

CloudMan and VMIC projects overview

Ulrich Schwickerath
Belmiro Moreira
Jinny Chien (ASGC)
Vineet Sharma (BARC)

- Developing these two projects in collaboration with other institutes
- These are not CERN specific projects
 - The implementation is generic for adoption by others sites

CloudMan

- BARC - The Bhabha Atomic Research Centre

VMIC

- ASGC - Academia Sinica Grid Computing

- **BARC**
 - The Bhabha Atomic Research Centre (BARC) is India's primary nuclear research facility based in Mumbai. It has a number of nuclear reactors, all of which are used for India's nuclear power and research programme.



What is CloudMan?

- Allows automated and centralized resource configuration at a high level (“resource shares”)
- It is a high level and generic graphical management portal for IT resources
 - Graphical web front-end
 - Pluggable back-ends
- Interacts with resource allocation tools

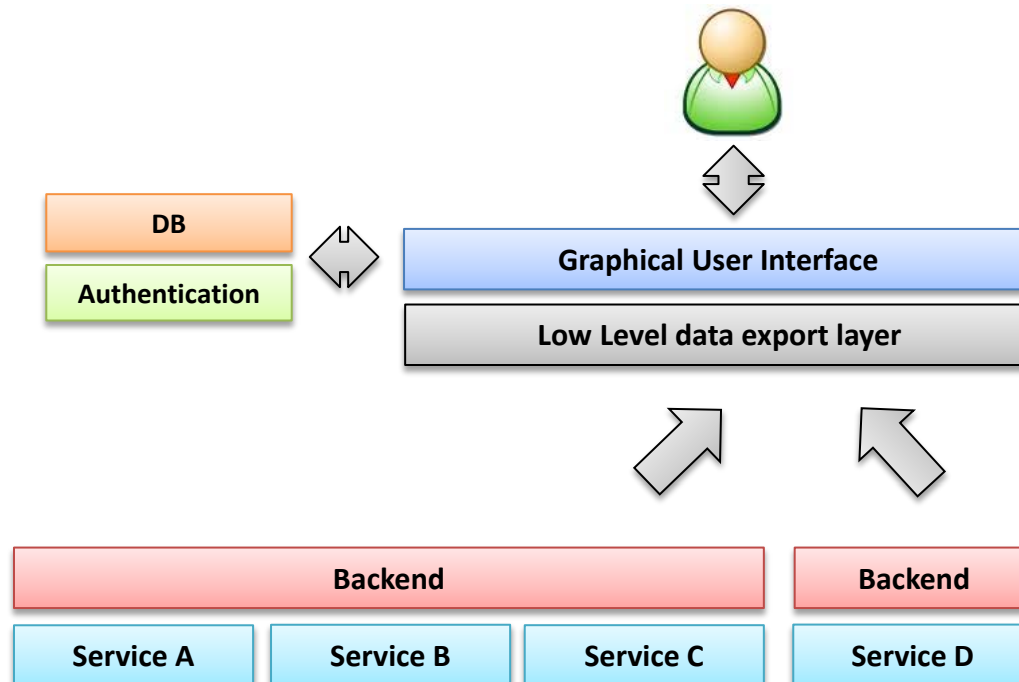
- LSFweb - used by CERN since 2008 to allocate resource shares between different groups

The screenshot shows the LSF Web Group list interface. The table displays the following data:

Name	Share	Used Share (%)	# of Subgroups	
u_ALEPH	1 kS12K (public)	62%	2	
u_ALICE	13500 kS12K (public)	100%	4	
u_AMS	1000 kS12K (public)	100%	2	
u_AMSP	1 kS12K (public)	100%	1	
u_ATLAS	7500 kS12K (public)	100%	11	
u_ATLASCAT	20 kS12K (public)	100%	1	
u_ATLASCSTORACL	1 kS12K (public)	13%	13	
u_ATLDEDICATED	1 kS12K (public)	11%	11	
u_C2	10 kS12K (public)	100%	3	
u_CAST	50 kS12K (public)	0%	0	
u_CMS	14500 kS12K (public)	97%	9	

Resource shares distribution example


- CloudMan – architecture



- Region – describes the physical location of resources



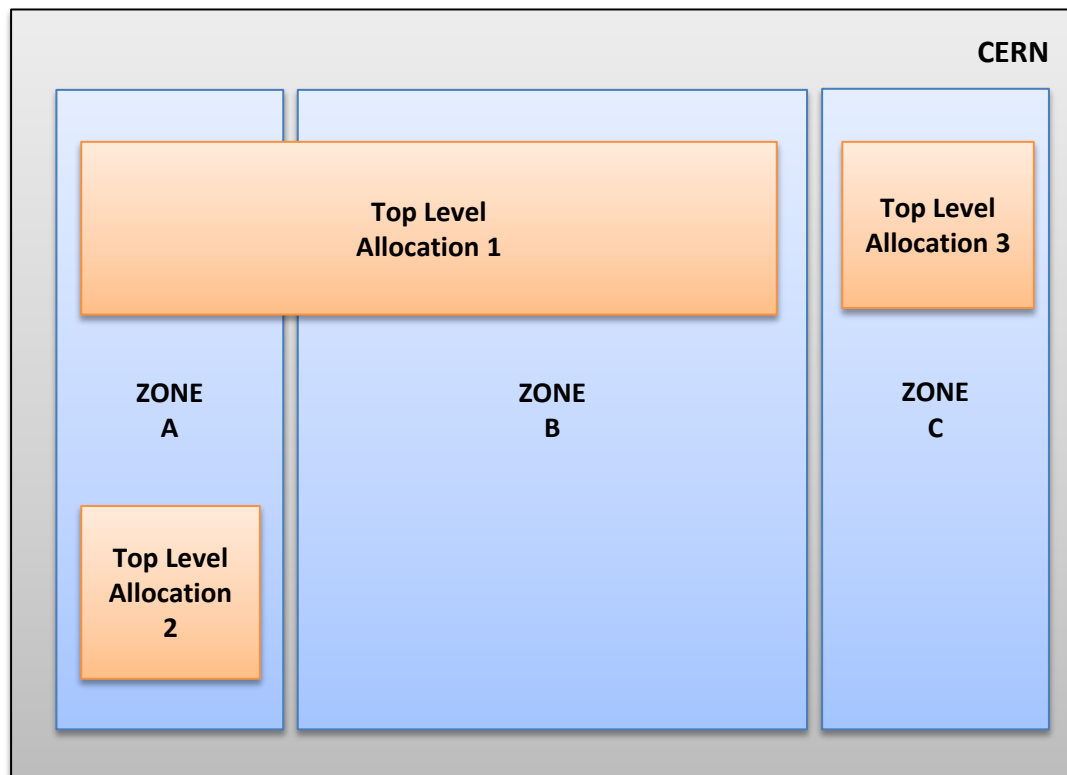
- Zone – describes a set of resources in a specific region



ZONE A
Total capacity:
- HEPSPC06: 5000
- Memory: 2000 GB
- Storage: 250000 GB

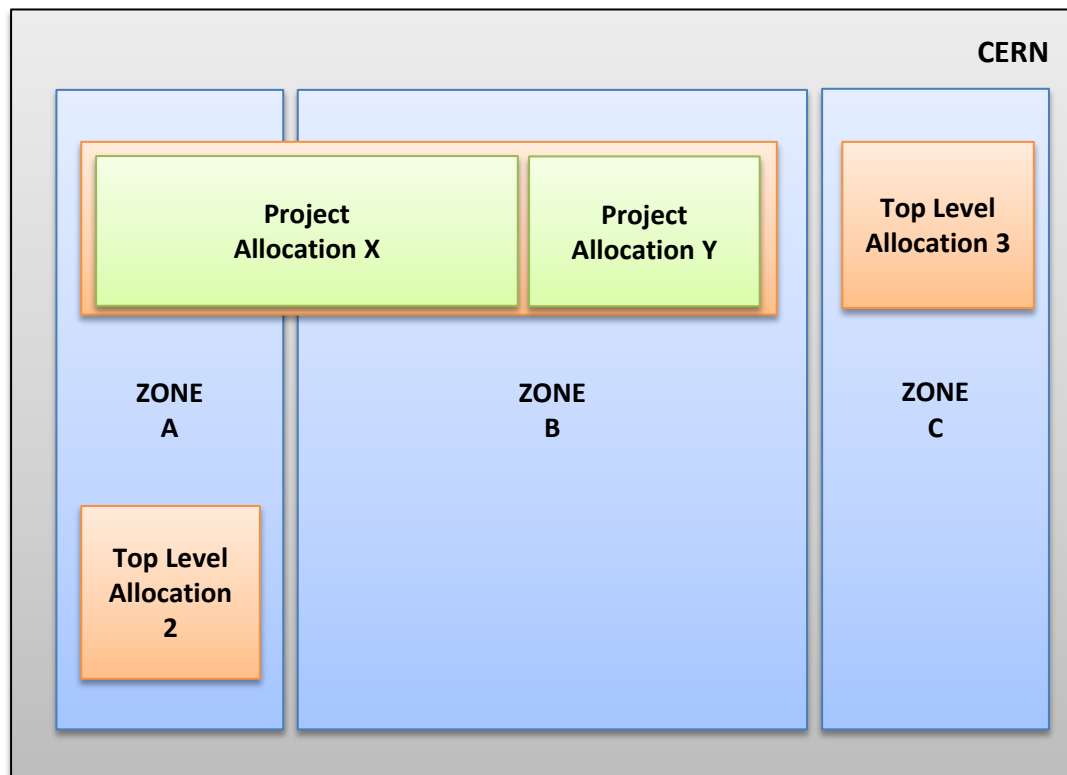
Allowed resources:
- Cern.vm.small
- Cern.vm.medium

- Top Level Allocation – Share of resources from one/different zones allocated to a group of users



Top Level Allocation 1
Allocated capacity: (zone A)
- HEPSPROC6: 2500
- Memory: 1000 GB
- Storage: 100000 GB
Allocated capacity: (zone B)
- HEPSPROC6: 4000
- Memory: 1500 GB
- Storage: 50000 GB

- Project Allocation – Share of resources allocated to a particular project




Project Allocation X

- Project: Batch
- HEPSP06: 5000
- Memory: 3000 GB
- Storage: 100000 GB
- Cern.vm.small

Project Allocation Y


- Project: EC2
- HEPSP06: 4000
- Memory: 1500 GB
- Storage: 50000 GB
- Cern.vm.small
- Cern.vm.medium

- Group Allocation – For each project the share of resources for each group of users



- A group allocation is associated to a project allocation
- Resource shares are defined using the same characteristics

- Group Allocation – For each project the share of resources for each group of users



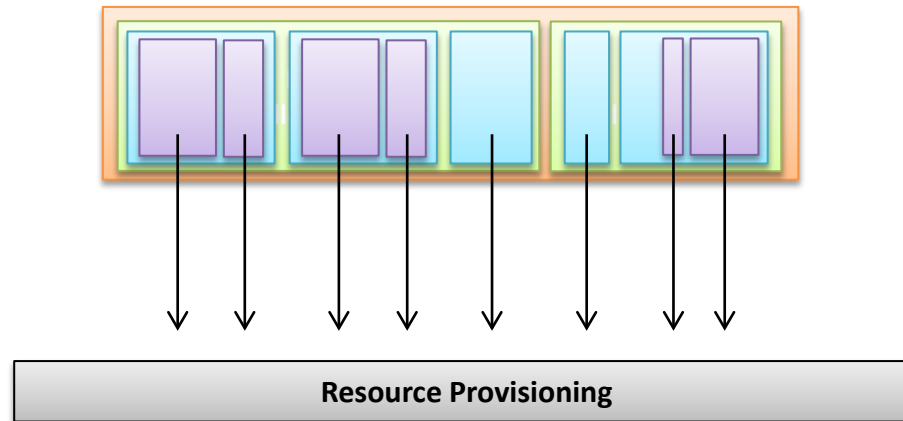
- A group allocation is associated to a project allocation
- Resource shares are defined using the same characteristics
- A group allocation can be nested

- Group Allocation – For each project the share of resources for each group of users



- A group allocation is associated to a project allocation
- Resource shares are defined using the same characteristics
- A group allocation can be nested

- Resource provisioning



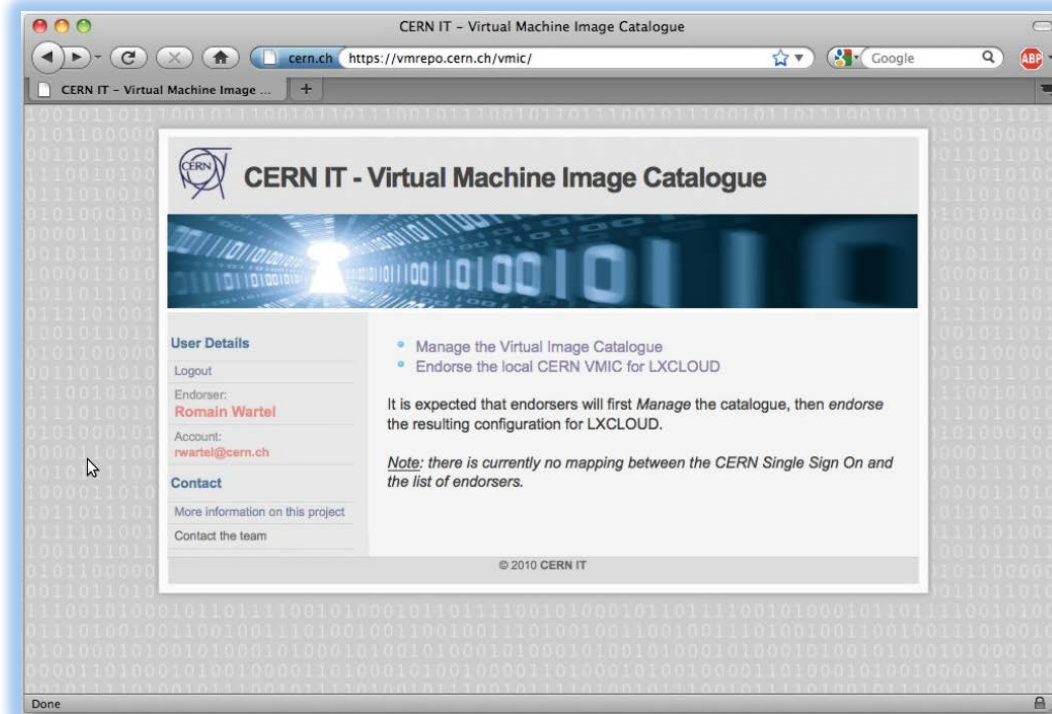
- CloudMan
 - High level and generic graphical management portal for IT resources
 - Adapted to the cloud computing concept
 - Different levels of resource allocation shares
 - Implementation started

- ASGC – Academia Sinica Grid Computing
 - It supports research activities in a wide variety of disciplines, ranging from mathematical and physical sciences, to life sciences, computing, and to humanities and social sciences

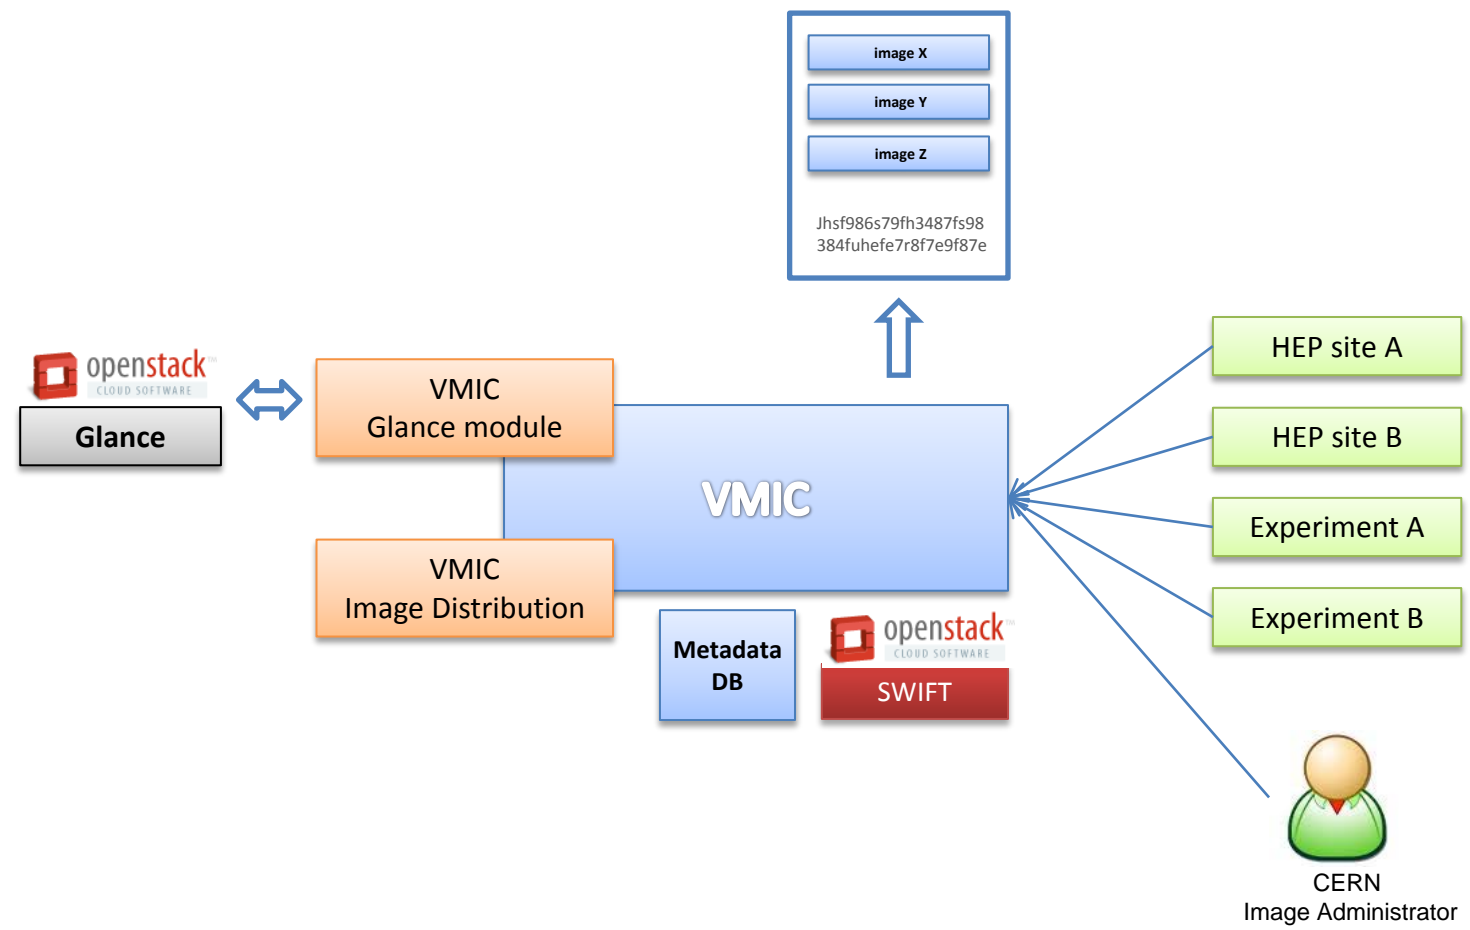


- What is VMIC?
 - VMIC – Virtual Machine Image Catalogue
 - VMIC is a catalogue of virtual machine images
 - Can subscribe images list from other sites (HEPiX specifications)
 - Can easily export an image list for other sites
 - Virtual machine images files and metadata are preserved for traceability
 - Interaction with Glance and BitTorrent image distribution

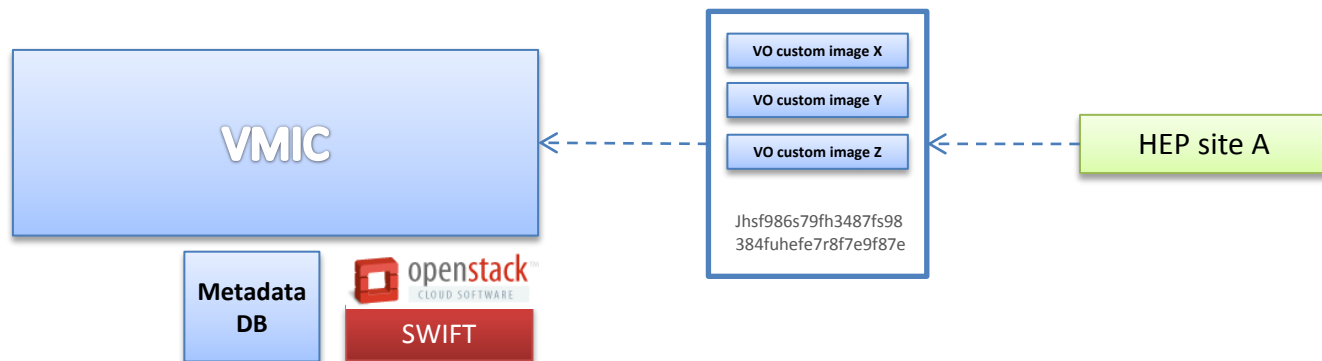
- VMIC – virtual machine image catalogue



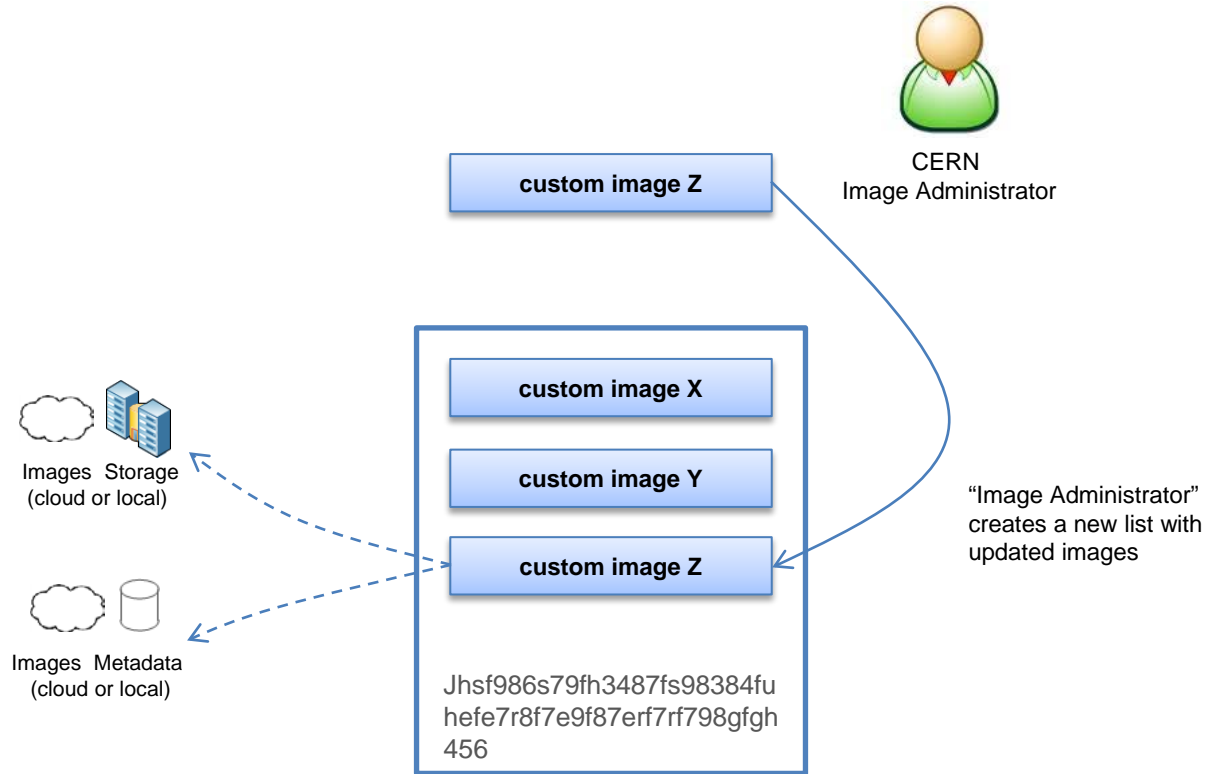
VMIC – architecture and concepts



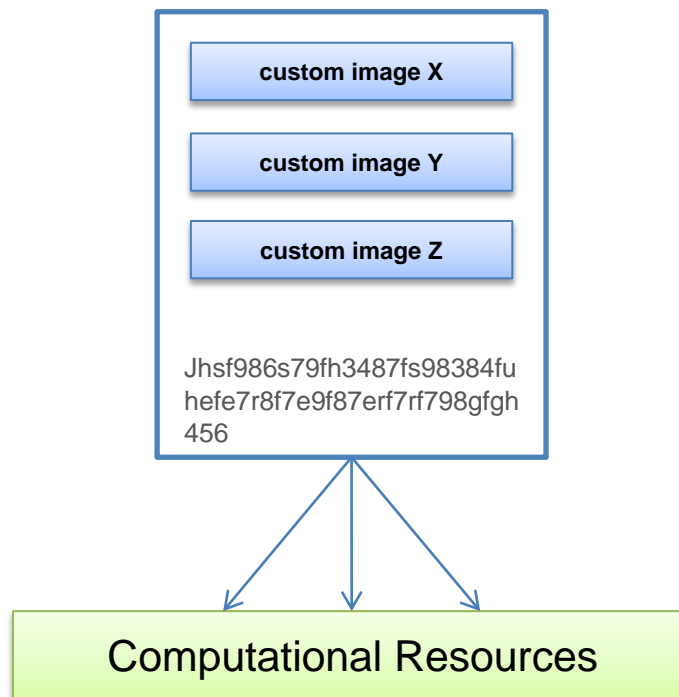
- Import images from other sites
 - Only one image
 - Subscribe an image list



Publish image list for other sites



Module – BitTorrent Image distribution

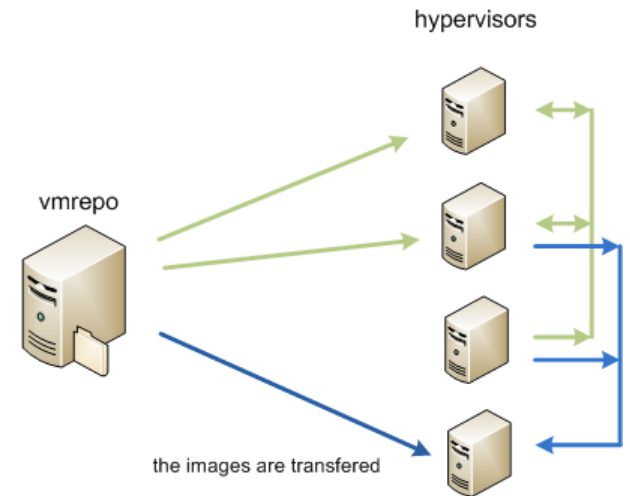


- Image Distribution

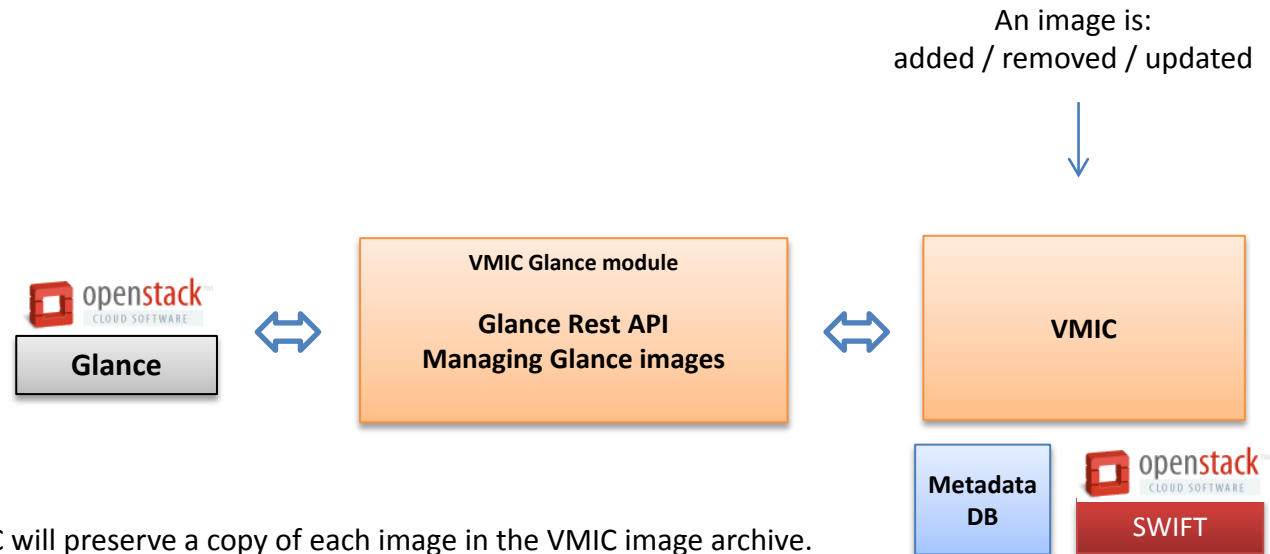
- BitTorrent protocol for image distribution

- Different steps

- The compute node query VMIC for the image list
- The image integrity is verified
- Check which images are for the node
- Download the torrent files for the desired images
- rtorrent downloads the images
- Image management software in the nodes deploys the new images



Module - Glance



VMIC will preserve a copy of each image in the VMIC image archive. The image archive will use SWIFT.

For Glance, when the image is registered it is created a new entry in the image service and a new image copy in the glance store method.

- VMIC
 - Virtual Machine Images management
 - HEPiX specifications
 - Traceability
 - Supports BitTorrent image distribution
 - Integration with OpenStack Glance
 - Implementation started

- HEPiX Virtualization Working Group
- HEPiX Virtualization Tools (Owen Syngé)
 - <https://github.com/hepix-virtualisation>
- HEPiX spring and autumn meetings 2010



www.cloudtweaks.com – David Fletcher