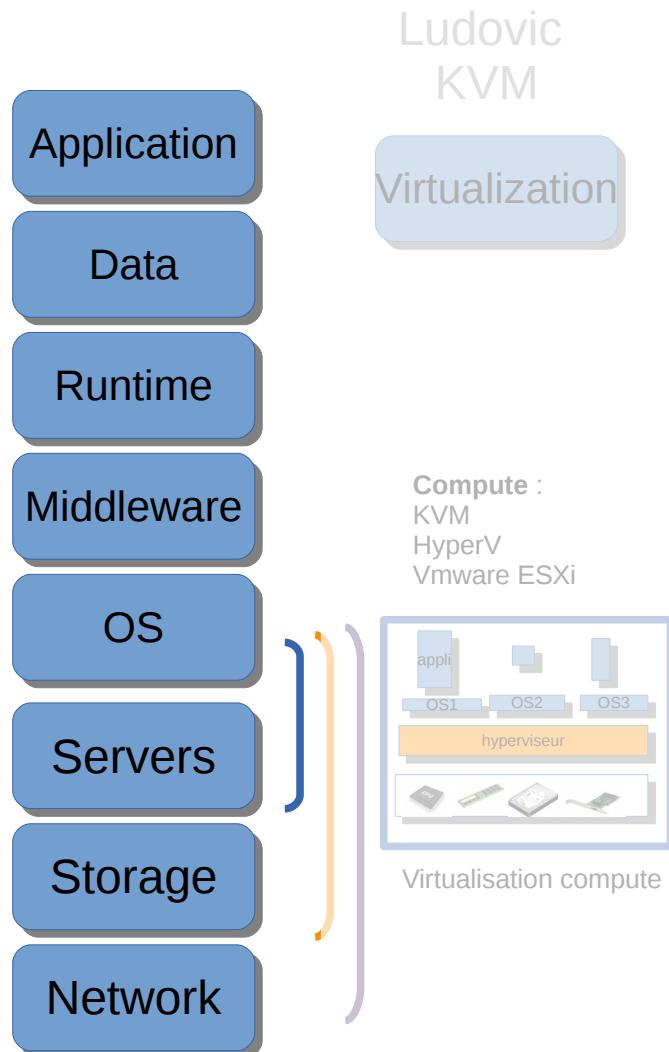


Introduction

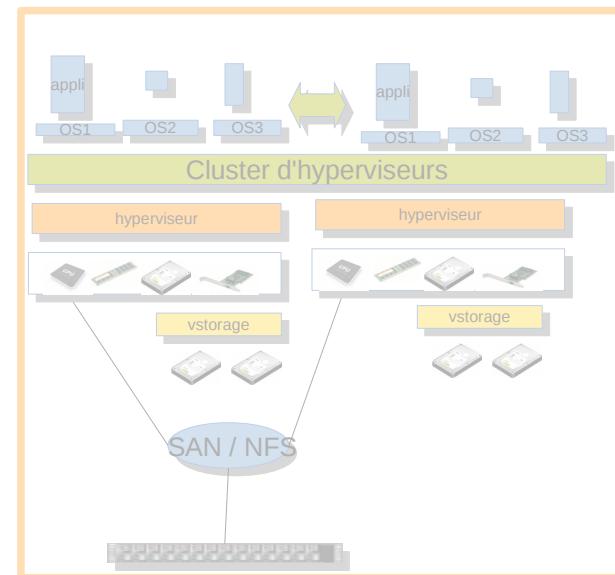
Du modèle physique au IAAS



Philippe O.
PROXMOX + CEPH

Clusters

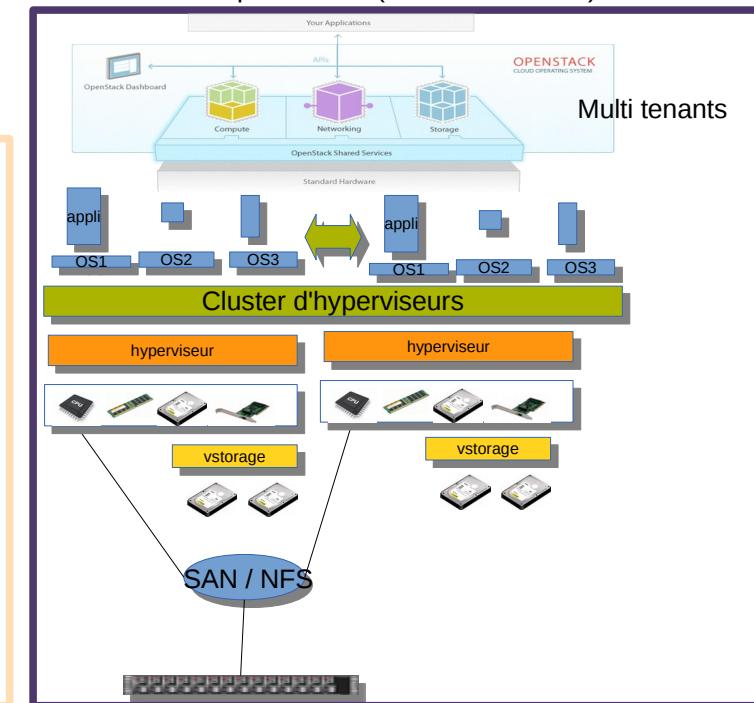
Compute : KVM, ESXi
Storage : SAN, vSAN, CEPH
Network : vswitch (L2)



Philippe S.
OPENSTACK + CEPH

IAAS

Compute : PROXMOX, Vmware VSphere
Storage : SAN, vSAN, CEPH
Network : Openvswitch (L2 + L3 + service)



Openstack et CEPH

Sommaire

- De la virtualisation au Cloud
- Introduction à Openstack
- Notre implémentation
- Un projet collaboratif ?

De la virtualisation au cloud

Introduction

Du modèle physique au cloud

Pourquoi ?
1/ CAPEX vers OPEX
2/ perte force vive
3/ changement de périmètre politique
4/ mode

CAPEX (investissement - capital expenditure)
OPEX (exploitation - operational expenditure)

On-Premise
(sur place)

Application

Data

Runtime

Middleware

OS

Virtualization

Servers

Storage

Network

vous
tenant

Infrastructure
as a service

Application

Data

Runtime

Middleware

OS

Virtualization

Servers

Storage

Network

Equipe ASR
Projet



openstack

Développeur
Équipe / projet

VM debian

provider
autres



Platform
as a service

Application

Data

Runtime

Middleware

OS

Virtualization

Servers

Storage

Network

Utilisateur
Equipe/ projet

sqlld
httpd

Software
as a service

Compte mail
Espace stockage

Application

Data

Runtime

Middleware

OS

Virtualization

Servers

Storage

Network

Niveaux de services du cloud

Les trois niveaux de services Cloud Computing

- SaaS (Logiciel-service / Software as a Service)

Mise à disposition par Internet d'applications informatiques (logiciels) comme un service. Dans le cadre d'un abonnement, les données sont elles aussi stockées sur un serveur de l'opérateur SaaS. C'est en quelque sorte la partie visible du Cloud Computing pour l'utilisateur final, qui n'a plus besoin d'installer l'application sur son poste, et qui accède à son compte par le Web, sur un environnement 100% sécurisé.

- PaaS (Plateforme-service / Platform as a Service)

Mise à disposition d'un environnement de développement et d'exploitation de logiciels sur Internet. La PaaS regroupe la partie développeur et système du Cloud Computing. Elle propose des fonctions de base, afin que le développeur, par exemple, ne doive pas se soucier de la gestion des utilisateurs ni des questions de disponibilité.

- IaaS (Infrastructure-service/ Infrastructure as a Service)

Infrastructure matérielle, louée à la demande : stockage, machines virtuelles, OS, etc. L'entreprise peut, dans ce cas, disposer sur demande d'une capacité de traitement pour n'importe quel type d'application.

Type de cloud

Les trois modèles du « Cloud Computing »

- Le Cloud Public (ou Public Cloud)

L'infrastructure est mise à la disposition du grand public ou d'un grand groupe d'entreprises via Internet. Cette infrastructure est gérée par un prestataire externe

Les ressources peuvent être partagées entre plusieurs entités clientes

- Le Cloud Privé (ou Private Cloud)

L'infrastructure est exploitée par l'entreprise. Elle peut être hébergée par l'organisation ou par un tiers

Dans le cas d'un fournisseur tiers, l'infrastructure est accessible via des réseaux sécurisés de type VPN.

- Le Cloud Hybride (ou Hybrid Cloud)

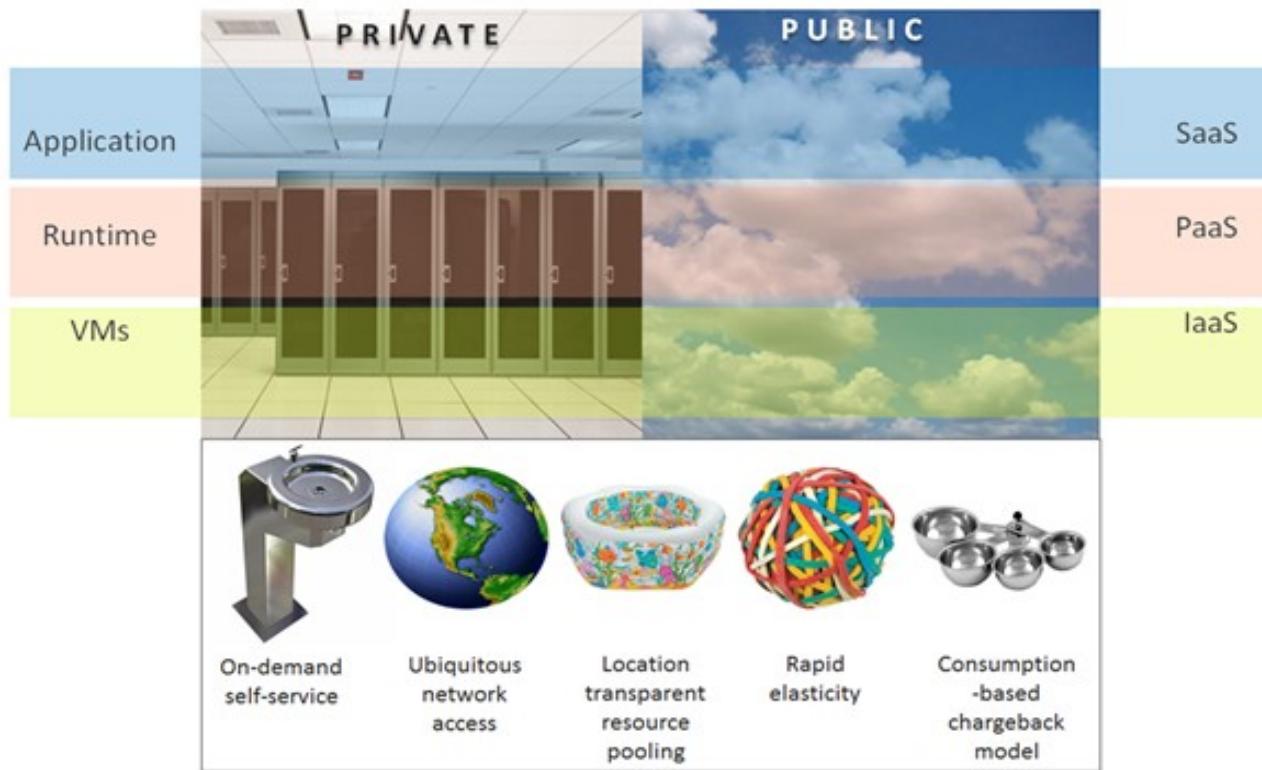
L'infrastructure est composée de deux ou plusieurs nuages (public ou privé). Ces nuages demeurent des entités indépendantes.

Ils sont toutefois liés par l'entremise de technologies standards ou propriétaires

Objectif: permettre la portabilité des données et des applications entre les nuages

Cloud « in a nutshell »

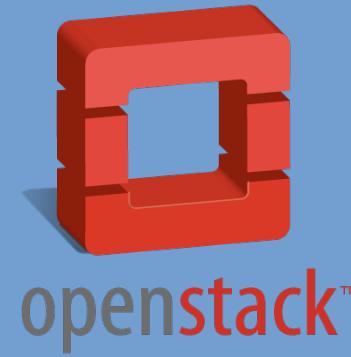
The 5-3-2 Principle of Cloud Computing



© 2011 Yung Chou
Source: <http://aka.ms/532>

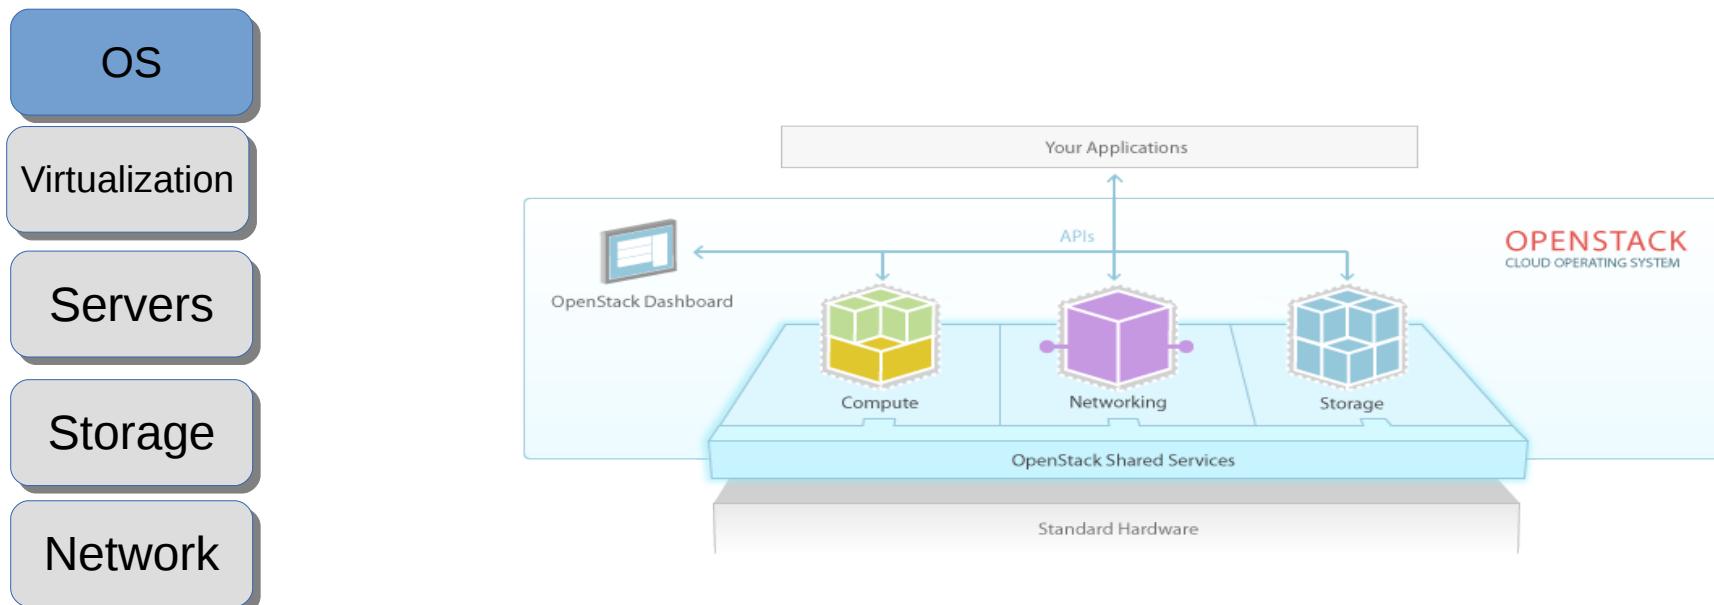
<https://blogs.technet.microsoft.com/yungchou/2011/03/07/cloud-computing-goes-far-beyond-virtualization-virtualization-vs-private-cloud-part-1/>

Introduction à Openstack



Openstack

- Sous licence Apache 2.0, le projet Open Source **IaaS (Infrastructure as a service)** , dév. en python
- OpenStack facilite et automatise la création et la gestion de ressources virtuelles en grand nombre dans un datacenter : compute , storage , network.
- Il orchestre les composants au niveau de du IaaS pour permettre à toutes ces applications d'interagir.
- La 13eme version d'OpenStack, baptisée Mitaka, 3000 développeurs de 345 entreprises, acteur majeur cloud
- Support
 - OS (sous jacent) : Ubuntu, Centos/Fedora/RHEL, Suse/Opensuse Debian
 - Hypervisor : KVM, Xen, HyperV, Vmware/ESX + conteneurs : LXC Docker
 - Network : Linux Bridge, OpenvSwitch, Cisco, NSX, OpenDaylight, OpenContrail
 - Storage : lvm, NAS/SAN, ceph



Openstack

Infrastructure
as a service

Application

Data

Runtime

Middleware

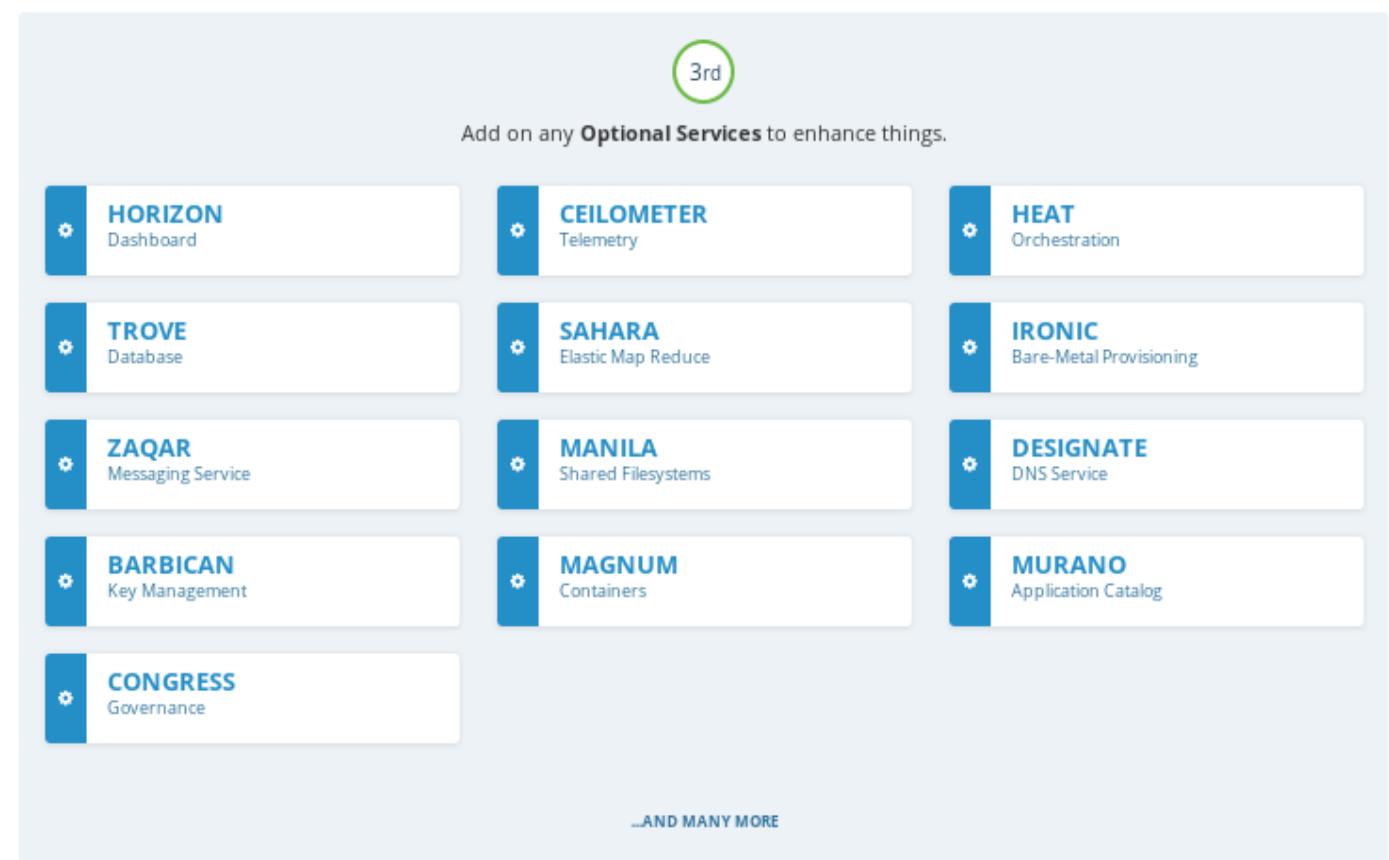
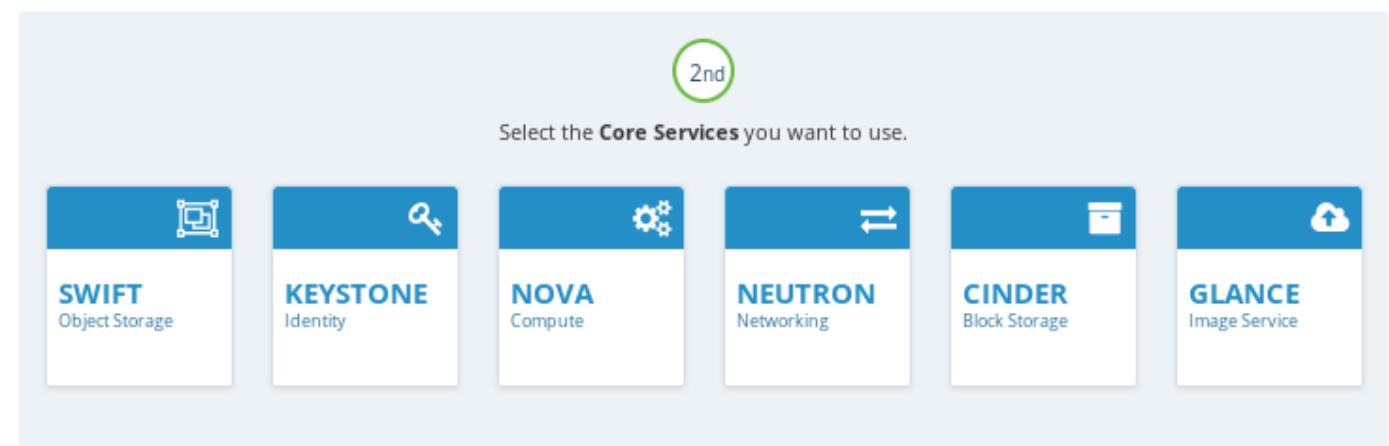
OS

Virtualization

Servers

Storage

Network



Openstack

Infrastructure as a service

Application

Data

Runtime

Middleware

OS

Virtualization

Servers

Storage

Network

Core Services (6 Results)

NOVA

Compute



Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning of machines on demand.



MORE DETAILS

NEUTRON

Networking



Enables network connectivity as a service for other OpenStack services, such as OpenStack Compute. Provides an API for users to define networks and the attachments into them. Has a pluggable architecture that supports many popular networking vendors and technologies.



MORE DETAILS

SWIFT

Object Storage



Stores and retrieves arbitrary unstructured data objects via a RESTful, HTTP based API. It is highly fault tolerant with its data replication and scale out architecture. Its implementation is not like a file server with mountable directories.



MORE DETAILS

CINDER

Block Storage



Provides persistent block storage to running instances. Its pluggable driver architecture facilitates the creation and management of block storage devices.



MORE DETAILS

KEYSTONE

Identity



Provides an authentication and authorization service for other OpenStack services. Provides a catalog of endpoints for all OpenStack services.



MORE DETAILS

GLANCE

Image Service



Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.



MORE DETAILS

Adoption is the percentage of production deployments running the project based on the latest biannual user survey results.

Maturity comes from looking at 8 distinct tags that indicate stability and sustainability. The current criteria includes whether or not the project has an install guide, whether it is supported by 7 or more SDKs, if the adoption percentage is greater than 75%, whether or not the team has achieved corporate diversity and whether or not there are stable branches.

Age is the number of years the project has been in development.

Openstack

Infrastructure as a service

Application

Data

Runtime

Middleware

OS

Virtualization

Servers

Storage

Network

Optional Services (13 Results)

| NAME | SERVICE | MATURITY | AGE | ADOPTION | DETAILS |
|------------|-------------------------|----------|-------|----------|------------------------------|
| Horizon | Dashboard | 6 of 8 | 5 Yrs | 86 % | More Details |
| Heat | Orchestration | 5 of 8 | 4 Yrs | 64 % | More Details |
| Ironic | Bare-Metal Provisioning | 5 of 8 | 3 Yrs | 20 % | More Details |
| Manila | Shared Filesystems | 5 of 8 | 3 Yrs | 11 % | More Details |
| Sahara | Elastic Map Reduce | 4 of 8 | 3 Yrs | 11 % | More Details |
| Trove | Database | 3 of 8 | 3 Yrs | 17 % | More Details |
| Zaqar | Messaging Service | 3 of 8 | 3 Yrs | 2 % | More Details |
| Designate | DNS Service | 3 of 8 | 3 Yrs | 17 % | More Details |
| Murano | Application Catalog | 2 of 8 | 2 Yrs | 15 % | More Details |
| Ceilometer | Telemetry | 1 of 8 | 4 Yrs | 59 % | More Details |
| Barbican | Key Management | 1 of 8 | 3 Yrs | 2 % | More Details |
| Magnum | Containers | 1 of 8 | 2 Yrs | 11 % | More Details |
| Congress | Governance | 1 of 8 | 2 Yrs | 1 % | More Details |

Adoption is the percentage of production deployments running the project based on the latest biannual user survey results.

Maturity comes from looking at 8 distinct tags that indicate stability and sustainability. The current criteria includes whether or not the project has an install guide, whether it is supported by 7 or more SDKs, if the adoption percentage is greater than 75%, whether or not the team has achieved corporate diversity and whether or not there are stable branches.

Age is the number of years the project has been in development.

Openstack



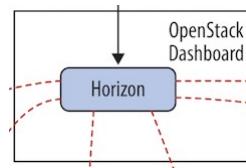
Keystone : centralise l'authentification

Neutron : réseau – pilotage architecture (Linux Bridge, OpenvSwitch, Cisco, NSX etc.) + LbaaS, FwaaS, VPNaas

Nova : gestion des instances

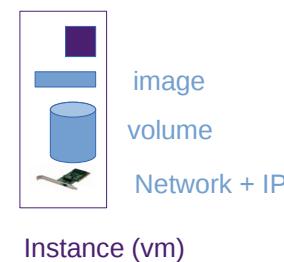
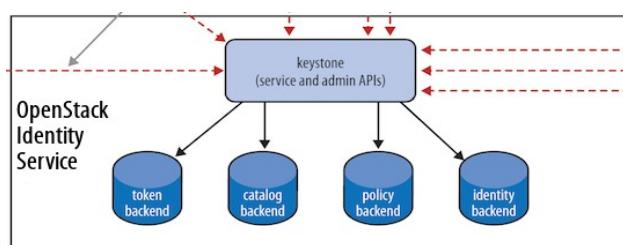
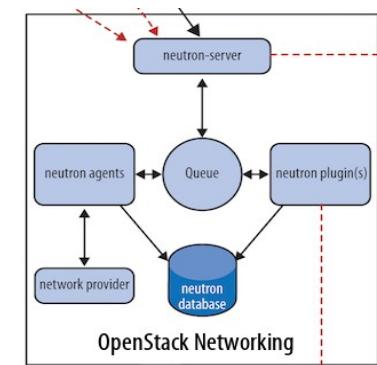
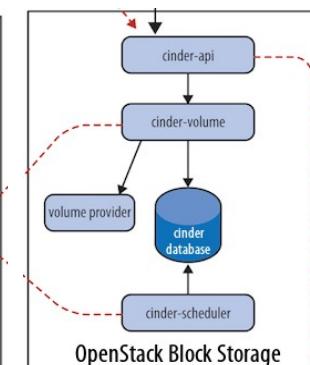
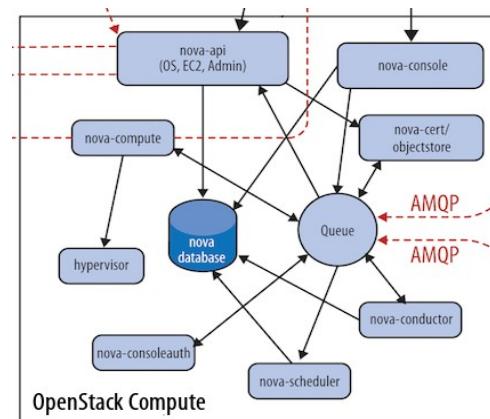
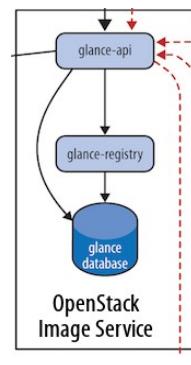
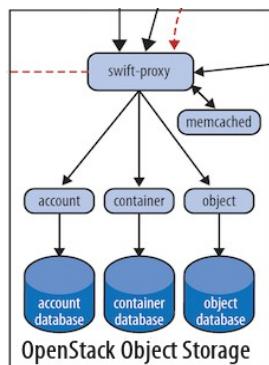
Glance : gestion des images (raw, iso, vmdk, qcow2 etc.) - disques locaux, NAS/SAN, ceph

Cinder : stockage éphémère par défaut, cinder = stockage persistent - lvm, NAS/SAN, ceph

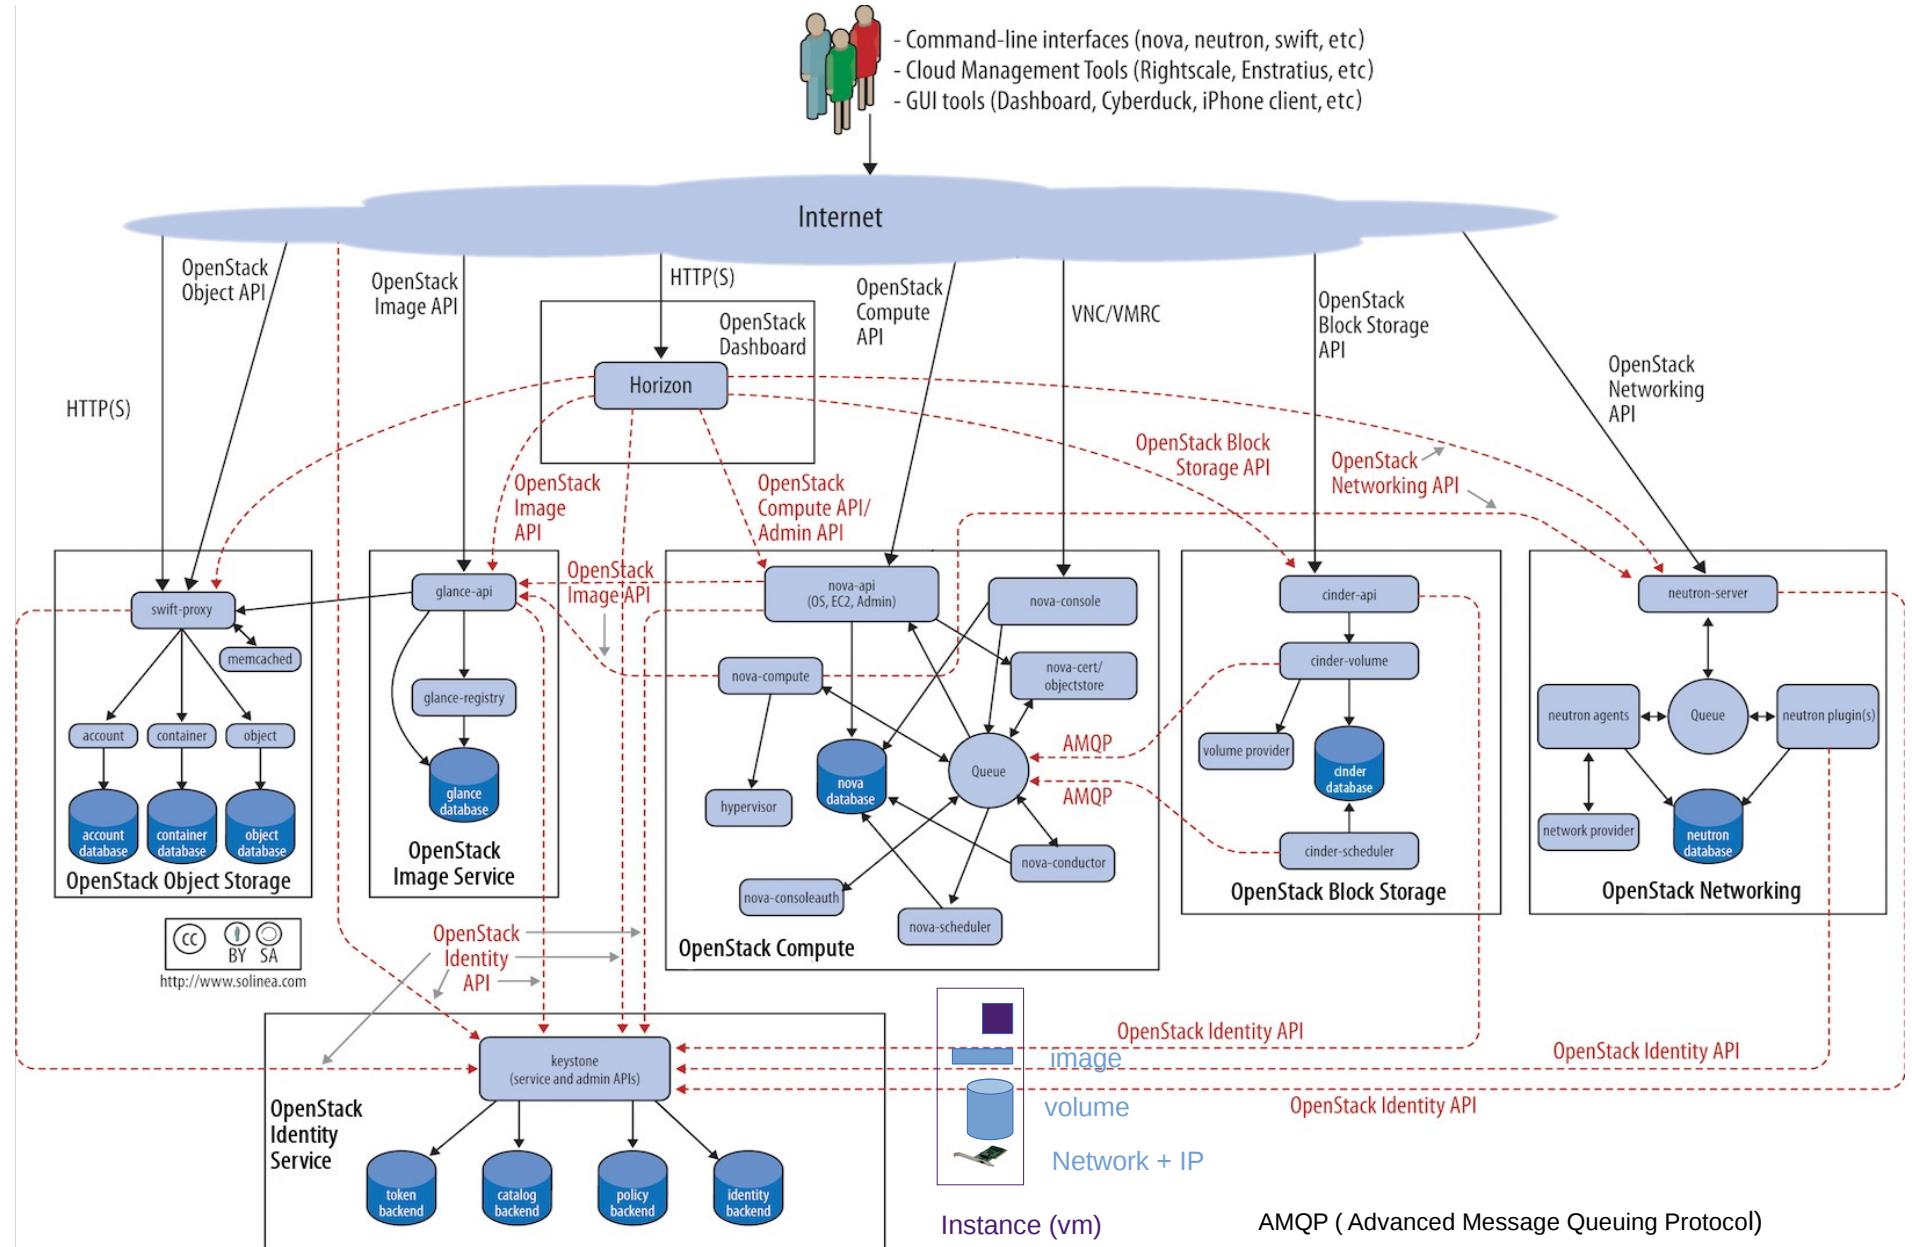


Communication entre composants

- API REST (representational state transfer) - stateless
- RPC via AMQP (Advanced Message Queuing Protocol) RabbitMQ
- SQL (interne à chaque composant)



Openstack

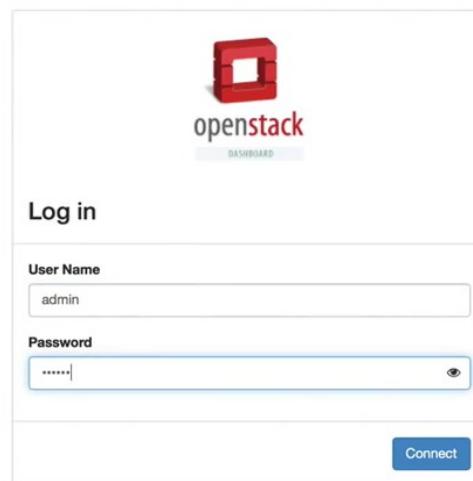


Openstack Logical Architecture

<http://docs.openstack.org/ops-guide/architecture.html>

Démo Horizon - newton

Openstack Démo Horizon



OpenStack Newton Release Demo

<https://www.youtube.com/watch?v=z6ftW7fUdp4>

Published on Oct 5, 2016

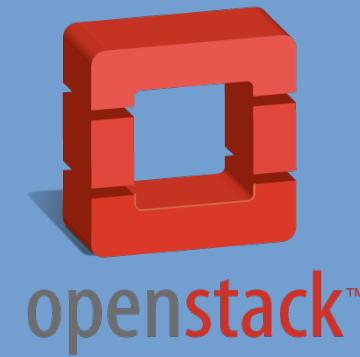
Category

Science & Technology

License

Standard YouTube License

Notre implémentation Openstack + CEPH



Notre implémentation

- Découvrir le concept cloud avec un financement faible pour :
 - Apporter de la souplesse et de la liberté à nos utilisateurs (DEVOPS)
 - Proposer en // de notre infra vmware une infrastructure de type cloud
- Choix d'openstack assez naturel
- Association avec un cluster CEPH déjà en test

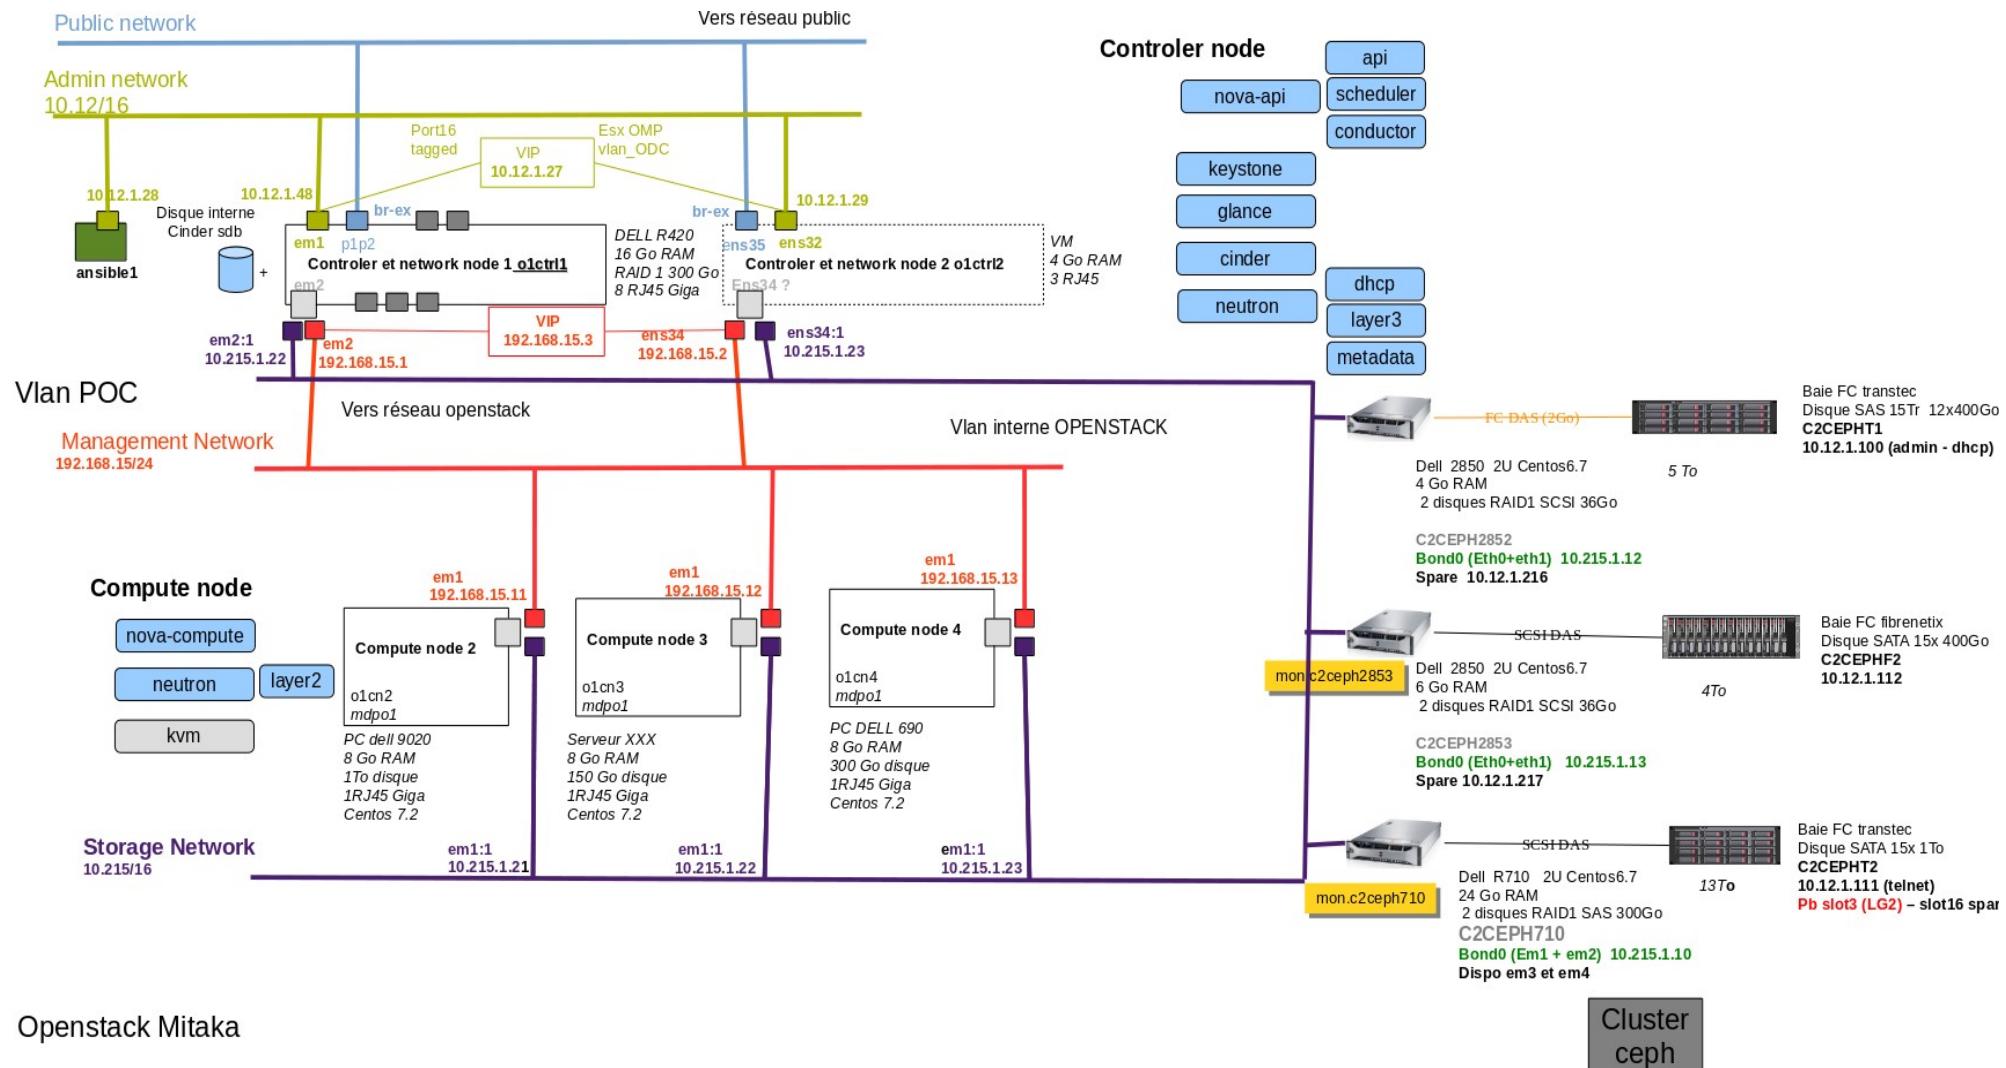
Faire un POC OPENSTACK en s'appuyant sur un stockage objet

- POC

Fin 2015, POC avec l'aide de la société ObjectifLibre

- Déploiement avec Ansible via une machine dédiée + playbook (leur choix – difficile de le faire « à la main »)
- Deux contrôleurs redondants (un physique et un virtuel)
- Trois compute node hétérogènes
- Utilisation de CEPH pour volumes, instances
- Version Mitaka (OS), Hammer (ceph) et linux CentOS 7.2 (les deux)

Notre implémentation



Notre implémentation

Public network

Admin network
10.12/16

ansible1
10.12.1.28
Disque interne
Cinder sdb

+
em1 p1p2
em2

em2:1
10.215.1.22
em2
192.168.15.1

Port16
tagged

VIP
10.12.1.27
Esx OMP
vlan_ODC

Vers réseau public

10.12.1.29

DELL R420
16 Go RAM
RAID 1 300 Go
8 RJ45 Giga

ens35 ens32

Controller et network node 2 o1ctrl2

Ens34 ?

ens34
192.168.15.2
ens34:1
10.215.1.23

in POC

Management Network
192.168.15/24

Vers réseau openstack

Vlan interne OPENSTACK

Compute node

nova-compute
neutron
layer2
kvm

em1
192.168.15.11
Compute node 2
o1cn2
mdpo1

PC dell 9020
8 Go RAM
1To disque
1RJ45 Giga
Centos 7.2

em1
192.168.15.12
Compute node 3
o1cn3
mdpo1

Serveur XXX
8 Go RAM
150 Go disque
1RJ45 Giga
Centos 7.2

em1
192.168.15.13
Compute node 4
o1cn4
mdpo1

PC DELL 690
8 Go RAM
300 Go disque
1RJ45 Giga
Centos 7.2

Storage Network
10.215/16

em1:1
10.215.1.21

em1:1
10.215.1.22

em1:1
10.215.1.23

Controler node

nova-api

keystone

glance

cinder

neutron

sd
co

mon c2cep

mon

Bilan du POC openstack + ceph ?

- Appréhender
 - Software defined Compute → ok
 - Software Defined Network → layer 2 (vswitch) – OK , layer 3
 - Software Defined Storage → ceph seul + connexion vers openstack (pool dédié)
- Découvrir IaaS et architecture multi-tenant
- Projet fédérateur, opensource et « multi partenaires »
- Projets disjoints mais complémentaires



Points d'efforts

- Courbes d'apprentissage un peu raide (cf workshop autour d'Openstack mai 2016 – OMP)
- Versions (6 mois ou 1 an), mises à jour (?)
- Jeunesse

Points forts

- Vrai projet collaboratif et prometteur
- Réalisations en production : CERN - 150 000 coeurs sur 5 000 serveurs hébergés dans deux datacenters basés à Genève et à Budapest, CloudWATT (cloud souverain orange), TETANEUTRAL (Laurent et Mehdi)
- Mettre un pied dans le monde du cloud
 - comme nous l'avons fait il y a dix ans pour la virtualisation
 - Opensource, machines en fin de garantie

Projet collaboratif ?

Un projet collaboratif ?

Pourquoi ne pas aller plus loin ?

- Mise à disposition de l'infrastructure à la communauté comme embryon : Openstack et ceph
- Pourquoi ?
 - Travail sur Openstack et/ou ceph
 - Pas de moyen ou de temps pour monter l'infra mais envie de l'utiliser
 - Apprentissage, expertise, comparaison, partage (entre infra)
 - Appréhender le « multi-tenant »
 - Plateforme pour POC (exemple gestion des logs, outils supervision etc.)
 - Le groupe travaille, monte le POC sur l'infra mutualisée et laisse une trace
 - Possibilité de reprise, de mise à jour et tests lors d'un changement de version, etc.
- Comment ?
 - Monter une équipe, transversale et hétérogène construite autour de la motivation, de l'échange et du partage
 - Définir des objectifs simples et atteignables avec un gain personnel et collectif
 - Faire quelques tours de roue et, dans un an, un bilan
- Actions
 - Porté par CAPITOUL ? (avis du comité pilotage)
 - Ecrire une « lettre d'intention », créer une équipe et fixer des objectifs

Références

Ne pas confondre Cloud et Virtualisation

<https://www.sfrbusiness.fr/room/cloud-et-hebergement/confondre-cloud-virtualisation.html>

OpenStack – IaaS

- <https://www.openstack.org/>
- <http://docs.openstack.org/>
- <http://www.silicon.fr/openstack-la-reference-de-liaas-ouvert-et-automatise-145916.html>

Cloud Computing: What is Infrastructure as a Service

- <https://technet.microsoft.com/en-us/library/hh509051.aspx>

VMware et Openstack

- Integrated OpenStack vous permet de contrôler les ressources de Cloud grâce à des API OpenStack
 - <http://www.vmware.com/fr/products/openstack.html>
 - <http://www.vmware.com/fr/products/vrealize-suite.html>
- / <http://www.vmtocloud.com/vrealize-automation-or-openstack-or-both>**

Discussion

Discussion

Infrastructure as a service

Application

Data

Runtime

Middleware

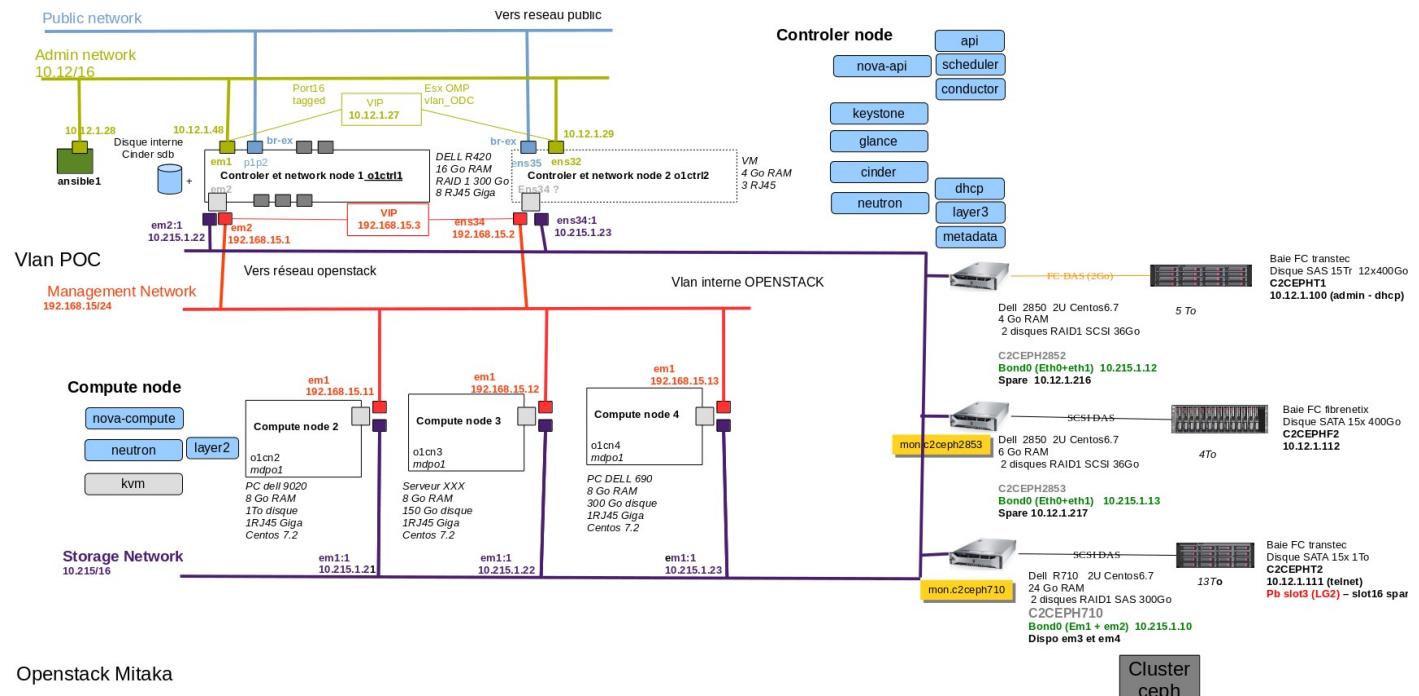
OS

Virtualization

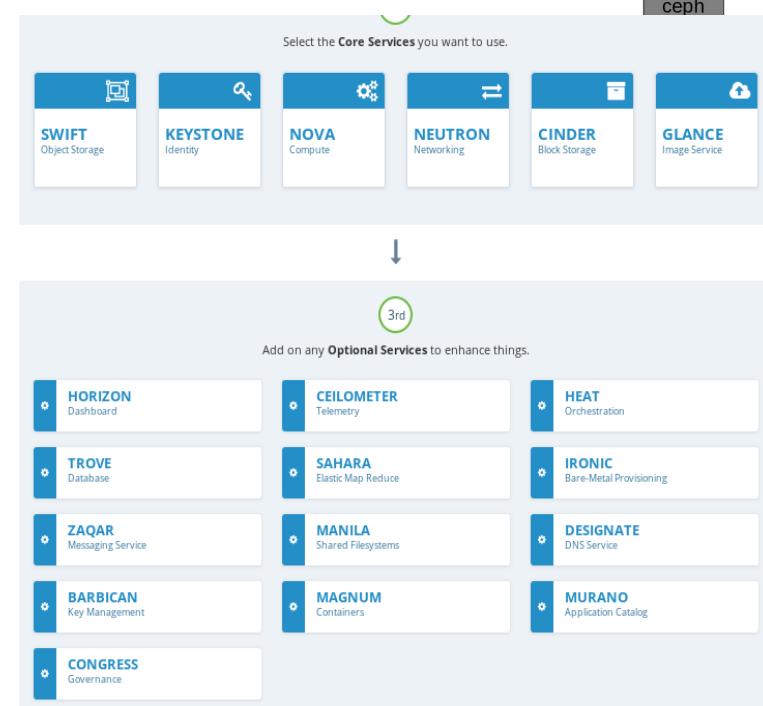
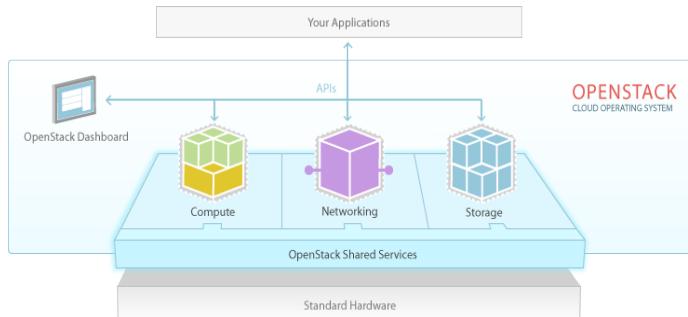
Servers

Storage

Network



Openstack Mitaka



Bonus

Openstack et hyperviseurs

<https://wiki.openstack.org/wiki/HypervisorSupportMatrix>

Compute Drivers

OpenStack Compute (Nova) has an abstraction layer for compute drivers. This is what allows you to choose which hypervisor(s) to use for your Nova deployment. The reality is that the support of each of the options is not equal. They are not tested the same amount. They also do not all support the same features. The purpose of this page is to describe the current support state of each compute driver in terms of testing and functionality.

Driver Testing Status

Group A

These drivers are fully supported. Test coverage includes: unit tests that gate commits functional testing that gate commits

Drivers in this group include: [libvirt \(qemu/KVM on x86\)](#)

Group B

These drivers are in a bit of a middle ground. Test coverage includes: unit tests that gate commits

functional testing provided by an external system that does not gate commits, but advises patch authors and reviewers of results in gerrit (the code review system).

Drivers in this group include: [Hyper-V](#) [VMware](#) [XenServer](#) [Xen via libvirt](#)

Group C

These drivers have minimal testing and may or may not work at any given time.

Use them at your own risk. Test coverage includes: (maybe) unit tests that gate commits

no public functional testing

Drivers in this group include. [Baremetal](#) [docker](#) [LXC via libvirt](#)

Vrealize service catalog

VMware vRealize® Automation

Welcome, Cloud Admin. | Preferences | Help | Logout

Home Catalog Items Requests Inbox Advanced Services Administration Infrastructure Business Management

Service Catalog

Browse the catalog for services you need.

All Services (26)

Business Group: All

Search

On behalf of: Cloud Admin (cloudadmin@corp.local)

All Services

- Agile Development
- Containers
- Infrastructure
- Platform
- Public Cloud
- Services

| Service Name | Provider | Description | Action |
|----------------------------|---------------------|---|---------|
| Apple Time Capsule | Rainpole Developers | Create a Time Capsule with VMware Photon! | Request |
| App Stack | Rainpole Developers | | Request |
| CentOS6 Test | Chef | Blueprint for deploying a CentOS Linux development server. | Request |
| CentOS6 | Chef | Blueprint for deploying a CentOS Linux development server. | Request |
| CentOS6.6 | Rainpole Developers | vCloud Air | Request |
| Centos6Kickstart | Rainpole Developers | | Request |
| CENTOS 6 w/ LogInsight ... | Rainpole Developers | Blueprint for deploying a CENTOS Linux development server. | Request |
| CentOS7 | Rainpole Developers | | Request |
| Create a computer in an... | Rainpole Developers | Creates an Active Directory computer in an organization. | Request |
| DevLampStack | Rainpole Developers | MySQL and Apache | Request |
| DockerNode | Rainpole Developers | CentOS 6.6 with Docker | Request |
| Linux | Rainpole Developers | Latest build! | Request |
| Linux AWSWest1 | Rainpole Developers | | Request |
| New Container | Rainpole Developers | New Container | Request |
| Photon | Rainpole Developers | VMware's Linux Container Host | Request |
| PhotonwithContainer | Rainpole Developers | VMware's Linux Container Host. | Request |
| PuppetEnterprise | Rainpole Developers | PE 3.7 with RHEL 6 Demo | Request |
| Rainpole Financial Syst... | Rainpole Developers | Developer instance of the Rainpole Financial System. | Request |
| RHEL 6.5 No Approval | Rainpole Developers | Blueprint for deploying a Red Hat Linux development server. | Request |
| RHEL 6.5 w/Approval | Rainpole Developers | Blueprint for deploying a Red Hat Linux development server. | Request |
| RHEL6-AWSEast | Rainpole Developers | | Request |
| RHEL7 | Rainpole Developers | Latest Version | Request |
| SUSE Linux Enterprise | Rainpole Developers | Blueprint for deploying a SUSE Enterprise Linux development server. | Request |
| Windows | Rainpole Developers | | Request |
| WindowsServer2012 | Rainpole Developers | Windows Server 2012 with the latest updates and Patches. | Request |
| WindowsServer2008 | Rainpole Developers | | |

Copyright © VMware, Inc. All rights reserved.

Build 6.2.2-2754020 | Privacy Policy | Contact us

Etat des lieux

VRA 7

GUI / API

Lifecycle
Policy
Governance
Service Catalog

orchestration

Ext Cloud

VRA
Infra Temp
- Images
- Sec
- Net

Dev
Consumer
image,sec

IaaS

Openstack
Dev
create,consume
image
Net, Top
Security

- 1) Dev Tools
- 2) 3rd Party
- 3) VRA

API / GUI

Virt Infra Consumption

| | | |
|---------|-----|---------------------|
| Isolate | NSX | VMFS NFS VSAN |
|---------|-----|---------------------|

Virt Infra G

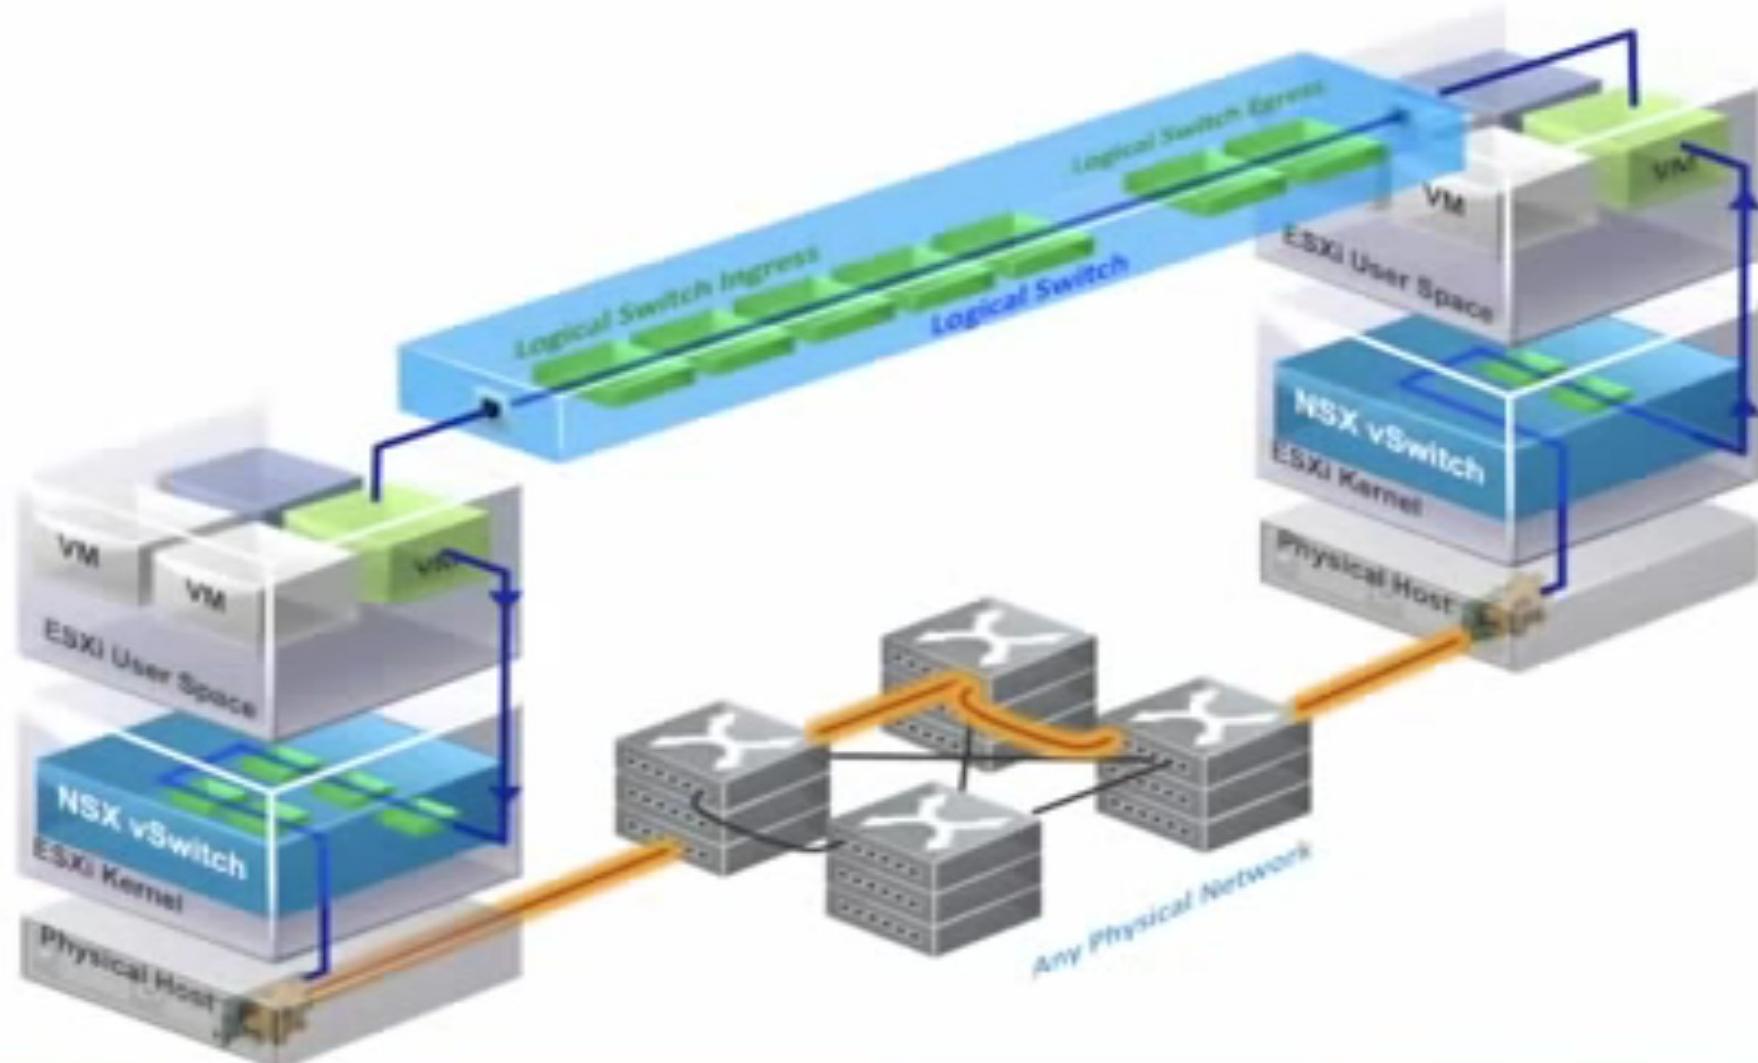


14:29 / 14:41



VmWare NSX

Virtual Network – How it works



Etat des lieux

Keep in mind...

OpenStack is NOT a Hypervisor

OpenStack Nova is a direct manager of hypervisors and built to manage multiple types of hypervisors

Flexibility is power

The flexibility around designing and deploying OpenStack is the power all Infrastructure admins want/need

Advancement thru automation

Full OpenStack capability can only be reached when some sort of automation is involved

Vrealize vs Openstack

vRA



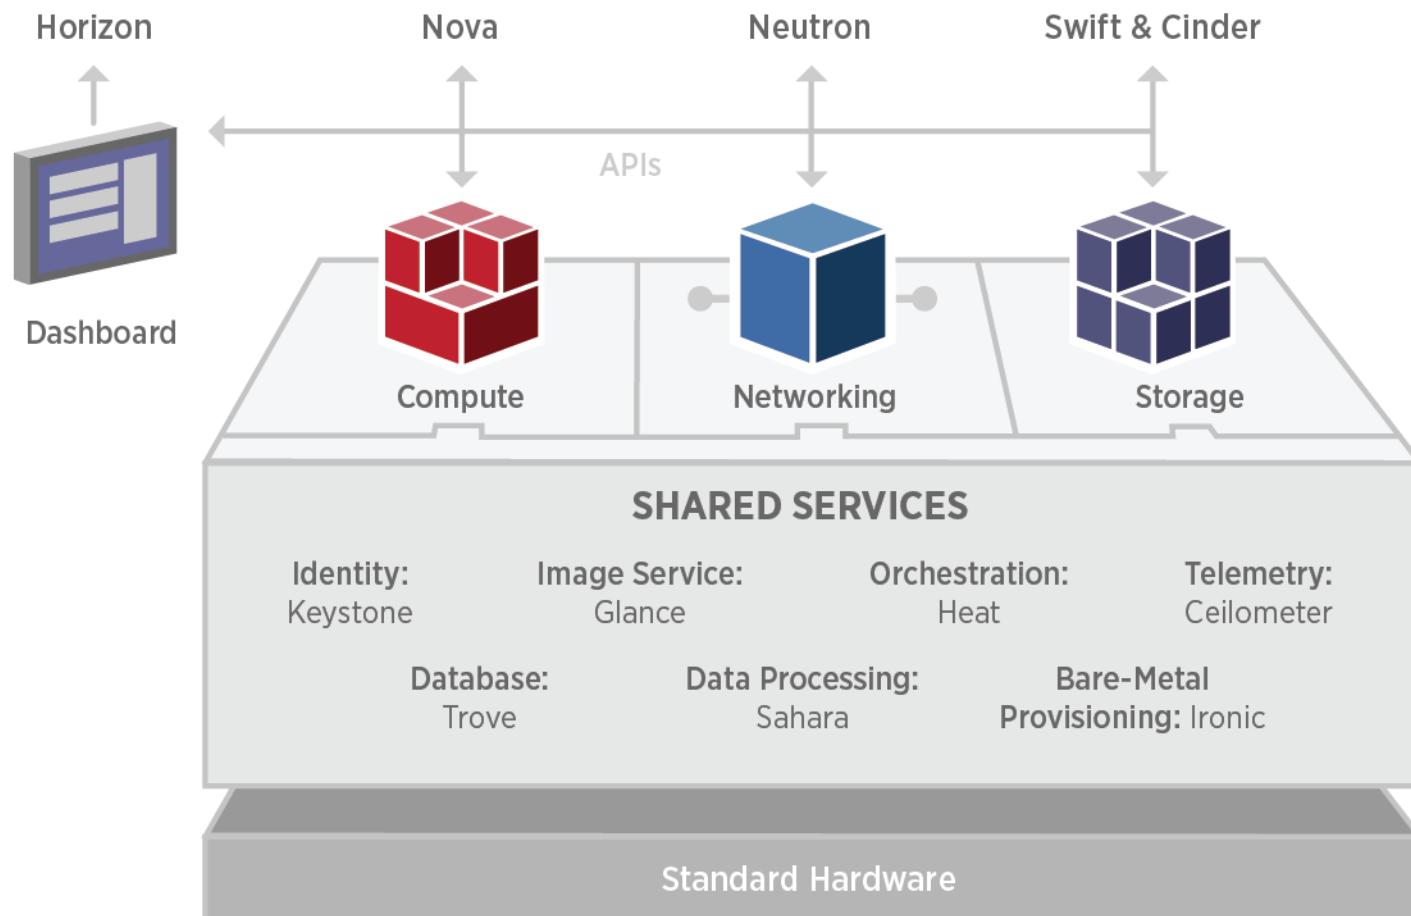
OpenStack



Self service catalog of pre-configured catalog items with IT standard policy and control

Self Service infrastructure where developers can pick and choose the items they want to form an IaaS request. Little to no IT policy and control.

Openstack



Virtualization vs Cloud

