



Unified A&A

(with Cloud, HPC, VMs and Containers)

Giuliano Taffoni

Italian National Institute of Astrophysics – OATs



Resource providing

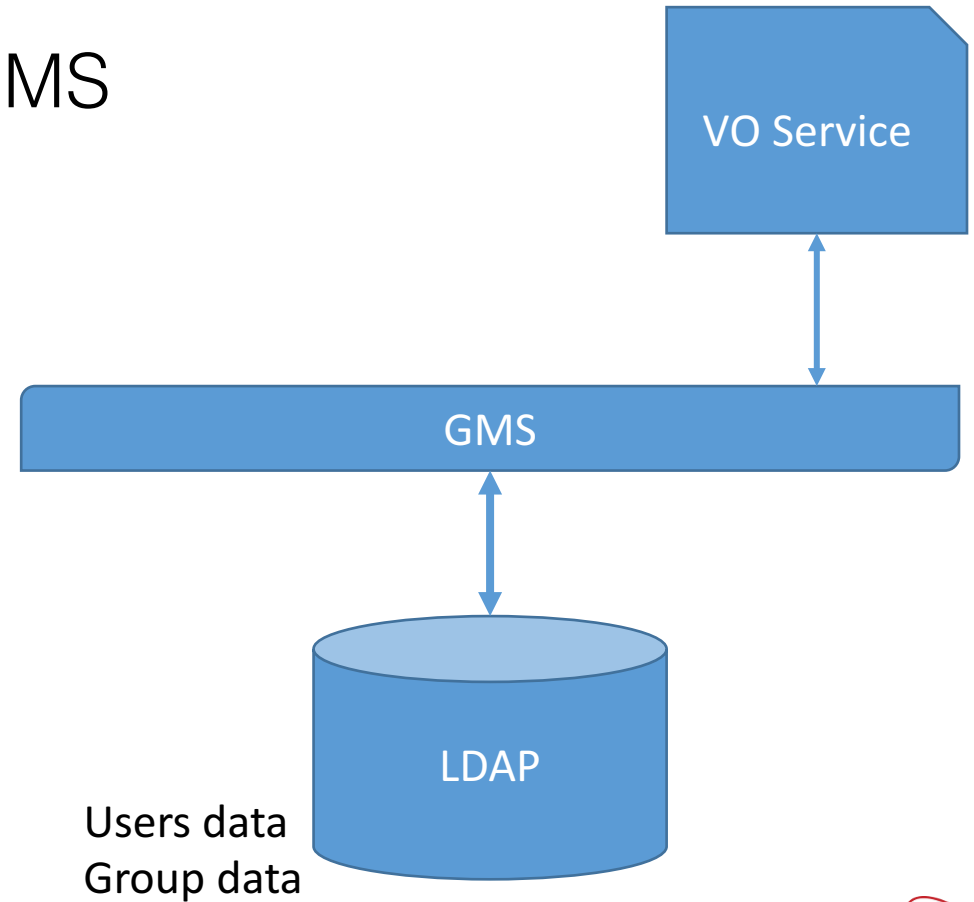
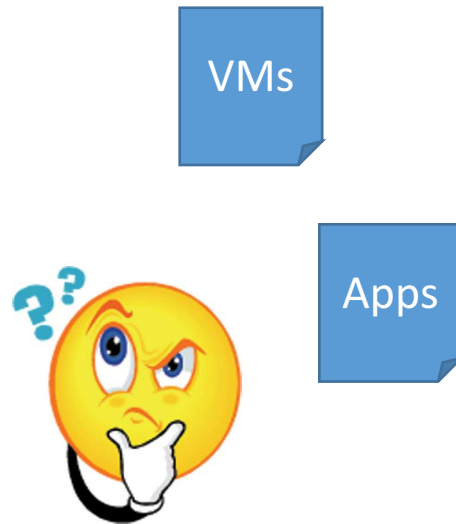
- To allow users to access computing resources for data reduction and analysis
- To allow Data Centers to expand resources on demand
- To provide a virtual laboratory for numerical astrophysics
- To implement a 'simple' way to compare theory with observations

OATs ecosystem

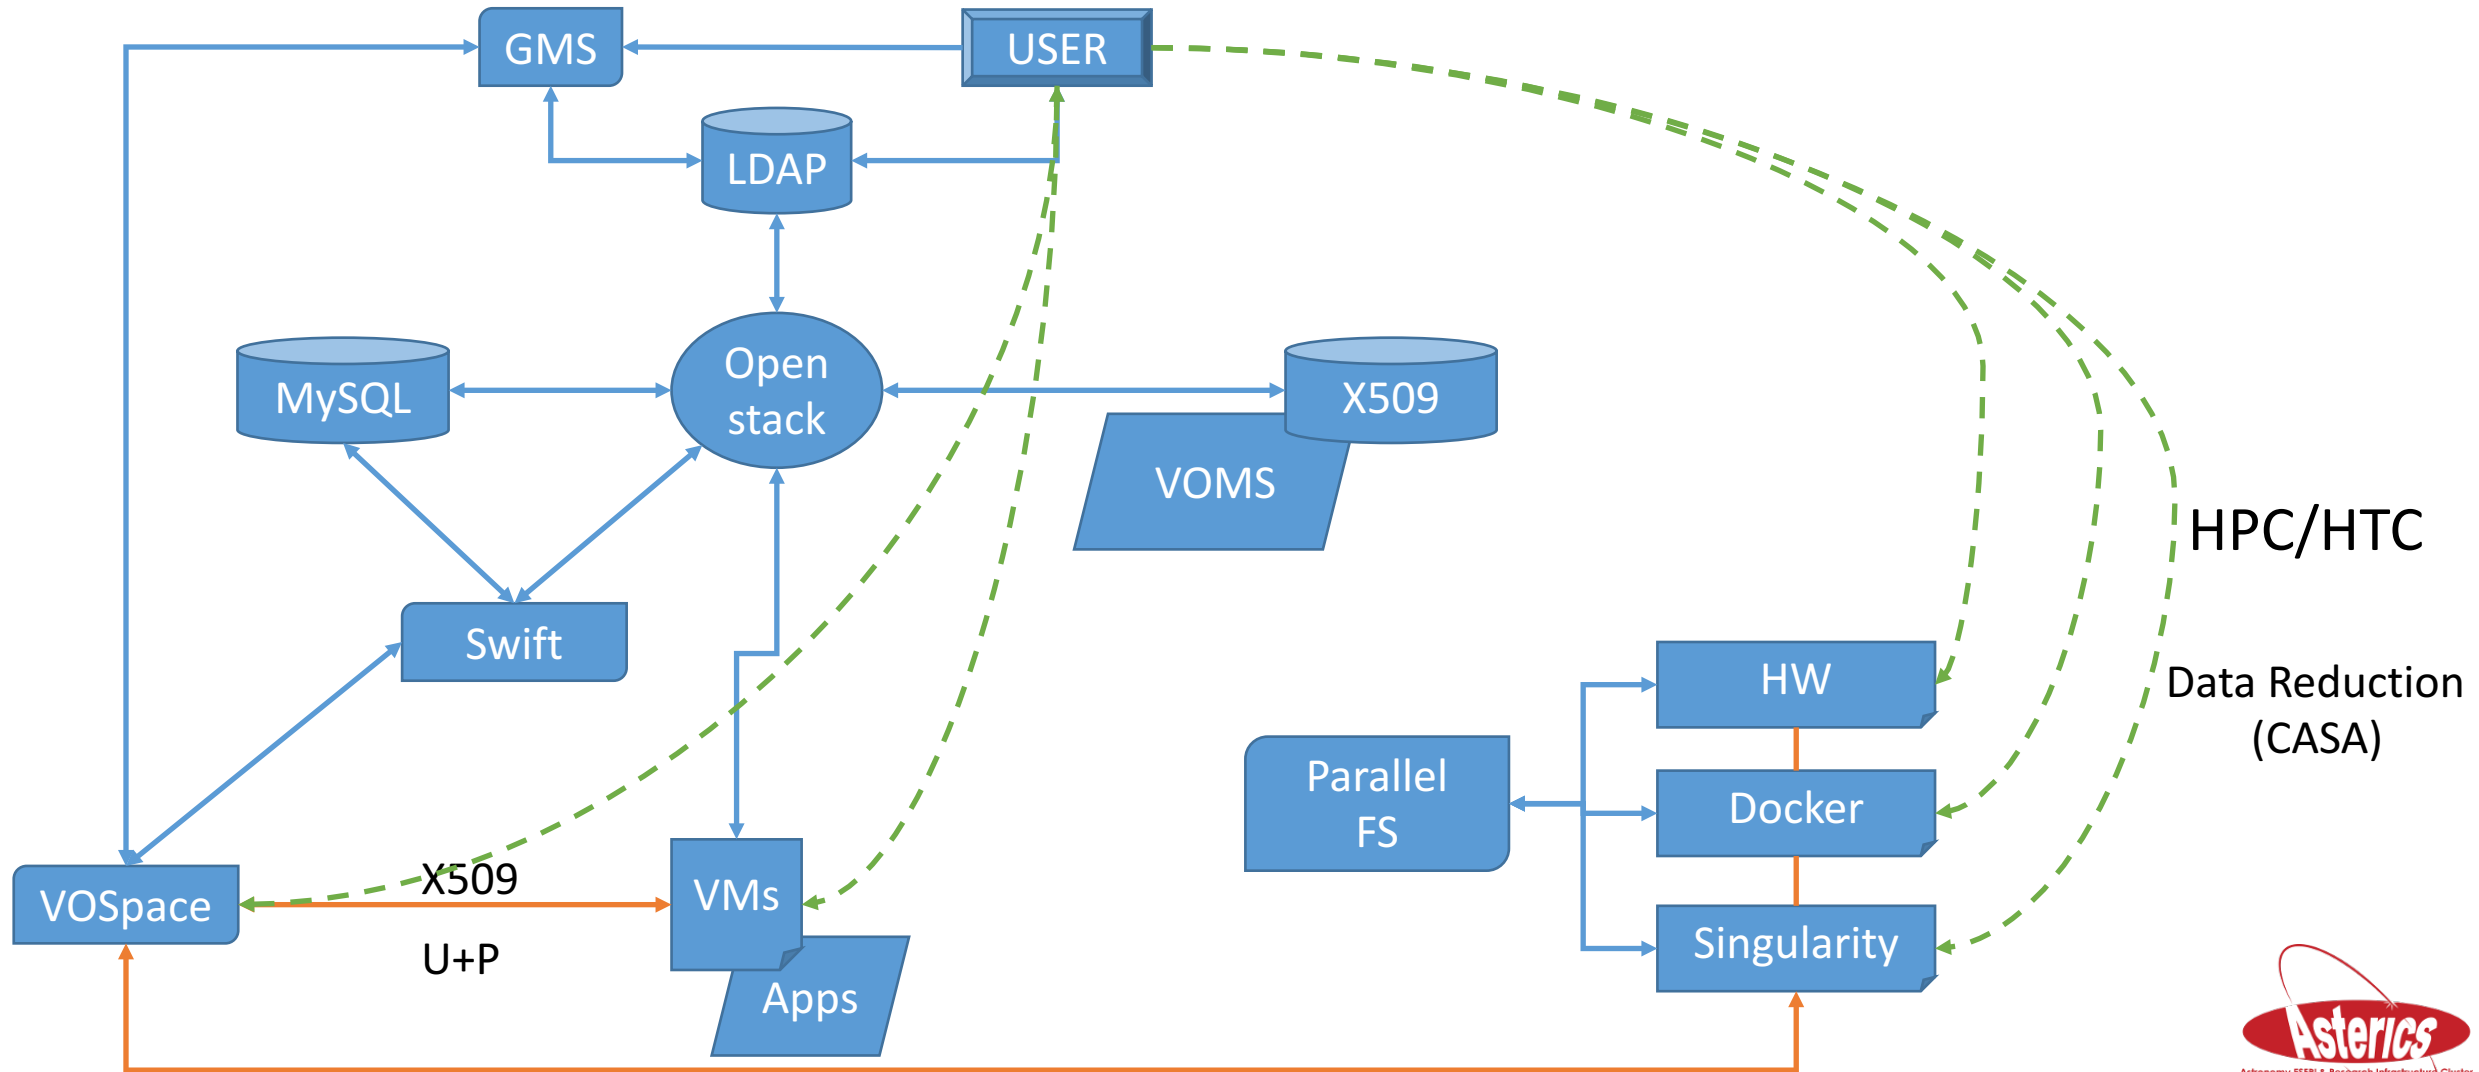
- **HOTCAT** (an HPC and HTC computing resource) 1600 Cores
- HPC interconnect (IB 56Gbs)
- Parallel Filesystem (BeeGFS)
- Software and tools for computing and data analysis
- Containers
- **CLOUDCAT** (OpenStack cloud - Mitaka) 200Core
- HPC interconnect (IB 10Gbs)
- Swift Object Storage
- VM images for computing and data analysis

IVOA services

- Authentication and Authorization: GMS
- VOSpace
- UWS
- CDP




The over all picture



OpenStack Domains

- “The intent of domain is to define the administrative boundaries for management of Keystone entities. A domain can represent an individual, company, or operator owned space.”
- Each domain may have a different A&A configuration (LDAP, MySQL, X509, SAML...)
- LDAP domain connected to GMS LDAP →
 - users authenticate to Stack using GMS credentials



Active domains

- EGI Domain: Based on X509 Auth and VOMS Authz
 - Each VOMS VOs → Tenant
 - VOMS plug-in
- INAF: LDAP (is going to be SAML). Tenant == INAF institutes
- CANFAR: LDAP (GMS LDAP)

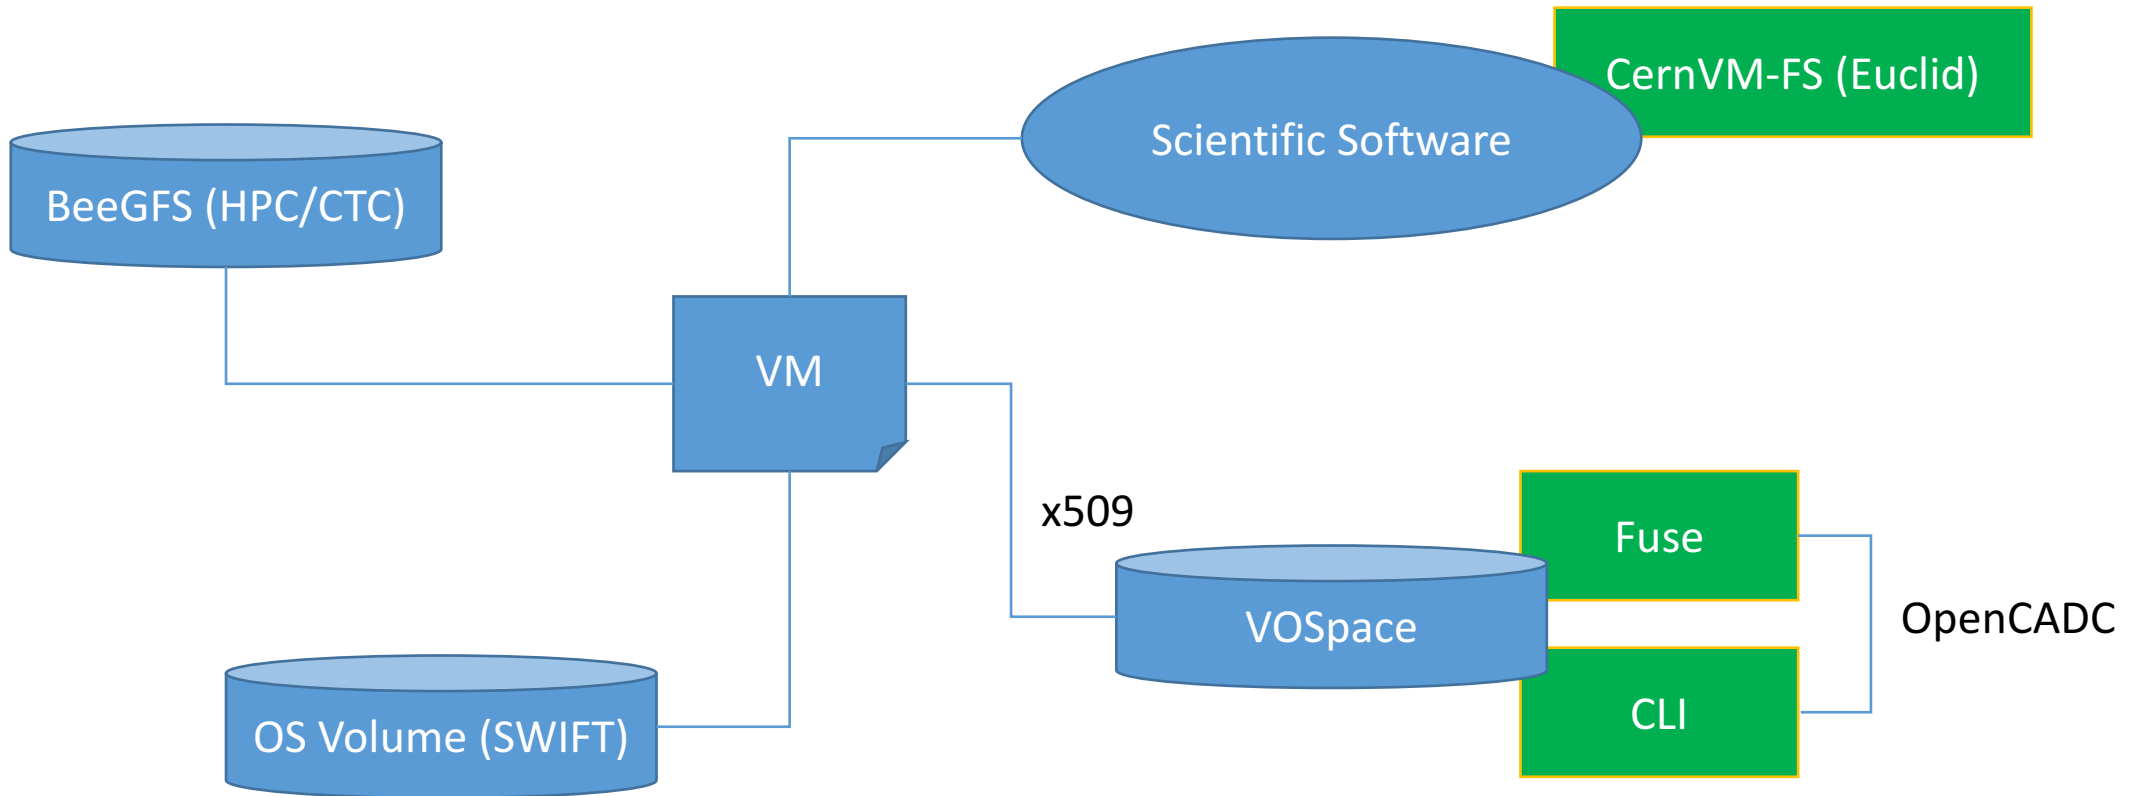
```
$ curl --cert $X509_USER_PROXY -d '{"auth":{"voms": true, "tenantName":  
"planck"}}' -H "Content-type: application/json"  
https://cloud.oats.inaf.it:5000/v2.0/tokens | python -m json.tool
```

Virtual Machines AA

- A set of VM images configured with various software (VOSpace client, data reduction, data analysis, visualization...)
- Internal authentication is based on LDAP (PAMLdap)
- When a VM is created you can add your X509 proxy

```
# (openstack) server create --flavor m1.medium --image "canfarfprod" --nic net-  
id=4b7e33c7-af8a-48a9-9e2d-29225cc15f47 --security-group default --key-name egi-  
sshkey --file /tmp/x509up_u1000=/tmp/x509up_u10000 --user-data mydata.cfg canfar-  
prod
```


Virtual Machine

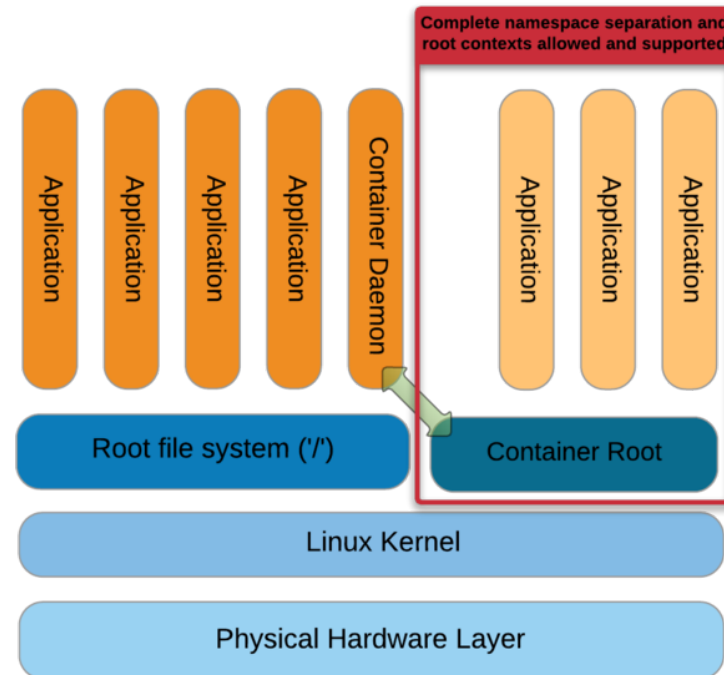


HPC/HTC

- Scope: provide a flexible environment for different applications, “integrated” with OpenStack
- We want to be able to move applications from system to system to make it close the the best HW or to the data.

- Common AA infrastructure
(I need only one user/passwd)

- Containers
 - Docker
 - Singularity
 - Shifter



Docker

- Docker well known: the first one,
- Very convenient for developers.
- Some security issues
- “... ‘simple’ Linux versioning as well as a few other issues crop up with weighty parallel file systems.”
- We integrate Docker with PBS-Professional: submit program inside the container that mounts the FS (including the parallel one) providing full access to HW.
- PBS bound docker container to the requested resources (cgroups)
- PreBuild Images for users

```
qsub -lselect=1:ncpus=2:mem=4GB -v DOCKER_IMAGE="oatsCentos" <job script>
```

Docker and Parallel IO

- Docker storage driver impact performances

Storage driver	Commonly used on	Disabled on
overlay	ext4 xfs	btrfs aufs overlay zfs eCryptfs
overlay2	ext4 xfs	btrfs aufs overlay zfs eCryptfs
aufs	ext4 xfs	btrfs aufs eCryptfs
btrfs	<i>btrfs only</i>	N/A
devicemapper	direct-lvm	N/A
vfs	debugging only	N/A
zfs	<i>zfs only</i>	N/A

Docker on Beegfs

- Docker Volume beegfs plugin to create persistent optimized volumes (<https://github.com/RedCoolBeans/docker-volume-beegfs>)
- Performances similar to direct IO



Singularity

- Application oriented: Singularity examines a program (executable) for its dependencies and add only them to the “container”
- Based on .sspec files (same idea of Dockerfile)
- it can be used for virtually any application in which mobility of compute, application appliances, or workflow distribution is important.

Singularity Usecase

- Euclid data reduction pipelines
- Development and production (actually Euclid Challenges are performed only using singularly containers)
- Does not need any driver for BeeGFS (performance close to HW)

Thanks for your attention

谢谢