

OpenStack

overview of the project

Belmiro Daniel Rodrigues Moreira

CERN IT-PES-PS

January 2011

Disclaimer: This presentation reflects the experience and opinion of the author which could be different from the CERN vision.
Some information and images were taken from the OpenStack website and presentations.

What is OpenStack?

“to produce the ubiquitous open source cloud computing platform that will meet the needs of public and private cloud providers regardless of size, by being simple to implement and massively scalable.”

What is OpenStack?

- Open source project;
- Apache 2.0 license;
- No paid “enterprise version”;
- Open design process;
- Modular design for deployment flexibility via APIs;
- Compute Engine – nova;
- Storage Engine – swift;

What is OpenStack?

- March 2010
 - Rackspace open source their cloud software;
- May 2010
 - NASA open source the Nebula software;
- June 2010
 - OpenStack formed;
- October 2010
 - First release “Austin”;

What is OpenStack?

- OpenStack compute
 - Software to provision virtual machines on standard hardware at massive scale;
- OpenStack object storage
 - Software to reliably store billions of objects distributed across standard hardware;



OpenStack hype?

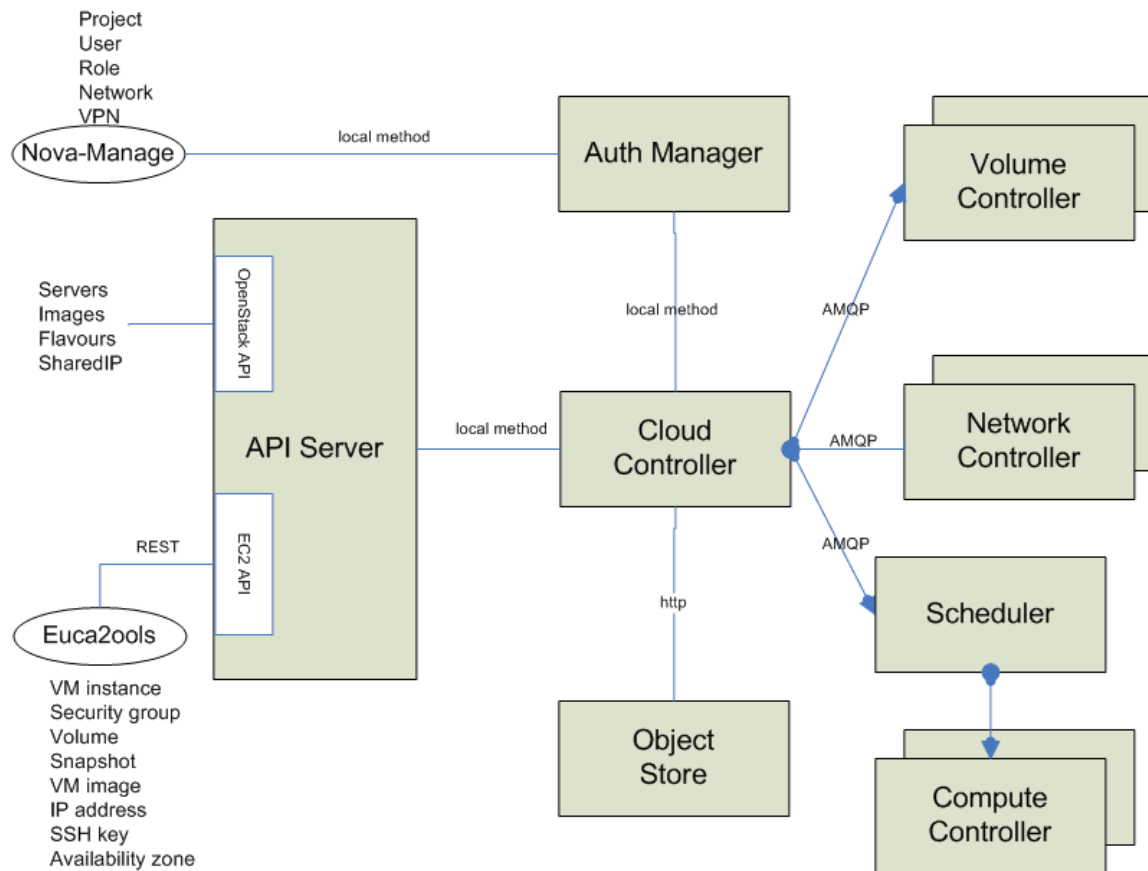


Compute Engine - nova

- IaaS
 - Similar to Amazon EC2 and RackSpace Cloud Servers;
- Supports various hypervisors:
 - XEN, KVM;
- APIs:
 - EC2;
 - OpenStack / Rackspace;

Compute Engine - nova

- Architecture



Compute Engine - nova

- Message queue
 - brokers the interaction between compute nodes, volumes, the networking controllers, API endpoints, the scheduler;
- Computing controller
 - manages computing instances on host machines;
- Network controller
 - manages the networking resources on host machines;
- Volume controller
 - interact with iSCSI storage to manage LVM-based instance volumes

Storage Engine - swift

Based on the Cloud Files offering from Rackspace;

- Concepts:

- Account
- Container
- Object



- Good for storing unstructured data;
- Object name as lookup key;
- No central point of failure;
- Keeps multiple replicas of each object;
- (...)

Storage Engine - swift

- Limitations:
 - Objects must be < 5 GB;
 - It does not use the typical POSIX filesystem semantics;
 - No user quotas;
 - No directories hierarchies;
 - No writing to a byte offset in a file;
 - (...)

Playing with nova

- nova CLI

```
nova-api          nova-compute      nova-dhcpbridge  nova-manage
nova-network      nova-objectstore nova-scheduler
```

- euca2ools

```
euca-add-group
euca-describe-instances
euca-delete-group
euca-release-address euca-allocate-address
euca-describe-regions
associate-address
euca-describe-snapshots
euca-delete-volume
euca-run-instances euca-authorize
euca-detach-volume
image
euca-disassociate-address
euca-describe-availability-zones
euca-upload-bundle euca-confirm-product-instance
euca-describe-groups
euca-version euca-create-snapshot
euca-modify-image-attribute
euca-describe-images
euca-delete-bundle
euca-register euca-add-keypair
euca-describe-keypairs
euca-delete-keypair
euca-reset-image-attribute euca-
euca-delete-snapshot
euca-revoke euca-attach-volume
euca-describe-volumes
euca-deregister
euca-terminate-instances euca-bundle-
euca-describe-addresses
euca-unbundle euca-bundle-vol
euca-download-bundle
euca-get-console-output
euca-describe-image-attribute
euca-create-volume
euca-reboot-instances
```

Playing with nova

- Create an user

```
dani@ubuntu:~$ sudo nova-manage user admin belmiro
2011-01-27 10:26:01,276 AUDIT nova.auth.manager [-] Created user belmiro (admin: True)
export EC2_ACCESS_KEY=ea178d0d-68a4-4617-b1f2-429b4adb3f34
export EC2_SECRET_KEY=041f4c80-53ca-49aa-acfe-4bd01d28b47f
```

- Create a project

```
dani@ubuntu:~$ sudo nova-manage project create cern belmiro
2011-01-27 10:26:59,474 AUDIT nova.auth.manager [-] Created project cern with manager
Belmiro
```

Playing with nova

- Upload a sample image to the “cloud”

```
dani@ubuntu:~$ euca-bundle-image -i images/aki-lucid/image -p kernel --kernel true
dani@ubuntu:~$ euca-bundle-image -i images/ari-lucid/image -p ramdisk --ramdisk true
dani@ubuntu:~$ euca-upload-bundle -m /tmp/kernel.manifest.xml -b mybucket
dani@ubuntu:~$ euca-upload-bundle -m /tmp/ramdisk.manifest.xml -b mybucket
dani@ubuntu:~$ euca-register mybucket/kernel.manifest.xml
dani@ubuntu:~$ euca-register mybucket/ramdisk.manifest.xml
dani@ubuntu:~$ euca-bundle-image -i images/ami-tiny/image -p machine --kernel ami-
fcbj2non --ramdisk ami-orukptrc
dani@ubuntu:~$ euca-upload-bundle -m /tmp/machine.manifest.xml -b mybucket
dani@ubuntu:~$ euca-register mybucket/machine.manifest.xml
```

Playing with nova

- Launching an image

```
dani@ubuntu:~$ euca-run-instances ami-2vovoyhw --kernel ami-sga4bsyz --ramdisk ami-rsds1avs -k mykey
```

```
dani@ubuntu:~$ euca-describe-instances
RESERVATION  r-01slpbtm          cern
INSTANCE     i-00000003             ami-2vovoyhw          10.0.1.3  10.0.1.3  pending  mykey
(cern_project, ubuntu)  0                m1.small  2011-01-27 09:19:25.166793
              nova
```

```
dani@ubuntu:~$ ssh -I mykey.priv root@10.0.1.3
```

Next Releases

- Bexar release (February 3, 2011)
 - OpenStack Compute (nova) ready for enterprise private cloud deployments and mid-size service provider deployments;
 - Enhanced documentation;
 - Easier to install and deploy.
- Cactus release (April 15, 2011)
 - OpenStack Compute (nova) ready for large service provider scale deployments;
 - This is the “Rackspace-ready” release.

Bexar release

- (nova) Boot VMs from raw images;
- (nova) Use the Glance image service;
- (nova) Hyper-V support;
- (nova) Security groups;
- (swift) Unlimited object size;
- (swift) Exposure of most of the S3 API;
- (...)

Conclusion

- nova is not production ready (Austin);
- Lacks documentation;
- Short life cycle (at least for now);
- Major players supporting the project;
- RackSpace will use OpenStack (nova) in their production environment;
- Supports/uses EC2 API.

Next steps?...

- Bexar release (February 3, 2011)
 - Dedicate 3? machines for this...
 - Multi-server installation;
 - Use the normal hypervisor configuration;
 - Run the usual SLC5 image in the VMs;
 - LVM support?
 - Attach external volumes? use swift?
 - Contextualization model?

Questions?