

Low cost private cloud with *OpenNebula*



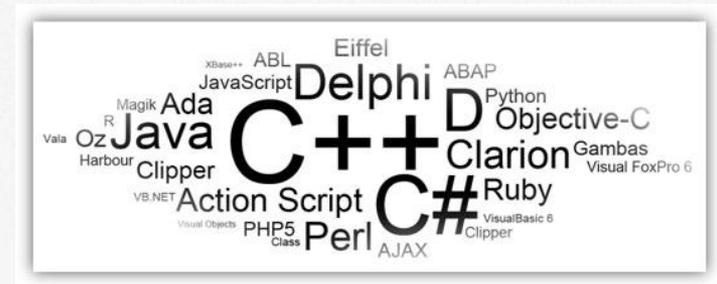
Federico Zani - INFN Roma Tor Vergata (federico.zani@roma2.infn.it)
Luca Mazzaferro - INFN Roma Tor Vergata (luca.mazzaferro@roma2.infn.it)

Tor Vergata case history...

Few money

Heterogeneous (*and quite old*) hardware

Several different analysis requirement



I have a dream...

Take advantage of every single core

High availability for core services

Easy farm management

Scalability

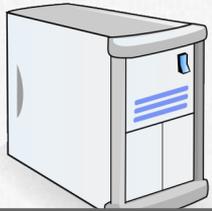
Quick integration of new hardware



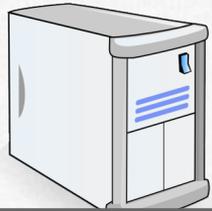
The beginning of time...



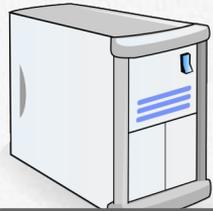
One physical host for each service



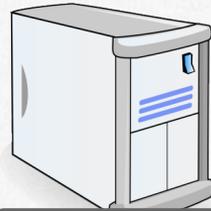
Grid worker node



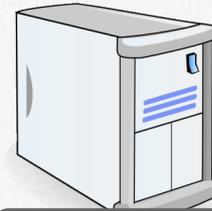
Grid worker node



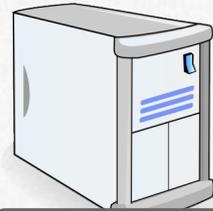
Local analysis



Storage frontend



Bastion host



Web server

The beginning of time...



PROS



(STILL SEARCHING FOR SOME...)

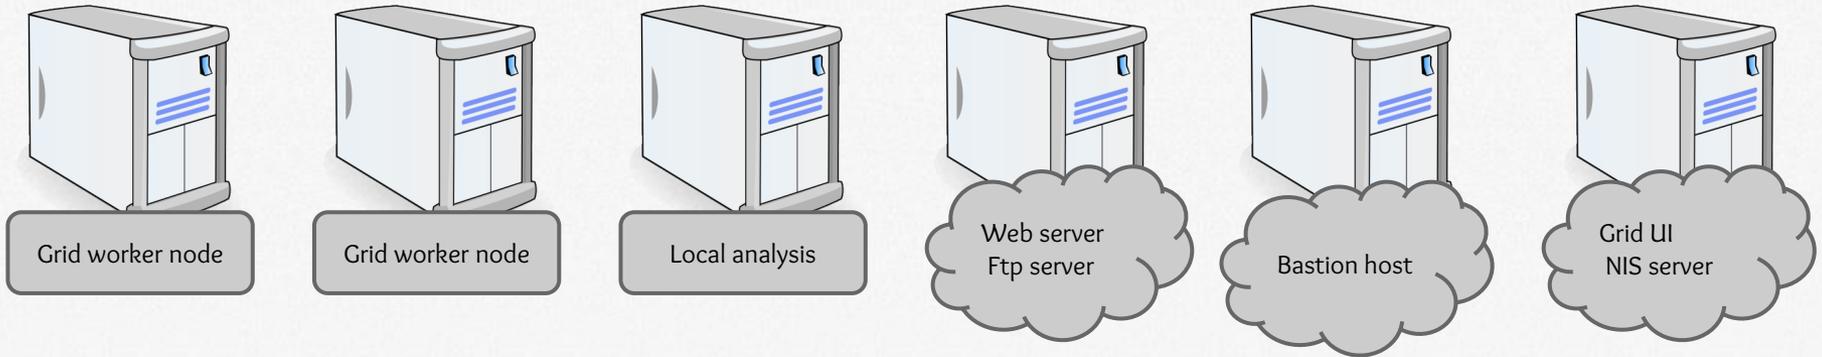
CONS

- NO CORE SERVICE HIGH AVAILABILITY
- NO EASY SCALABILITY
- COMPLEX MANAGEMENT
- PRETTY "OLD-STYLE" NODE CLONING...
- ...AND BACKUP
- SLOW INTEGRATION OF NEW HARDWARE
- [...] (YOU NAME IT)

The middle age...



Computing nodes on physical hosts
Virtualization of core services



The middle age...



PROS

- IMPROVED SERVICE AVAILABILITY
- EASIER CORE SERVICES BACKUP
- REDUCED DEPLOYMENT TIME FOR NEW SERVICES

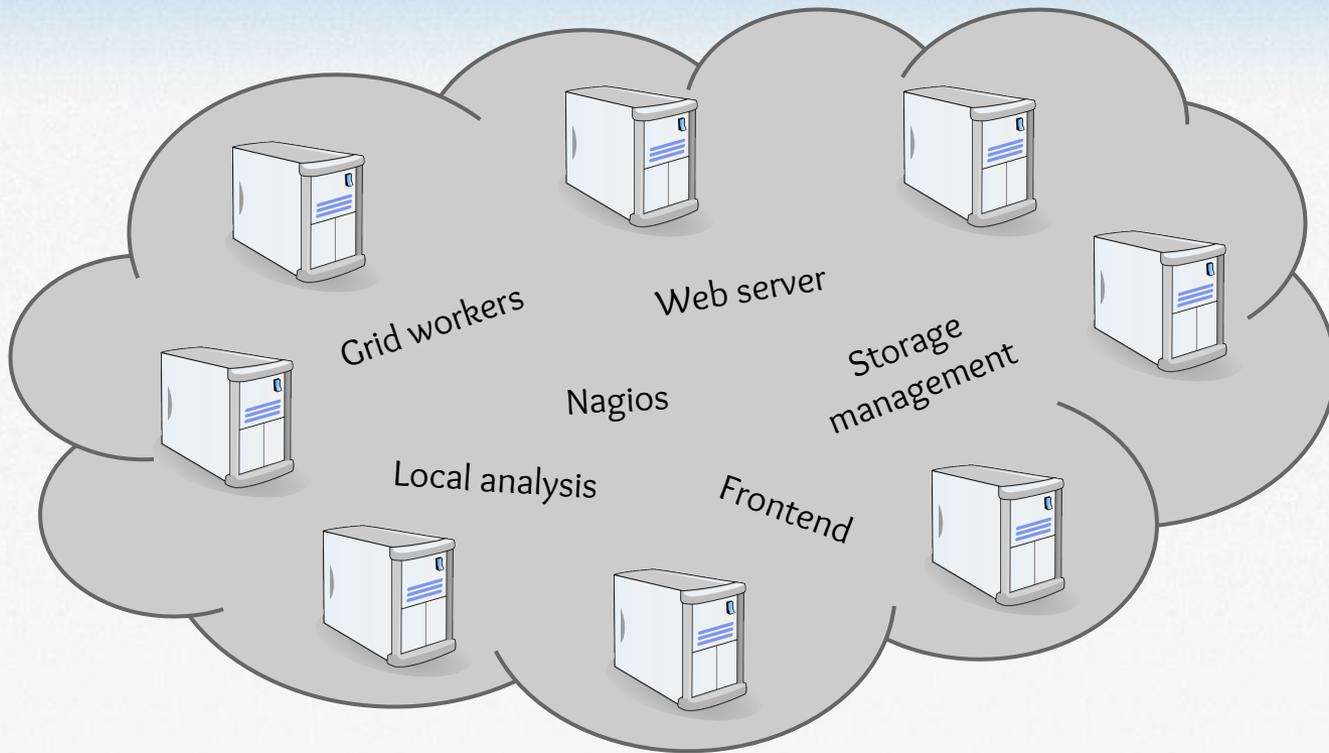
CONS

- STILL MISSING EASY MIGRATIONS, BACKUP AND NODE SCALABILITY
- NO FRONTEND TO MANAGE THE WHOLE INFRASTRUCTURE

Modern era...



OpenNebula private cloud



Modern era...



PROS

- EASY SERVICE DEPLOY
- CENTRALIZED FARM MANAGEMENT
- NODE CLONING WITH
CONTEXTUALIZATION
- INTEGRATION OF NEW HARDWARE IS
AS QUICK AS INSTALLING BASE OS

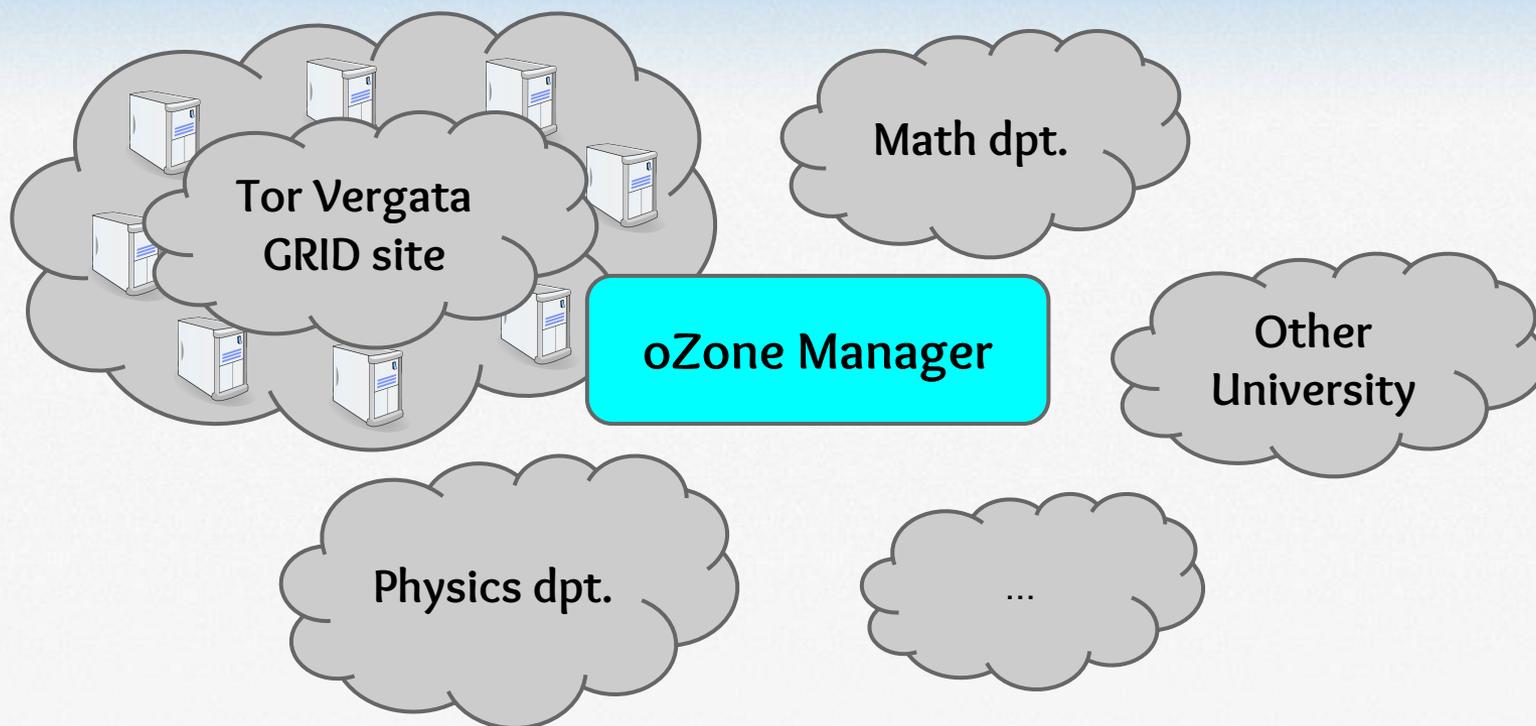
CONS

- SHARED FILE SYSTEM MIGHT BECOME
A SINGLE POINT OF FAILURE

The future...?



OpenNebula distributed datacenter with oZones



Some annoying tech details...



Physical hosts

- 11 physical hosts
- 126 virtual cores
- 106 GB ram



Network layout

- 1 gigabit ethernet for external connectivity
- 1 gigabit ethernet for private communication and storage access



Storage

- 1 SAN with 4TB dedicated to VM disk images, attached via iScsi to cloud frontend
- 1 NAS with 32TB dedicated to analysis data, mounted via nfs on all computing nodes

More annoying tech details...

Non shared file system

- Higher performance, but no live migration
- Avoid single point of failure
- Does not overload private network, already used for analysis

Cloud frontend/scheduler is a VM itself

- Avoid (another) single point of failure
- Easily migrate frontend on another physical host in case of hw issue

Contextualized nodes

- Easily deploy multiple worker node instances, using OpenNebula virtual networks and ip leases

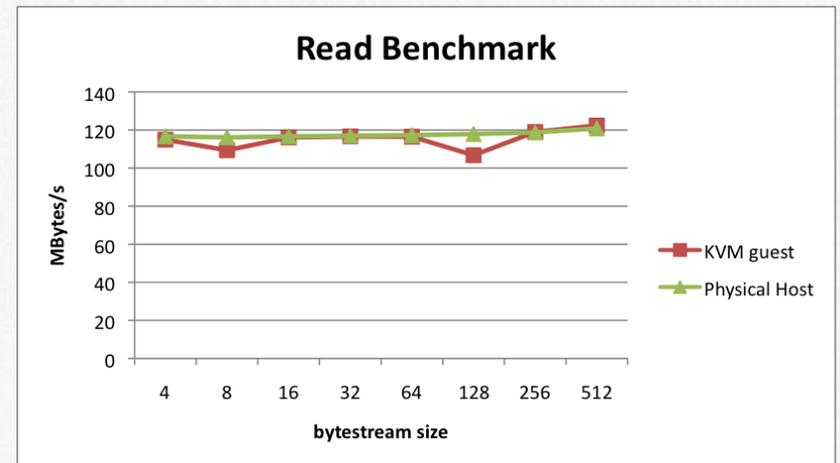
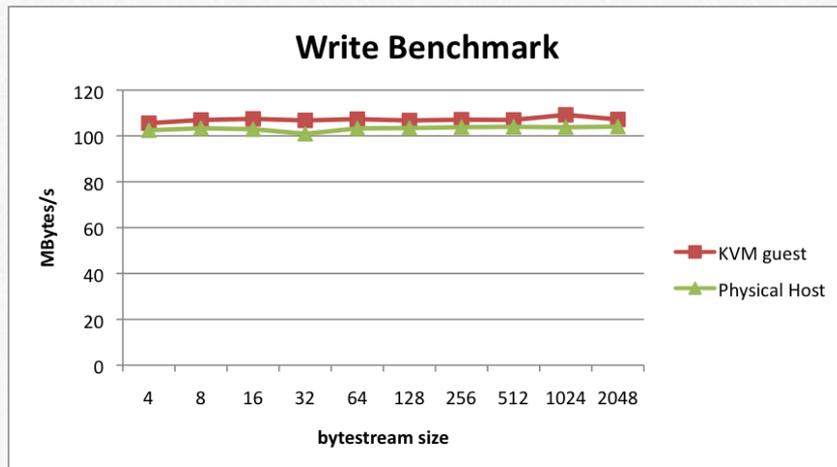
Some (*even more*) annoying benchmark...

CPU benchmark

| | Intel Xeon E5504 @ 2.00GHz | KVM guest (same CPU) |
|------------------|-----------------------------------|-----------------------------|
| OpenSSL | 44.7 AU | 43.2 AU |
| SciMark 2 | 478.36 Mflops | 472.71 Mflops |
| GCrypt | 4833 Microseconds | 4680 Microseconds |

Some (even more) annoying benchmark...

NFS mounted I/O benchmark



Question time *(and applauses)*...

