

OpenStack day 2018

Andrea Dell'Amico <andrea.dellamico@isti.cnr.it>

Sep 21th, 2018



CNR-ISTI InfraScience laboratory

From pure PV Xen and AoE to OpenStack and Ceph: a journey



D4Science.org

Integrated technologies that provide elastic access and usage of data and data-management capabilities

D4Science web site: <https://www.d4science.org>

- Virtual Research Environments (VRE) that give access to multiple services
- Data discovery, accessing, analysis, and transformation in standard format
- Powered by gCube: <https://www.gcube-system.org/>

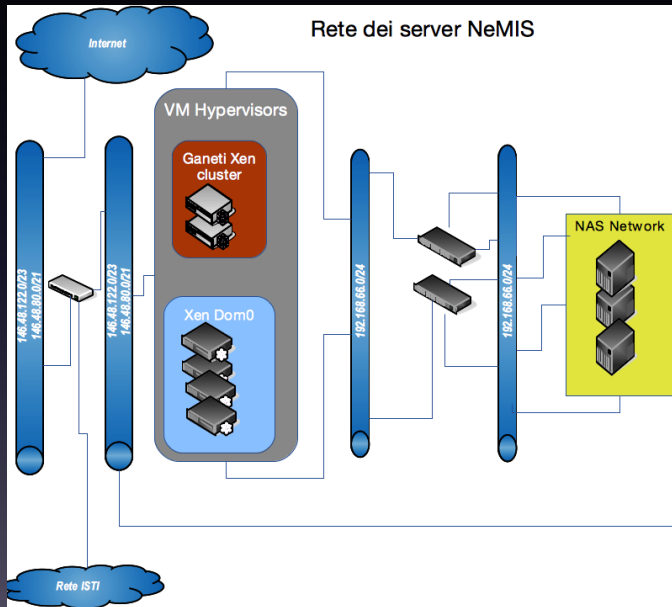


OpenAIRE, <https://www.openaire.eu>

- Operate a pan-European (and global) network for Open Science to articles and research data across countries and across research communities
- Definition and dissemination of guidelines for sharing scholarly products and links between them
- Provide services for populating and provide to the public an information graph of interlinked scholarly entities
- Provide services for assessing research impact of funders (the Commission in primis) and monitoring of open access trends
- Powered by D-Net:
<http://www.d-net.research-infrastructures.eu/>



InfraScience Data Center, 2015



InfraScience DC: Virtualization servers, 2015

27 servers that host virtual machines

PV Xen on all the servers.

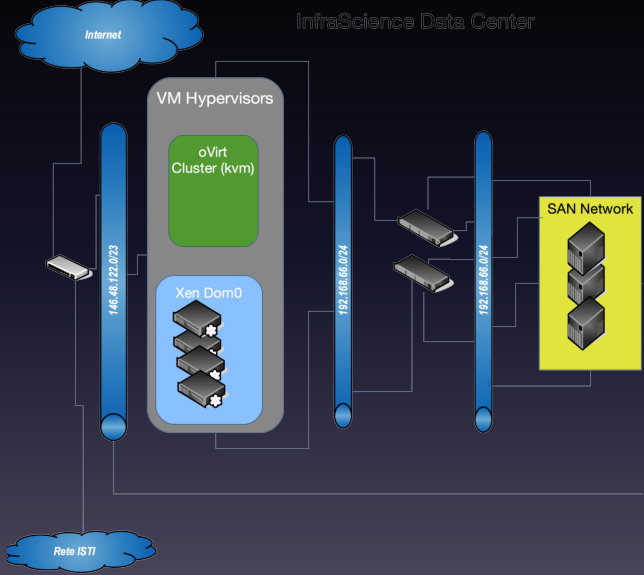
- Circa 250 Xen PV VMs
- Most of the VMs run on the newer and bigger servers.
- The older servers did not support hardware virtualization

InfraScience DC: Storage, 2015

10 servers used as storage area network (SAN)

- block devices exported by ATA over Ethernet (AoE)
- some block devices locally redundant using software raid
- some block devices have no redundancy

InfraScience DC, early 2018



InfraScience DC: Virt servers, early 2018

19 servers that host Xen based virtual machines

PV Xen on all the servers.

- More than 400 Xen PV VMs
- 4.8 TB of RAM, 960 CPU cores

5 servers that host **oVirt** (KVM) based virtual machines

- They host the *corporate* services
- Hyperconverged setup, gluster and hypervisors run on the same hosts
- The oVirt Manager runs on top of gluster too



InfraScience Data Center: AoE Storage

11 servers used as storage area network (SAN)


- Block devices exported by ATA over Ethernet (AoE)
- Some block devices locally redundant using software raid
- Some block devices have no redundancy
- Network split into two bonding groups

InfraScience infrastructure, early 2018 - continued

2 OpenStack Regions on the GARR cluster

- 100 VMs
- local instances of ldap, DNS resolver and Prometheus
- 1.3 TB of RAM in each region
- 650 CPU cores in each region

InfraScience architecture when the migration started



Resources reserved to the migration experiment when we started

OpenStack hardware resources

- 3 physical hosts reserved to OpenStack

Ceph resources

- Spare disks on the current SANs to be used as Ceph OSD

InfraScience Data Center: Ceph Storage

7 servers

Ceph cluster

- Based on Luminous
- Block devices
- Object storage (testing)
- Posix FS (testing not started yet)
- Almost half of the disks still used by AoE

OpenStack: software choices

- **RDO** as the OpenStack distribution
- **Ceph** as storage
- Global GUI cloud management using **ManageIQ**



First steps, based on PackStack

- Install PackStack
- Configure the private network
- Configure the floating IPs network
- Add a compute node
- Repeat

TripleO quickstart

- Test the installation
- Experiment with the configuration options
- Test deploying the overcloud
- Add nagios checks



TripleO production installation

- Install on baremetal on three physical hosts
 - Undercloud
 - Overcloud controller
 - Overcloud compute node
 - (A fourth node, the block storage, installed on a ceph node)

Starting the migration

WMs that

- Can be reprovisioned from scratch
- Without significant local storage
- Already redundant (behind a load balancer)

Hardware migration

- When a hypervisor is emptied from the old VMs, convert it to a compute node
- Start a new batch of VMs migration

Still a work in progress. Some VMs are too old to be migrated, a couple of Xen hypervisors are going to be operational until we can dismiss those VMs

OpenStack upgrade

- Not tested yet
- Plan: install a new TripleO undercloud and test on it



Monitoring

- Nagios checks and Prometheus exporters on the VMs
- Plan: integrate Nagios and hopefully Prometheus into the OpenStack components



That's all

Questions?

