Infrastructure Aspects

Cofinanciado por:
COMPETE 2020
PORTUGAL 2020
UNIÃO EUROPEIA
Fundo Europeu de Desenvolvimento Regional
Lisb@20 20
FCT Fundação para a Ciência e a Tecnologia

High level view

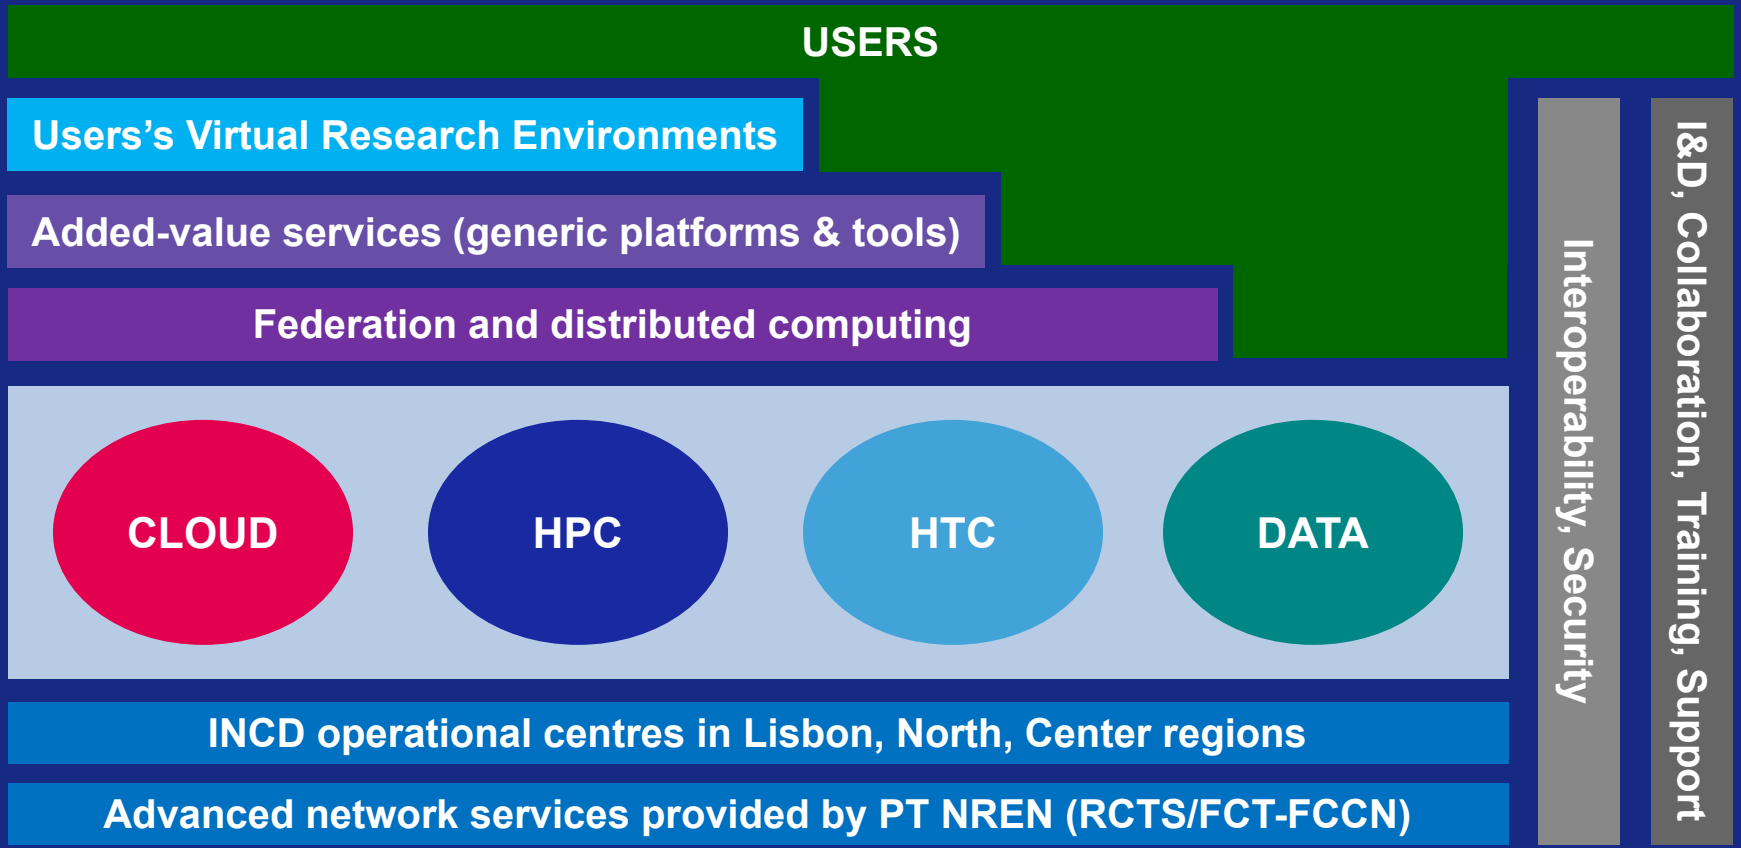

INCD operational centers

Lisbon Region I
(UPGRADE ONGOING)
HPC / HTC / Cloud / Federation
Stratus-A (cloud)
Cirrus-A (farm)
100GbE

North Region I
HTC / HPC
Cirrus-B (farm)

Center Region
(NEW TO BE DEPLOYED)
Tape storage expansion
10GbE

Lisbon Region II
Tape storage
10GbE

North Region II
(NEW TO BE DEPLOYED)
HPC / HTC / Cloud / Federation
Stratus-D (farm)
Cirrus-D (farm)
10GbE

Main services

- **HPC**
  - Slurm batch system
  - Infiniband network
  - Lustre
  - CVMFS
  - Access to GPUs A100/V100/T4
  - Containers
- **HTC**
  - Slurm batch system
  - Ethernet network
  - Lustre
  - CVMFS
  - NFS
  - Access to GPUs A100/V100/T4
  - Containers
- **Cloud**
  - Openstack
  - Virtual networks
  - Ceph block storage
  - Ceph object storage
  - Dashboard and APIs
  - Access to GPUs V100/T4

Expected compute capacity after ongoing upgrade

CPUs	RAM	Network	CORES	Place
AMD Epyc 96 CORE	512GB & 1TB	IB200 + 10GbE	4992	North
AMD Epyc 96 CORE (1)	512GB	10GbE + IB56	2304	Lisbon
AMD Epyc 64 CORE (2)	256GB & 512GB	10GbE + IB56	2112	Lisbon
INTEL Xeon 24 CORE	64GB	1GbE	2088	Lisbon
INTEL Xeon 16 CORE	32GB	IB56 + 1GbE	2560	North
Total	62912GB		14056	

(1) Nvidia A100
(2) Nvidia V100 and Nvidia T4

Expected storage capacity after ongoing upgrade

Storage Type	Capacity (RAW)	Place
Lustre	4680TB	Lisbon
Lustre	1920TB	North
Lustre	824TB	North
Ceph (block + object)	1392TB	Lisbon
Ceph (block + object)	2880TB	North
CVMFS	1000GB	*
Tape LTO9	20160TB	Center
Total	32324TB	

Computing Farms

- Infrastructure
  - Slurm batch systems
  - Lustre file system
  - CVMFS (EGI, LHC, INCD)
- Software
  - Modules
  - udocker and singularity
  - Some commercial sw e.g. Intel OneAPI
- Management
  - Nagios
  - Prometheus
  - Ansible
  - Own scripts

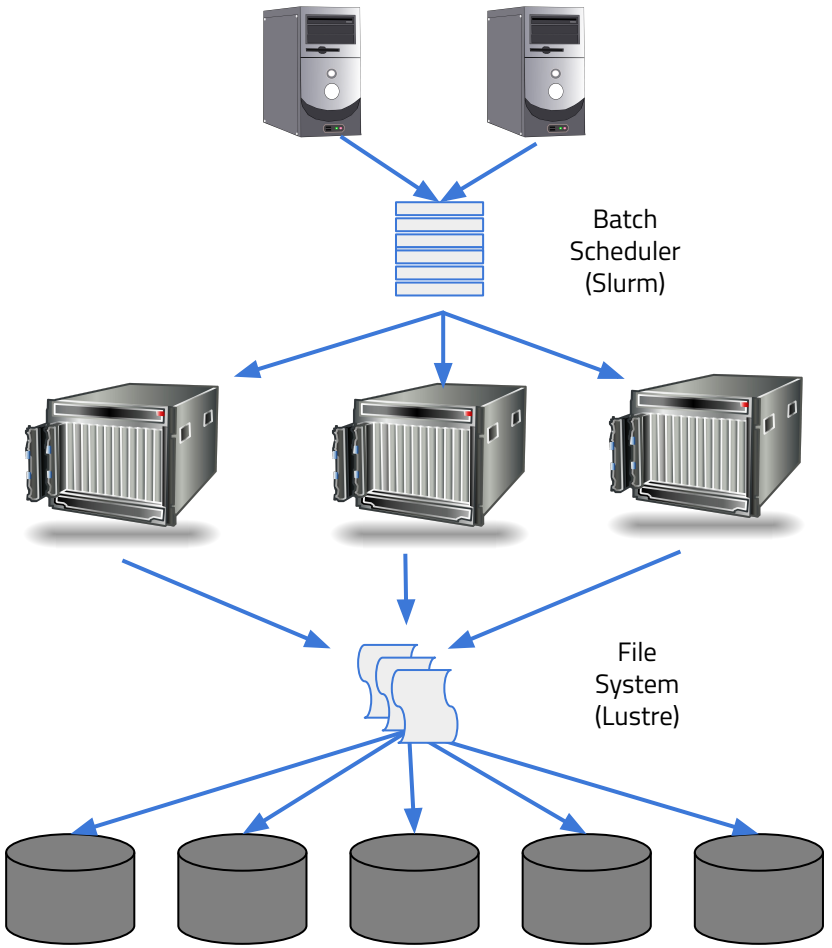

Grid

- Compute
  - ARC-CE for job management
- Storage
  - XRootD
  - StoRM (gridftp, webdav)
  - CVMFS
  - Squid (CMS)
- Management
  - Apel
  - Argus
  - BDII
  - Perfsonar
  - Netbox

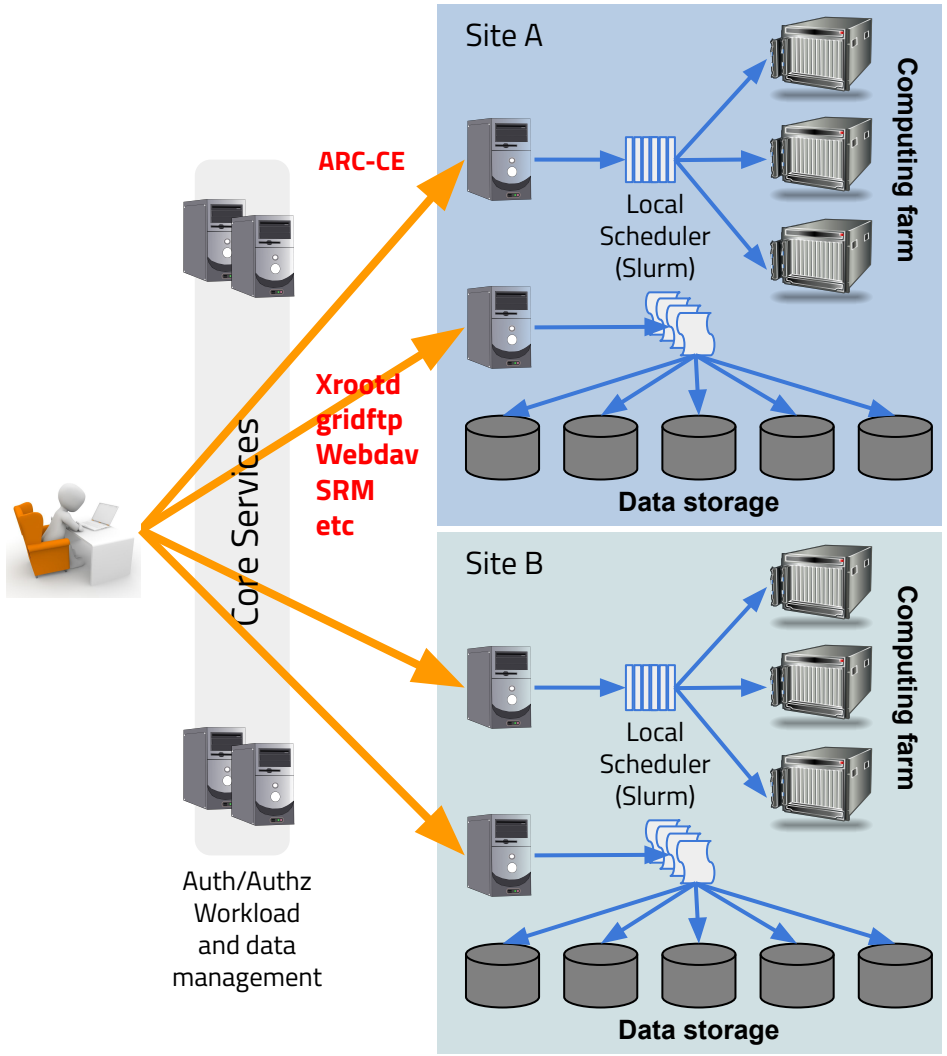

Access to GPUs

- GPUs are highly requested
  - Insufficient resources for the demand
- Nvidia V100 and T4
  - Mostly and preferably via the farm either HPC or HTC
  - Access either natively or using udocker or singularity
- Some older Nvidia GPUs
  - For training, testing, development and small projects
- Expecting Nvidia A100
  - Will be made available in the same way

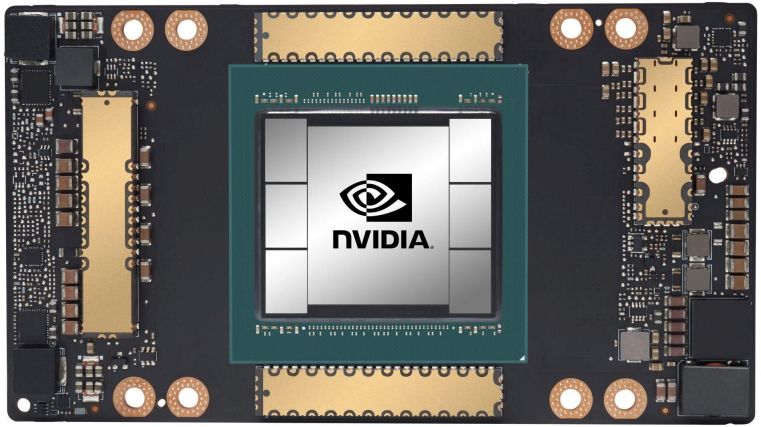

KVM

Virtualization

- Simple solution for virtualized datacenter services
  - KVM
  - LXC
  - DRBD

- KVM
  - hardware virtualization
- LXC
  - lightweight virtualization
- DRBD
  - Replicate volumes across hypervisor nodes
  - Provides redundancy
- Hearbeat
  - High availability

Virtualization

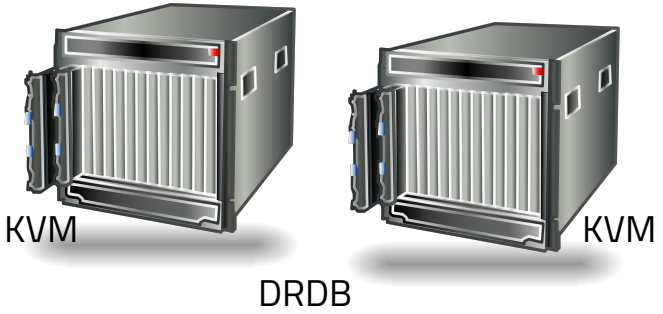

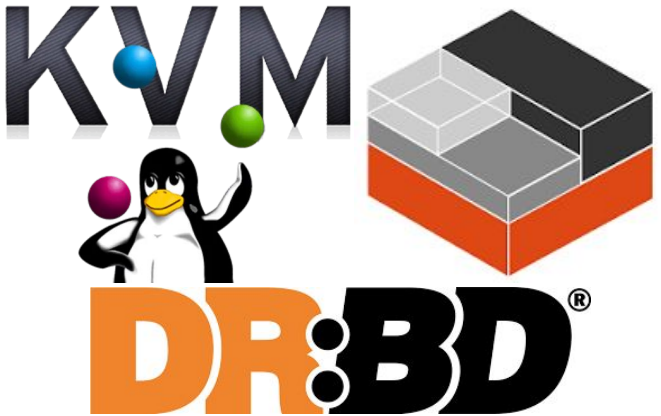

Cloud

- Openstack
  - Nova, Glance, Cinder, Neutron
  - Using KVM for virtualization
- Ceph
  - 3x replicas
  - Block storage
  - Object storage
- Network
  - Overlay GRE & Vxlan
  - Provider networks (VLANs)
  - SRIOV
- Monitoring
  - Nagios
  - Own tools for more complete accounting including storage
- Compute
  - AMD Epyc processors
  - Some GPUs mainly for testing

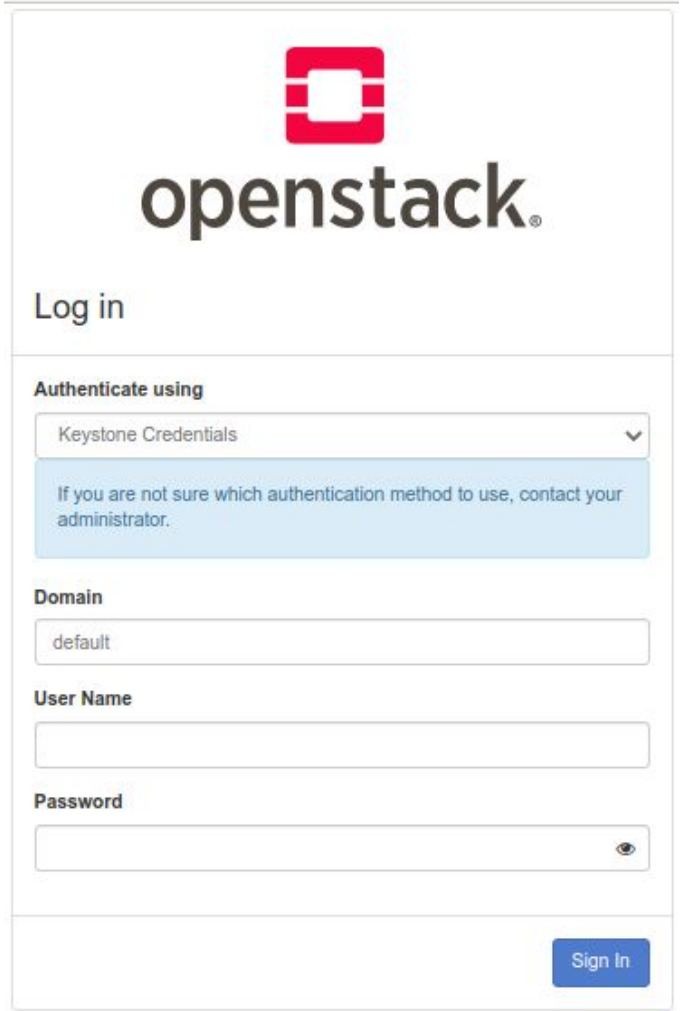

Cloud

- Openstack
  - Nova, Glance, Cinder, Neutron, Horizon, Keystone
  - Using KVM for virtualization
- Ceph
  - 3x replicas
  - Block storage
  - Object storage (S3 & Swift)
- Network
  - Overlay (GRE & Vxlan)
  - Provider networks (VLANs)
  - SRIOV
- Monitoring
  - Nagios
  - Own tools for more complete accounting including storage

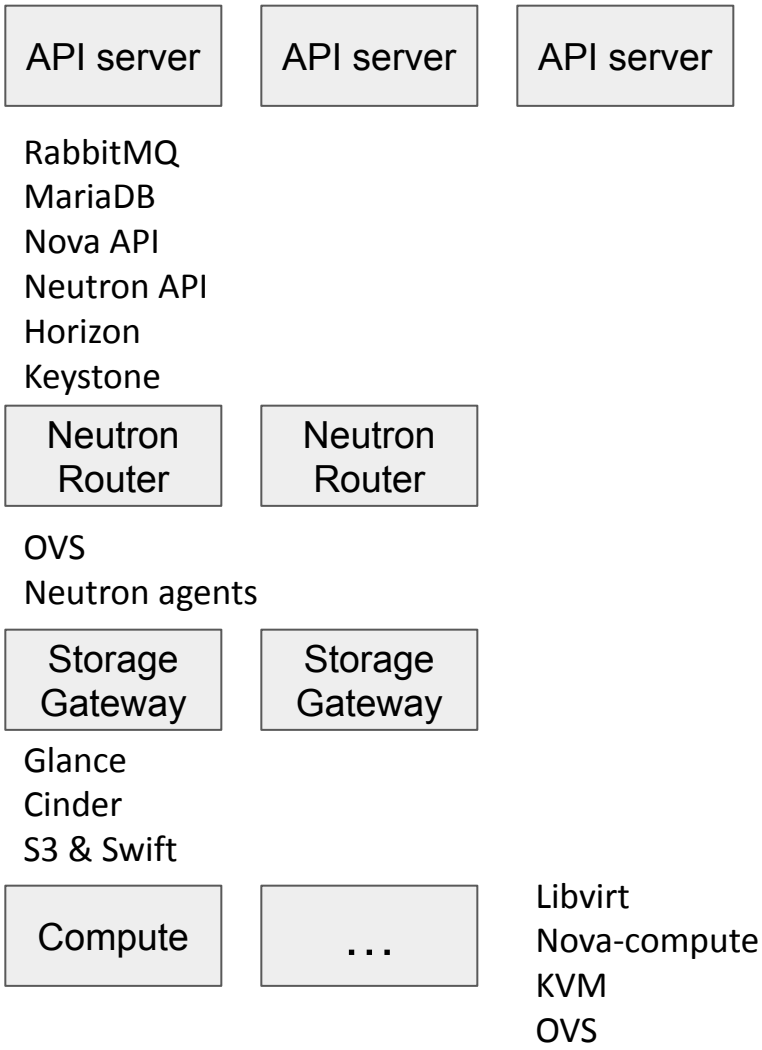

Federated cloud

- Information provider registered in GOCDB
- CASO + APEL for accounting
- Integration with the EGI Fedcloud
- EGI Check-in
- Using OIDC
- Also had support for the national NREN federation RCTS AAI

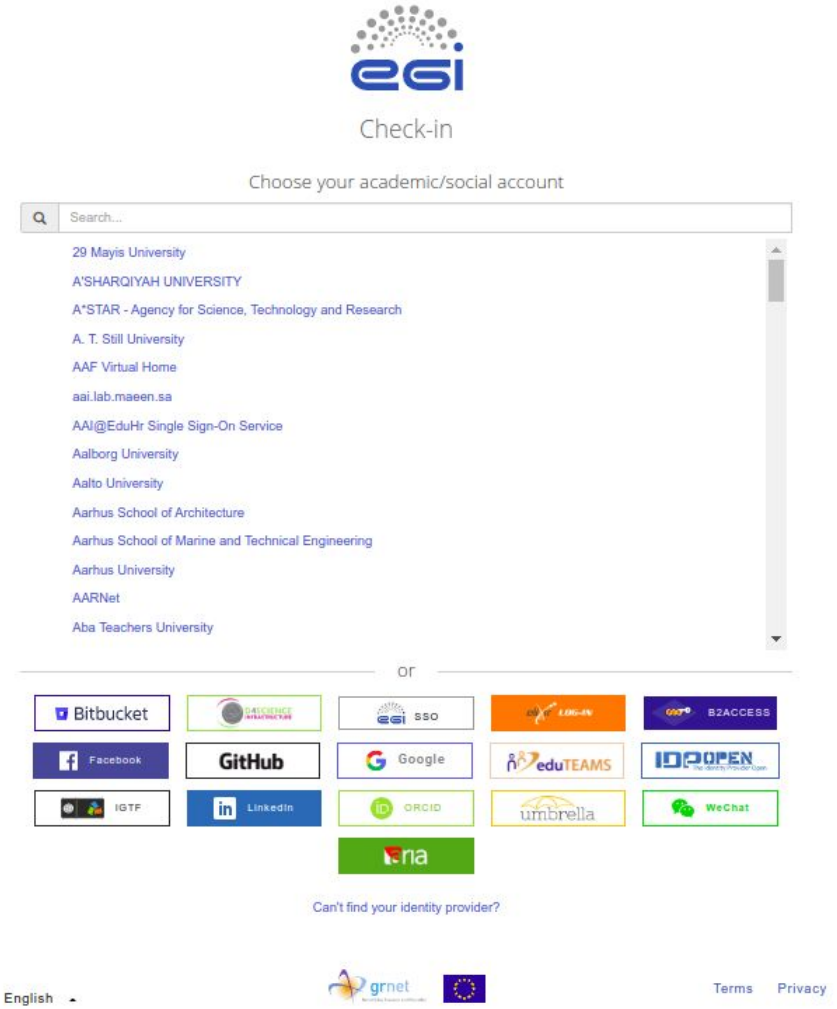

Containers

- Farm
  - udocker
  - singularity/aptainer
- Cloud
  - docker
  - docker-compose
  - kubernetes
- Management
  - portainer
  - kubernetes

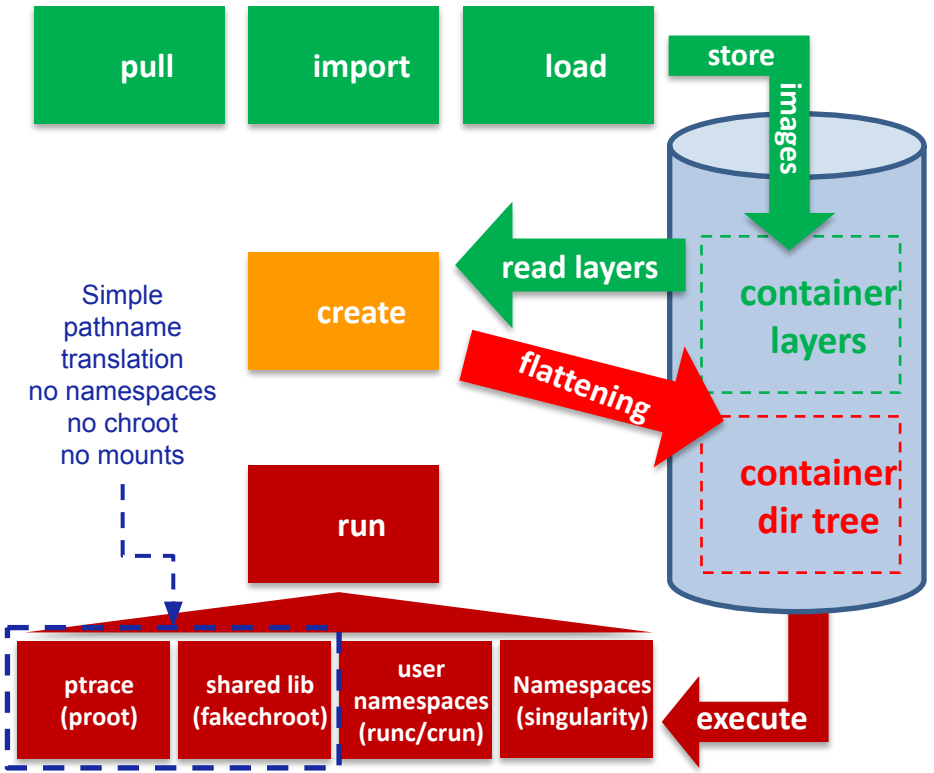

https://www.incd.pt

helpdesk@incd.pt

Thank you

Cofinanciado por:
COMPETE 2020
PORTUGAL 2020
UNIÃO EUROPEIA
Fundo Europeu de Desenvolvimento Regional
Lisb@20 20
FCT Fundação para a Ciência e a Tecnologia