Fermilab

Computing Sector

Francesco Ceccarelli

Cloud interoperability in FermiCloud

Contents

- Overview
- Image interoperability
- API interoperability
- Private cloud solutions

Overview

What is Cloud Computing?

From Wikipedia.org (up to Sept. 25, 2013):

Cloud computing, or the cloud, is a colloquial expression used to describe a variety of different types of computing concepts that involve a large number of computers connected through a real-time communication network, such as the Internet.

What is Cloud Computing?

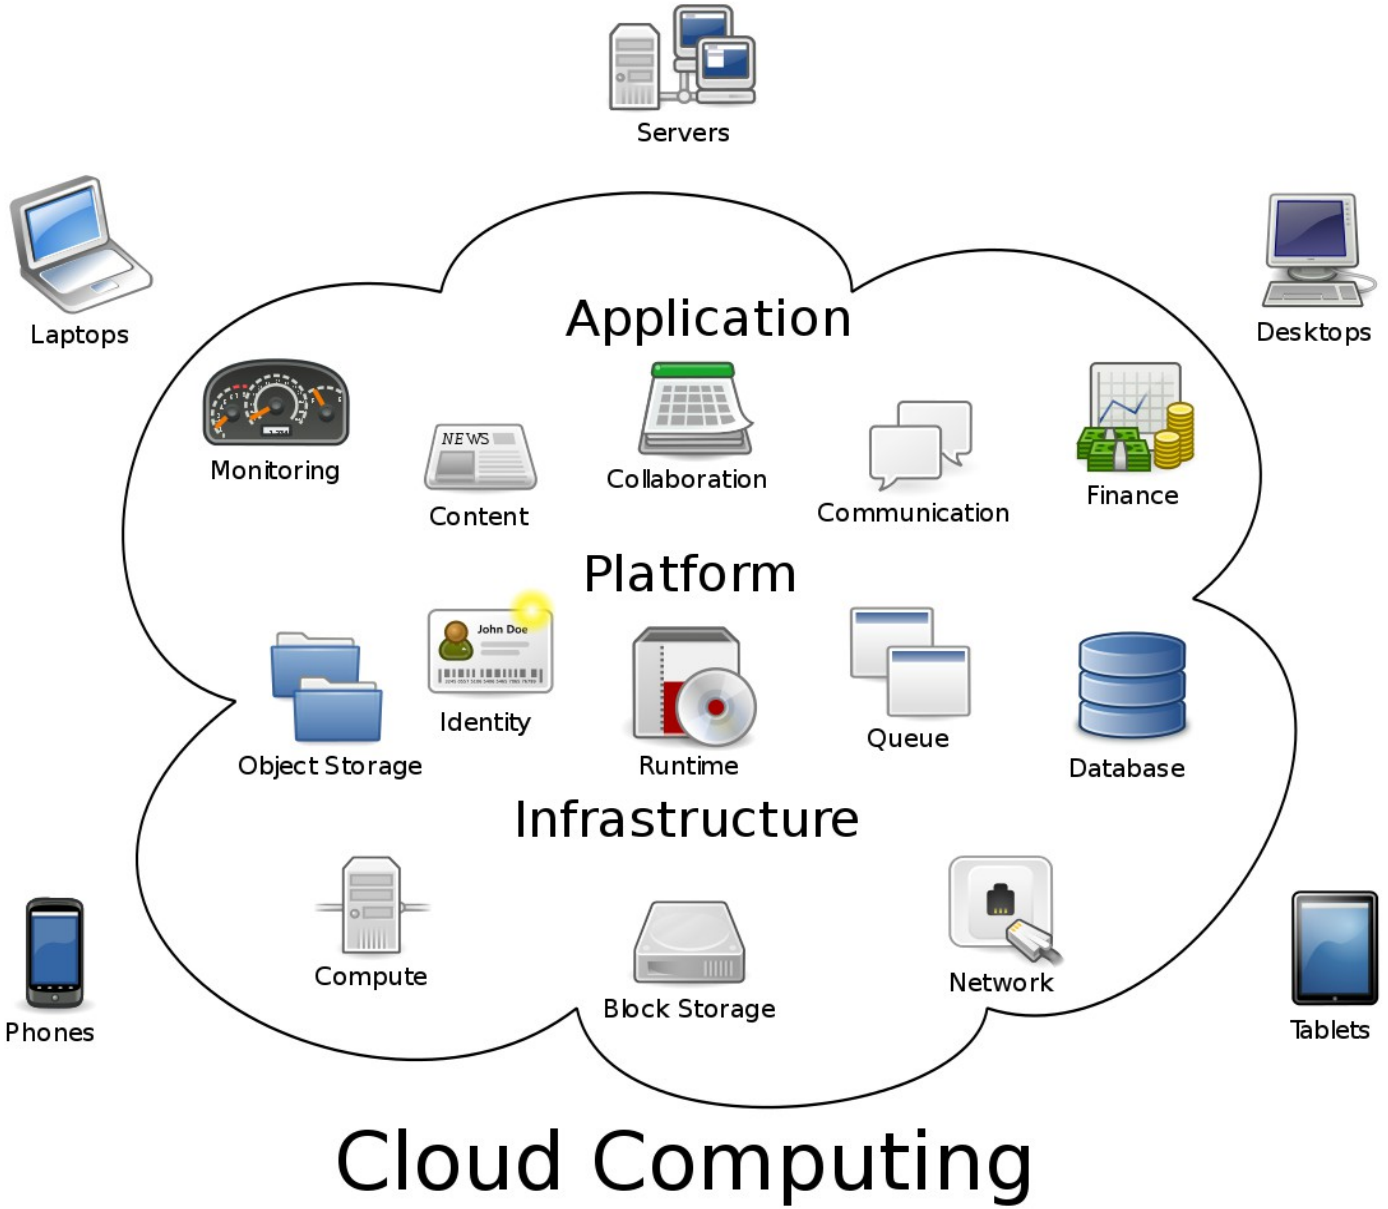

Virtualization

- It means to create a nested PC inside another PC (all the hardware is virtualized)
- Cloud uses virtualization for provisioning of computing resources
- Very popular in Desktop PC
- Common Desktop solutions:
  - VirtualBox
  - Vmware WorkStation
  - Parallels Desktop

Cloud at Fermilab

- In late 2009 the Grid Facilities Department established the FermiCloud Project
- It aims to develop and establish Scientific Cloud capabilities for the Fermilab Scientific Program
- Partnership with KISTI
- OpenNebula 3.2 based

Image interoperability

What is a disk image?

- Single file containing all the content of an hard disk
- Usually it contains also data structures (e.g. partition table)

Image portability

- In order to move a VM from a cloud solution to another we need to move the disk image
- Common problems:
  - Image format compatibility
  - Contextualization

Image formats

- FermiCloud standard format: RAW
- Other formats:
  - QCOW2
  - QCOW2 compressed
  - VDI
  - VMDK
  - ISO
  - VHD

QCOW2

- Qemu Copy and Write (2nd version)
- It comes also in a compressed version
- It supports copy-and-write
- Snapshot management
- Test result:
    - Big improvement in VMs deploying time
    - Same performances

Contextualization

- It aims to configure a VM that fit into the Fermilab environment
- It sets up:
  - Hostname
  - IP address and network parameters
  - SSH credentials
  - Fermilab Security Policy
- Managed installing the one-context RPM

Procedures developed

- The aim of this section is providing procedures to import/export VMs
- Documentation developed for the following cloud solutions:
  - Amazon EC2
  - OpenStack
  - OpenNebula 4.2
  - VirtualBox
  - GPCF (General Purpose Computing Farm)
  - VMware

To do

API interoperability

What is an API?

- Application Programming Interface
- It allows communication between two software components (very general concept...)
- No official standardization
- Amazon AWS is the “de facto” standard
- For example you can:
  - Describe images
  - Describe VMs
  - Run VMs

OpenNebula EC2 API

- It is possible to manage OpenNebula using an EC2 interface called ECONE (tested on ONE4.2 and ONE3.2)
- Not always 100% compliant with the EC2 interface

OpenStack EC2 API

- Very easy to use for dummies
- It doesn't support X.509 authentication
- It doesn't support SOAP APIs

OpenNebula and EC2 cloudbursting

- It is possible to launch VMs from ONE to an EC2 compliant public cloud (cloudbursting)
- You need to install Amazon Official Tools (or pretend to do that...)
- Very useful when there are no available local resources
- Still buggy (tested with OpenStack EC2 interface)

Private cloud solutions

OpenNebula 4

- Easy to deploy (single RPM)
- Easy to hack if you need customization (lots of ruby and bash code)
- Backward compatibility with ONE3.2
- Complete GUI (no need for CLI)
- It doesn't follow basic standards (e.g. LSB, EC2)

OpenNebula.org

OpenStack Grizzly

- Fair to deploy (puppet-automated installation)
- Compatible with a lot of image formats
- EC2 interface easy to configure and use
- Intuitive, but limited, GUI
- Contextualization using 3rd party tools (cloud-init), still buggy

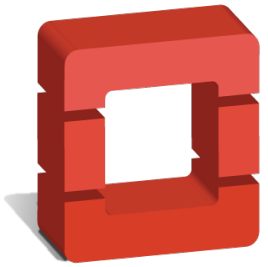

openstack™
CLOUD SOFTWARE

CloudStack

- No test deployment on site
- Friendly, scalable server architecture
- It is not possible to save/restore VM live status
- Monolithic design (harder to customize)

Eucalyptus

- No test deployment on site
- Best EC2 interface
- Aimed to be an EC2 test platform
- Lack of many basic capabilities constrained by AWS standard (e.g. save/restore/migrate VMs)