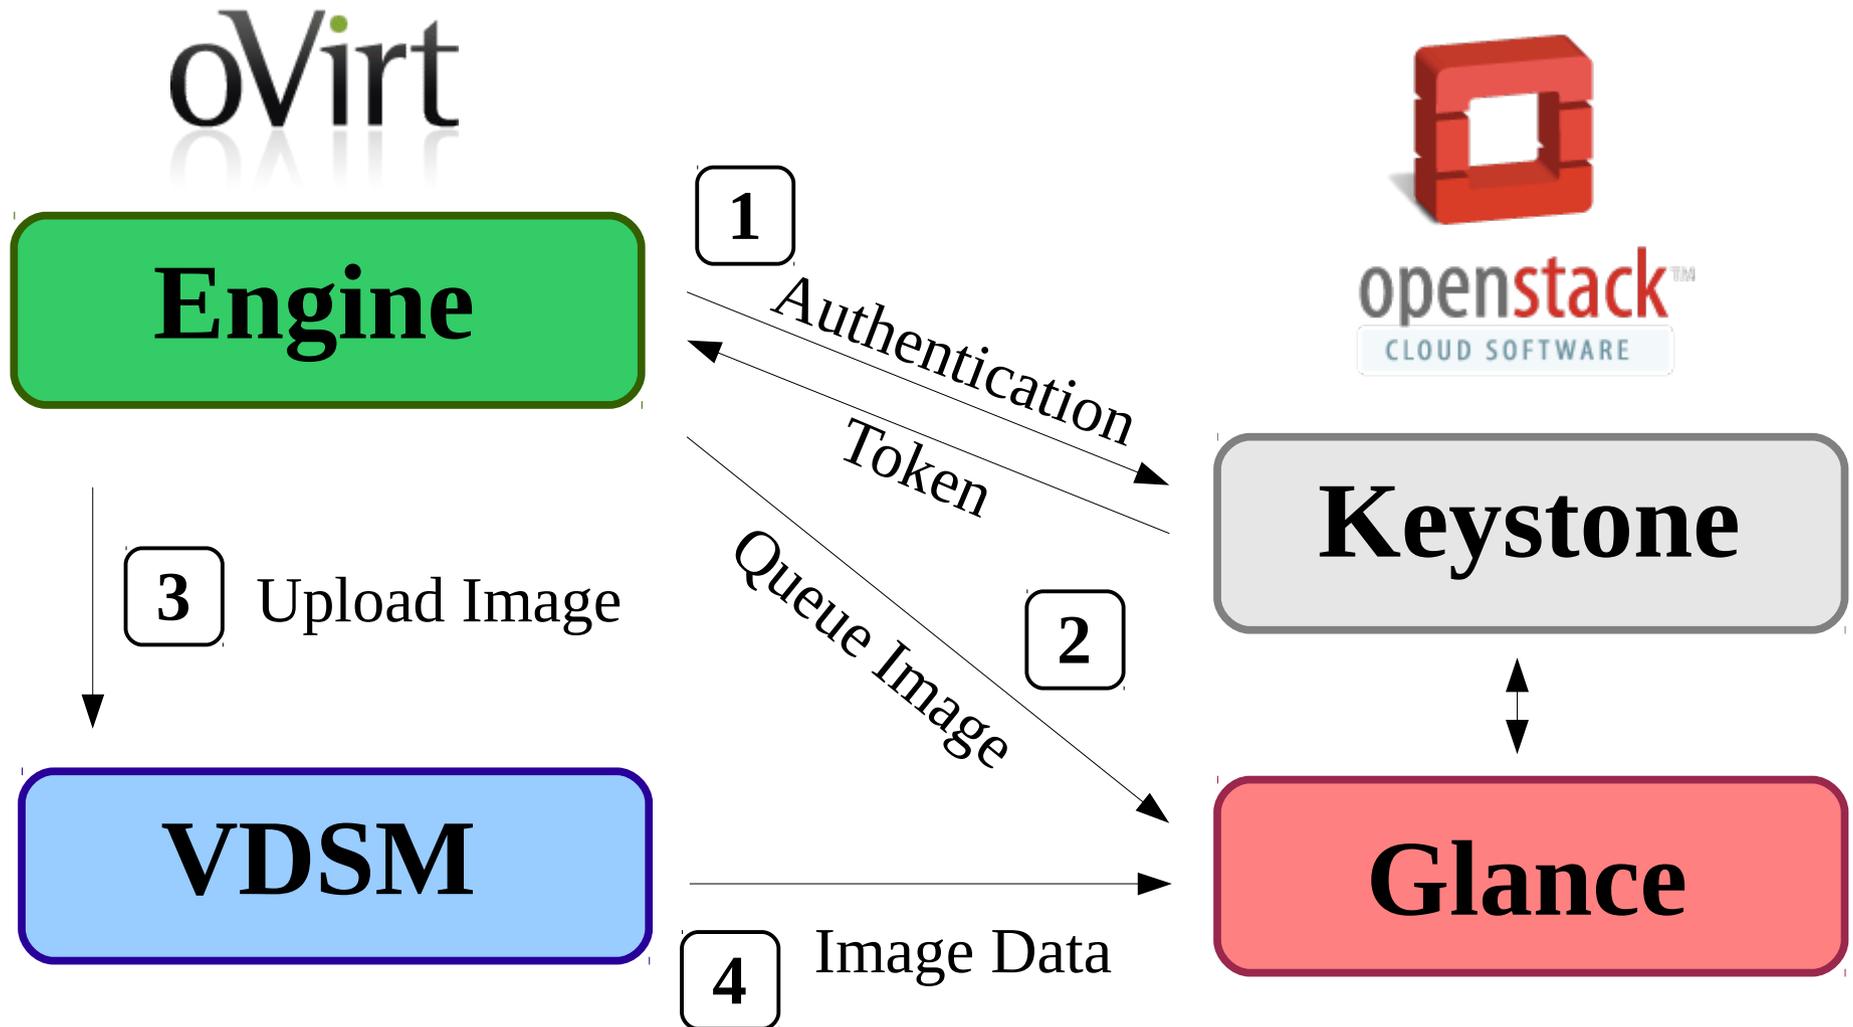
Alias: GlanceDisk-1477cd5
Description: Blank QCOW2 Image 20Gb (1477cd5)
ID: e0e3ca0f-d998-44fa-8f70-876f8c73cb47
Alignment: Unknown
Quota: Quota3

Exporting Disks to Glance

- To export to Glance select the relevant Disks and click on the “Export” button in the “Disks” tab
- Select the destination Glance Domain in the pop-up dialog



Exporting Glance Images Diagram



Exported oVirt Disk in Glance



```
# glance image-list
+-----+-----+-----+-----+
| ID           | Name                               | Disk  | Container | Size      | Status |
+-----+-----+-----+-----+
| 1477cd5a...  | Blank QCOW2 Image 20Gb           | qcow2 | bare      | 262144    | active |
| 7fdbd255...  | BlockDiskThin1                   | qcow2 | bare      | 134217728 | active |
+-----+-----+-----+-----+

# glance image-show 7fdbd255-4851-495e-8af2-300f9967c6b5
+-----+-----+
| Property      | Value                               |
+-----+-----+
| checksum      | 089092ee15690fe33ec12d426ac50e70  |
| container_format | bare                                |
| created_at    | 2013-07-24T16:41:11                |
| deleted       | False                               |
| disk_format   | qcow2                               |
| id            | 7fdbd255-4851-495e-8af2-300f9967c6b5 |
| is_public     | False                               |
| name          | BlockDiskThin1                     |
| owner         | 5b7d836142f9444a9e2bd3494669c831  |
| protected     | False                               |
| size          | 134217728                           |
| status        | active                              |
| updated_at    | 2013-07-24T17:07:15                |
+-----+-----+
```

- New API to download and upload images using predefined methods

```
downloadImage(methodArgs, spUUID, sdUUID, imgUUID, volUUID=None)
```

```
uploadImage(methodArgs, spUUID, sdUUID, imgUUID, volUUID=None)
```

- HTTP method implemented and fully compatible with Glance

```
methodArgs = {  
    'method': 'http',  
    'url': 'http://source-or-destination/path/to/image',  
    'headers': {'X-My-Header', 'MyHeaderValue'},  
}
```

- X-Auth-Token provided by Keystone is passed as one of the headers

VDSM Client Download Examples

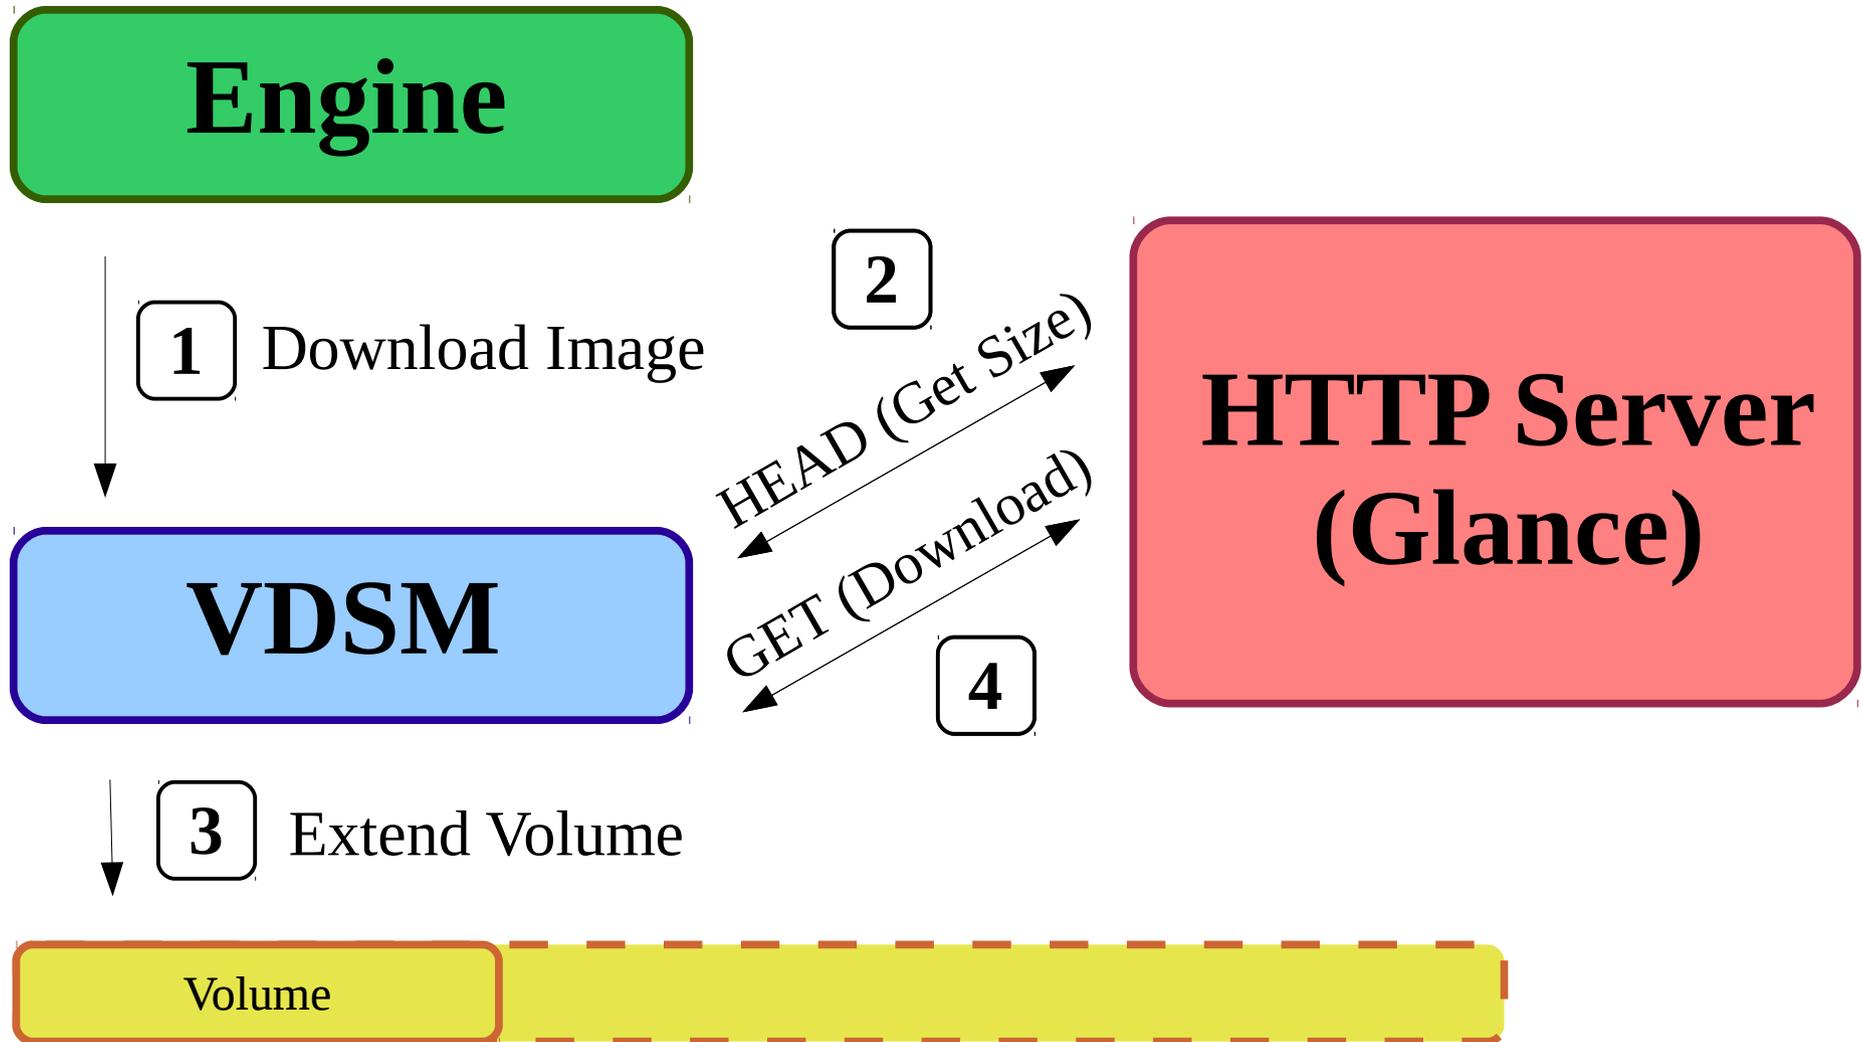


- VDSM gained support for Glance:

```
$ vdsClient 0 downloadImage http '{
  "method": "http",
  "url": "http://glance:9292/v1/images/320897a4-6924-4b31-a256-f51c97d15d95",
  "headers": {
    "X-Auth-Token": "1ad97ccf22aa47a89e082b6491fb4b70",
  }}' \
0b11dedd-6134-4775-8ccb-6a935d16320c d26915e8-9049-43a3-ba74-e403730875dc \
11400413-b99f-4eae-a0dd-24bd347bac18 16e71d6c-bf73-4e0f-a49d-8df45967d28c
```

- And for all other HTTP servers as well:

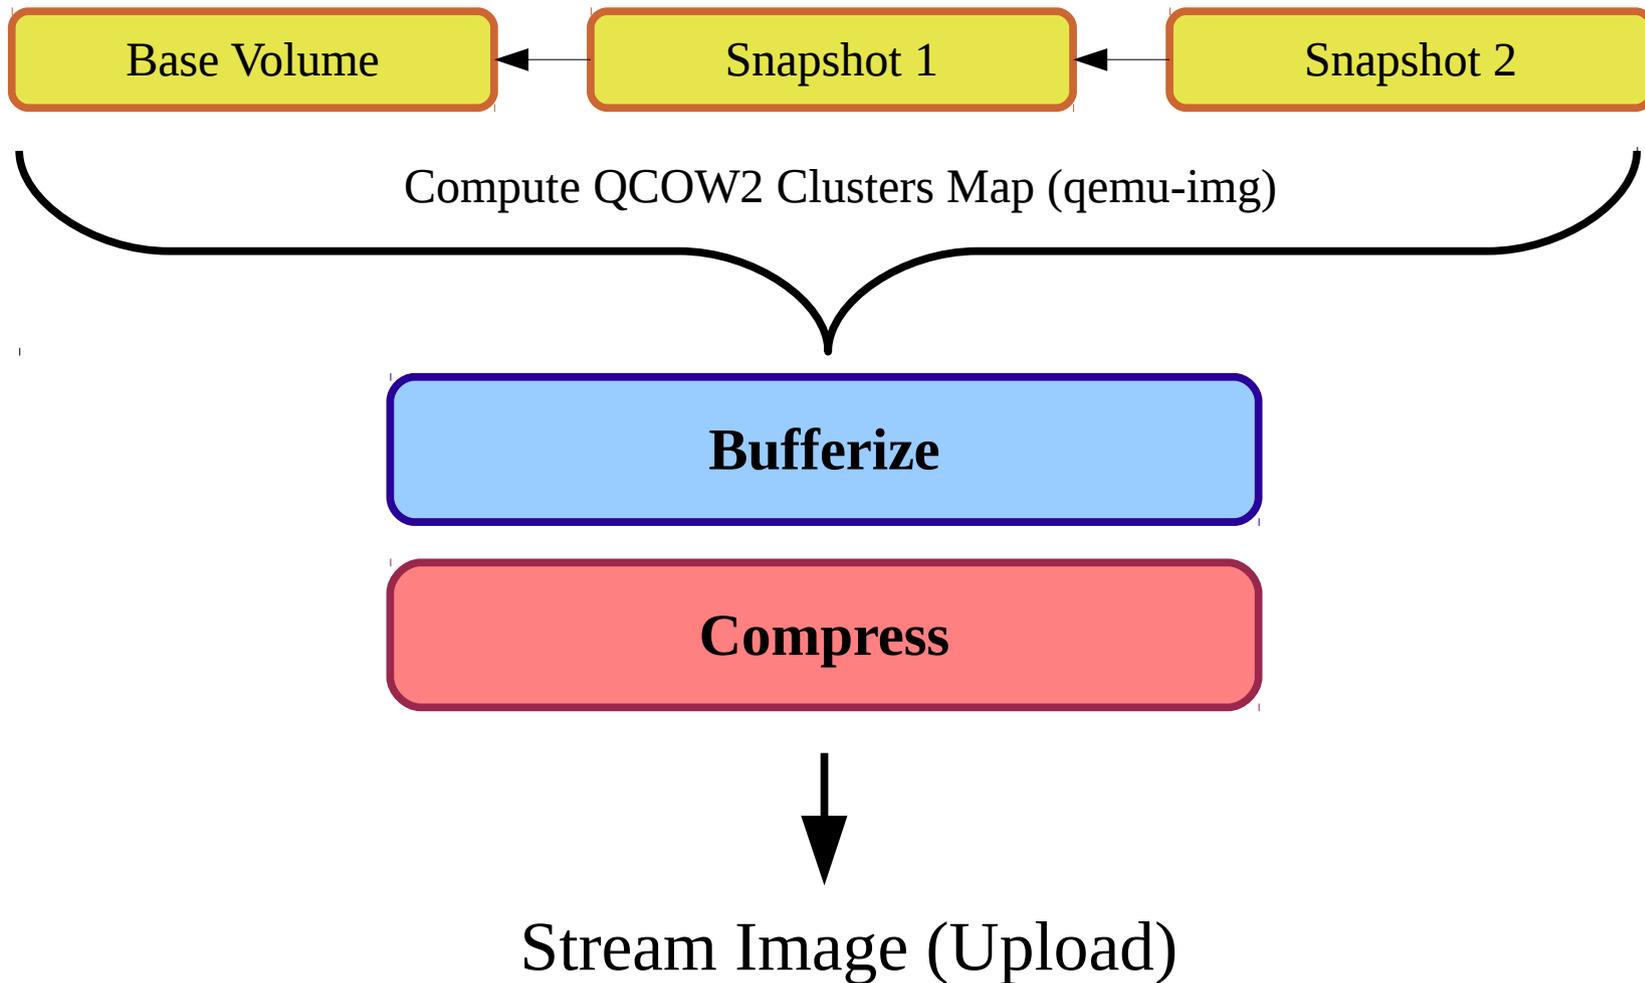
```
$ vdsClient 0 downloadImage http '{
  "method": "http",
  "url": "http://download.fedoraproject.org/.../Fedora-x86_64-19-sda.qcow2",
}' \
0b11dedd-6134-4775-8ccb-6a935d16320c d26915e8-9049-43a3-ba74-e403730875dc \
11400413-b99f-4eae-a0dd-24bd347bac18 16e71d6c-bf73-4e0f-a49d-8df45967d28c
```



- It's not currently possible to export images with multiple volumes:
 - Disks cannot have snapshots
 - Disks cannot be Thin-Provisioned on a Template
- There is no live VM disk export, VM should be down (but uploading Templates is always possible)
- No sparseness support (on file domains) for the raw format, imported images are always preallocated
- No resume option for partially downloaded images

Glance – Future Improvements

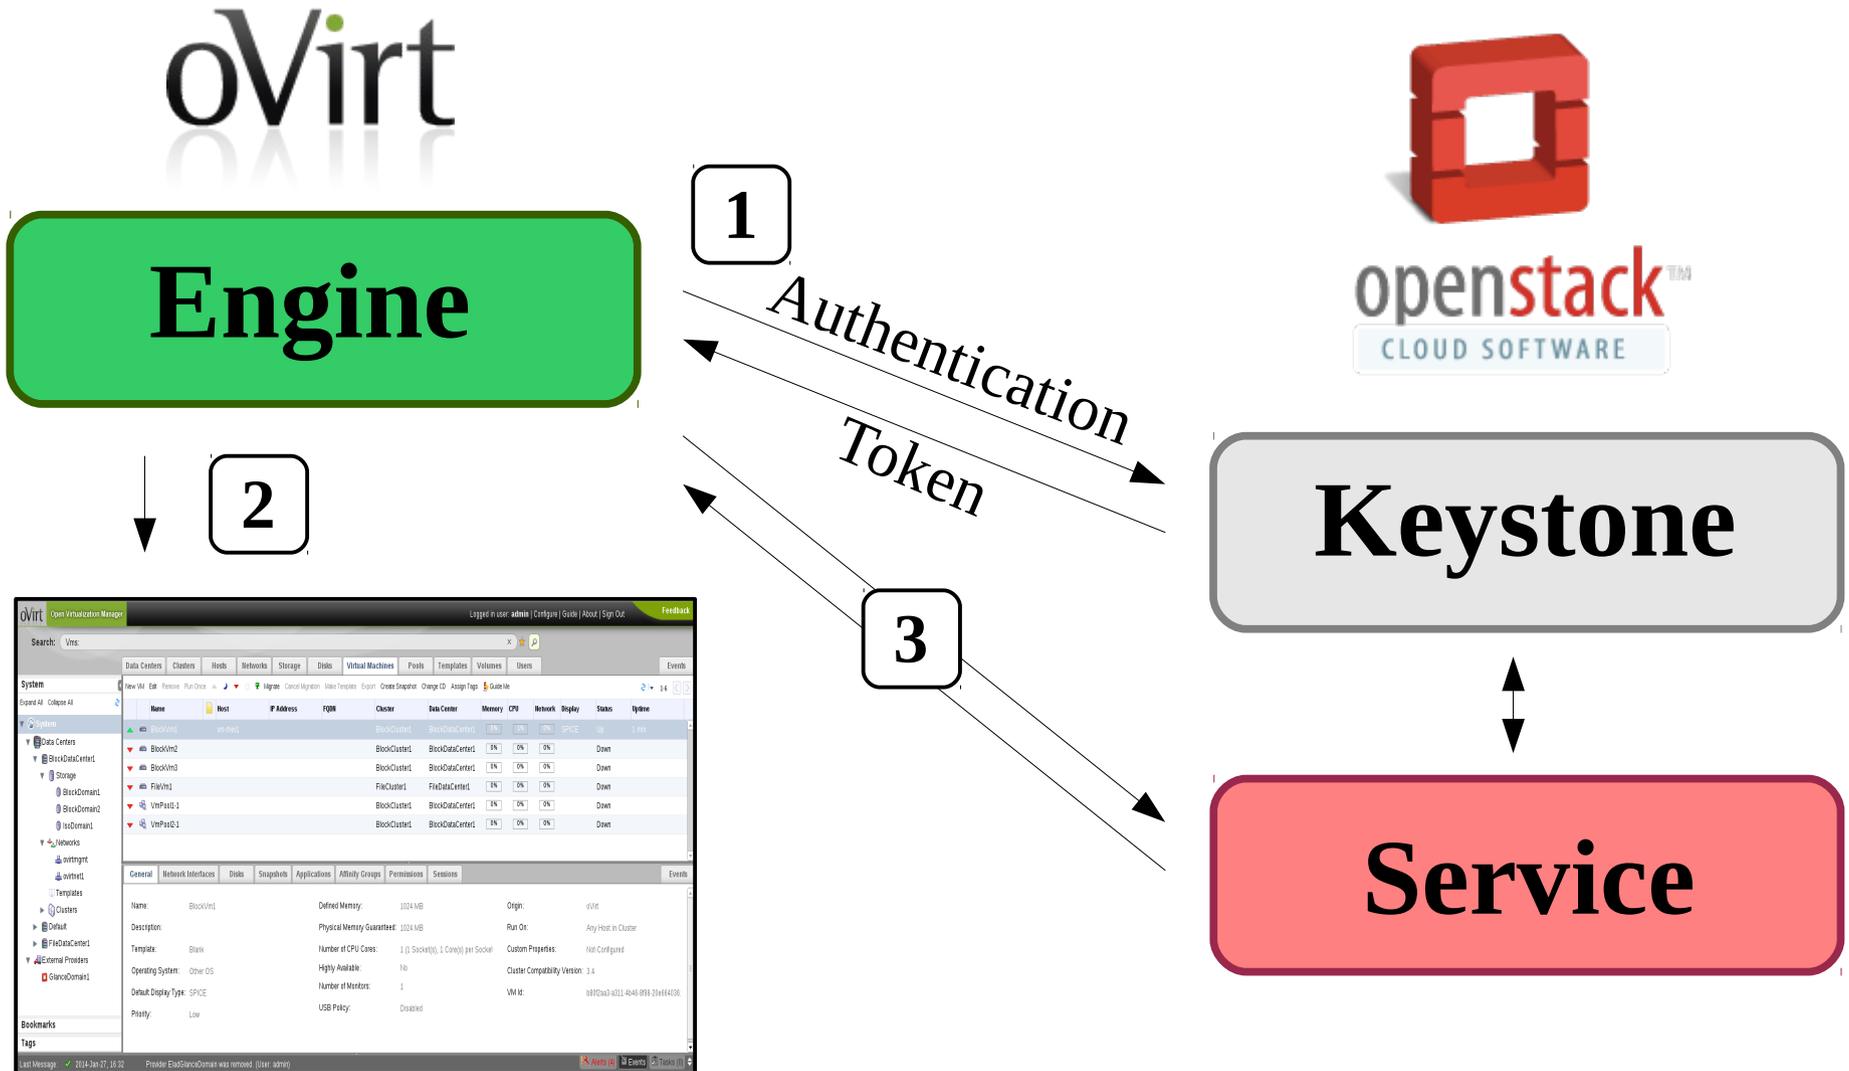
- **oVirt-Engine:** additional simple operations (e.g. Delete Image)
- **VDSM:** volume chain online squash (to support disks with snapshots)
- **VDSM/Glance:** upload Image compression (and sparseness)
- **oVirt-Engine:** integrated upload and download via WebAdmin and UserPortal
- **oVirt-Engine/VDSM:** Store ISO images in Glance and use them natively in the oVirt VMs
- **oVirt-Engine/VDSM/OpenStack:** Import and Export full VMs and Templates



- oVirt 3.4 introduced a large refactoring for Authentication and Directory services
 - Authentication confirms the identity of a User
 - Directory Services provide a set of information about Users
- Authentication and Directory services are mix-and-match (e.g. Kerberos Authentication with LDAP Directory Service)
- A Keystone Authentication backend has been sketched to verify the new Authentication/Directory implementation (anyone from the community is welcome to pick up this work and complete it)

- Consolidate Authentication in Keystone (oVirt and OpenStack data centers)
- Associate Keystone Users (and Tenants) to oVirt entities and assign specific Permissions and Roles
- Reuse the oVirt Keystone Identities when interfacing with other OpenStack services (glance, neutron, cinder, etc.)

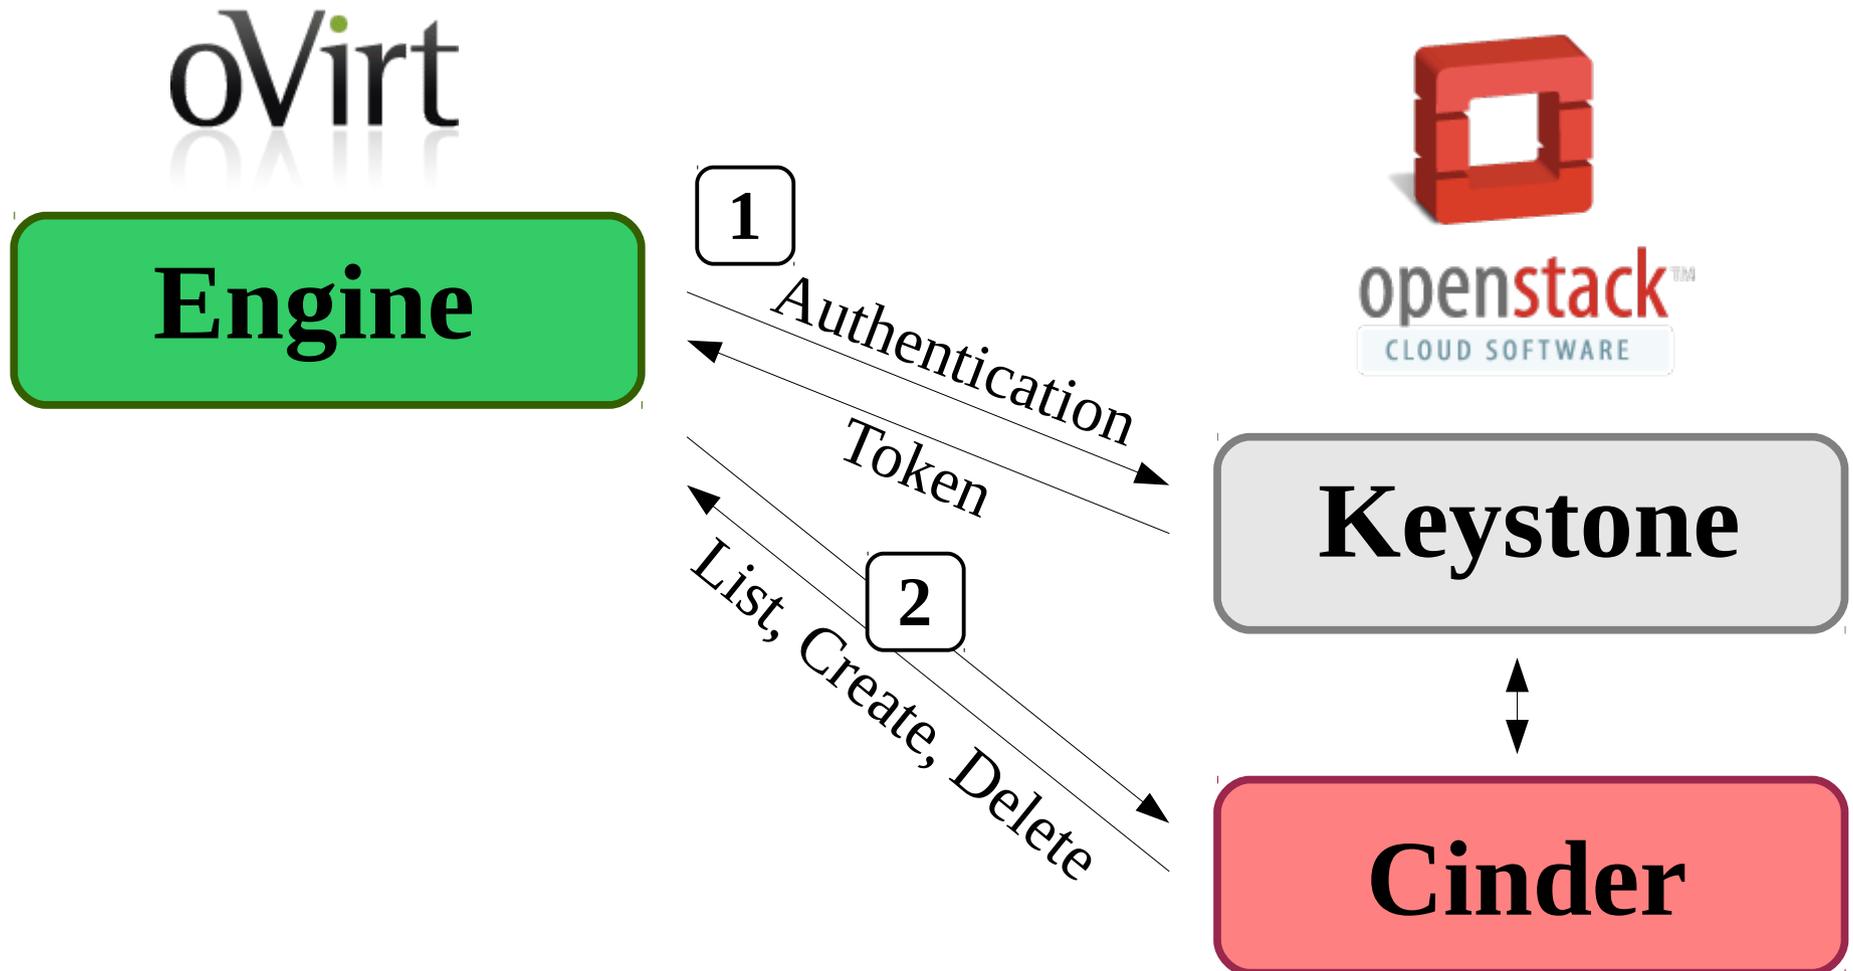
oVirt Authentication using Keystone



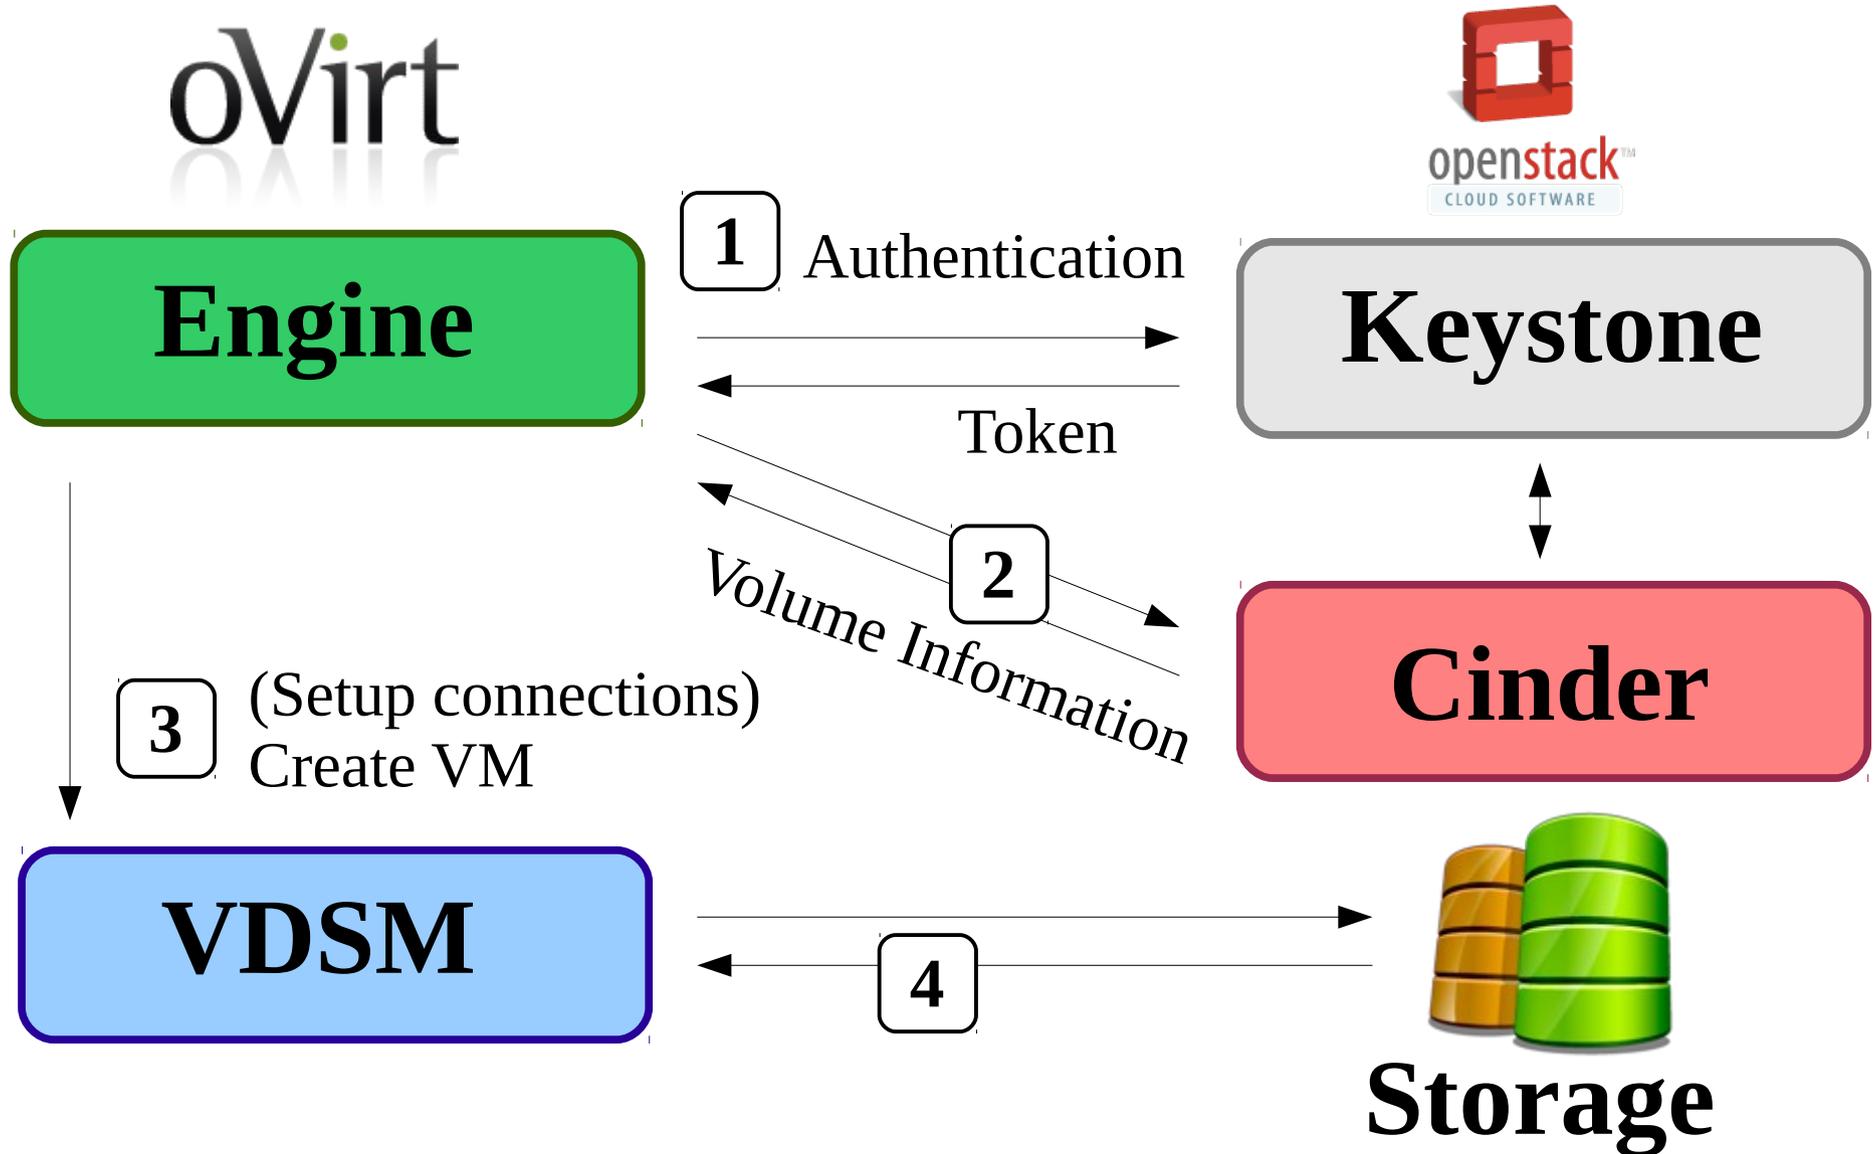
- Provides an infrastructure for managing volumes in OpenStack (since the Folsom)
- Volumes are persistent R/W block storage devices most commonly attached to the Compute Node through iSCSI
- Volumes have a lifecycle independent of VM instances
- Supports Snapshots (a read-only point in time copy of a volume)
- Snapshots can be used to create new volumes (and Glance images)
- Backups are archived in OpenStack Object Storage (Swift)

- Import and Export oVirt Images from and to Cinder
- Share Cinder Volumes for interoperability with OpenStack
- Leverage the Cinder Volume drivers for the supported storage appliances (IBM, NetApp, Dell, EMC)
- Gain storage offloading capabilities (Cinder Volumes)
- Unify snapshots and backup strategies in your data centers (oVirt/OpenStack)

Managing Cinder Volumes



Using Cinder Volumes



Useful Links and Mailing Lists

- **Work in Progress**
 - glance.ovirt.org (public glance images library)
- **Useful Links**
 - http://www.ovirt.org/Features/Glance_Integration
 - <http://www.openstack.org>
- **Mailing lists**
 - arch@ovirt.org users@ovirt.org
 - vdsm-devel@lists.fedorahosted.org
 - engine-devel@ovirt.org
- **IRC**
 - #ovirt on OFTC
 - #vdsdm on Freenode