



Portuguese Distributed Computing Infrastructure

Advanced Computing

INCD
Computing and Data for Science and Education

Jorge Gomes  

<http://www.incd.pt>

Cofinanciado por:





Portuguese Distributed Computing Infrastructure

Services: computing, data storage, data processing

Promote: resource sharing, common interoperable solutions

Target: scientific and academic communities, infrastructures and projects

Interface: International infrastructures (EGI, IBERGRID, WLCG, EOSC, EuroCC, ...)




Cloud Computing

cloud computing



HTC Computing

high throughput
computing (GRID)



HPC Computing

high performance
computing

INCD is a legal entity backed by a consortium: INCD Association. LIP, LNEC, FCT-FCCN

Researchers and Thematic Scientific Infrastructures

INCD
user facing services

Support, Training, Consulting	Linux Containers	Map reduce as service	Others SaaS, etc ...
Cloud IaaS federated & hybrid	Grid Computing	Data Files and Objects	Data Protection
IaaS Cloud direct access	farm HPC and HTC	databases	Specialised Hardware GPGPUs

INCD flexible
underlying
infrastructure

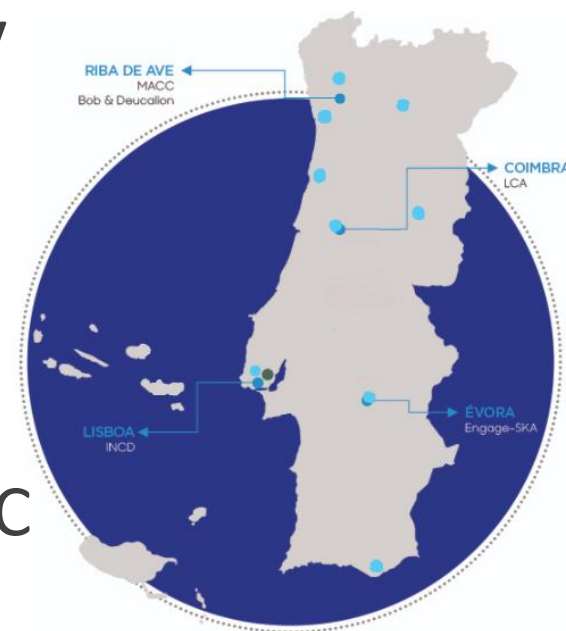
Processing Cloud + Farms	Data Storage and Data Access	Outsourcing & Integration
Physical infrastructure at multiple geographic locations Computing, Storage, Networks		

Integration and Access to International Digital
Infrastructures



Portuguese Advanced Computing Network

- Portuguese Advanced Computing Network
- Coordinated by the National Foundation for Science and Technology (FCT)
- Joins HPC centers in the country
- Supports the national calls for advanced computing projects
- Providers:
INCD (LX, North), LCA, HPC-UE, MACC





Collaboration with EuroCC NCC



ABOUT EUROCC ▾

HPC IN PORTUGAL ▾

NEWS AND EVENTS

TRAINING AND SUPPORT ▾

CONTACTS



EURO

Competence Centers for HPC

A source of information,
training and collaboration.

LIP in collaboration with INCD is one of the national organizations in the EuroCC National Competence Center



IBERGRID



IBERGRID





HPC Cirrus-A @ Lisbon

- Compute nodes
 - AMD EPYC based 75xx
 - 64 CPU cores 256GB RAM
 - 64 CPU cores 512GB RAM
 - 96 CPU cores 512GB RAM
 - Supermicro and Dell servers
 - Some nodes with GPUs
 - CentOS 7
 - Slurm workload manager
 - Infiniband
- Storage
 - Lustre on ZFS
 - Lustre on LDISKFS
 - CVMFS
- Access
 - Currently via SSH
 - Same cluster as grid but different partition
 - Conventional integration via UMD is possible
 - Compute nodes have outbound connectivity





HPC Cirrus-B @ Minho

- Compute nodes
 - Intel Xeon based
 - Dell C8000 / C8220X
 - 16 CPU cores 32GB RAM
 - E5-2680 @ 2.70GHz
 - CentOS 7
 - Slurm workload manager
 - Infiniband
- Access
 - Currently via SSH
 - Started to integrated in EGI
 - Compute nodes have outbound connectivity
- Status
 - Equipment will be moved to a new datacenter being built
- Storage
 - Lustre on ZFS
 - CVMFS





HPC Cirrus-C @ Coimbra

- Compute nodes
 - Intel Xeon based
 - Dell R720
 - 24 CPU cores 192GB RAM
 - E5-2695v2 @ 2.40GHz
 - CentOS 7
 - Slurm workload manager
 - Infiniband FDR
 - Commercial software
- Access
 - Currently via SSH
 - Not integrated in EGI
- Status
 - External cluster integrated
 - Collaboration with ISEC/LACED
- Storage
 - NFS
 - Infiniband IPoIB





HTC/HPC in the Cloud @ Lisbon

- Since we have an Openstack IaaS cloud infrastructure
 - Possibility of using Cloud nodes for HTC/HPC workloads
- Usually only for very specific situations:
 - e.g. many of our GPUs are in cloud compute nodes and can be made available as virtual compute nodes for Slurm
- Two approaches:
 - Slurm Compute nodes in the Openstack cloud (our preferred)
 - Instantiation of virtual infrastructure (by the user)
- Infiniband in the Cloud
 - We also have cloud compute nodes with Infiniband FDR
 - Mainly to experiment with HPC in the cloud not used in production



Access - login

- Login on head nodes
 - Currently login via SSH with SSH keys
 - Looking at the EOSC-Synergy approach for federated login with SSH (see more later in Marcus presentation)
- Currently recommending
 - openssh
 - sshuttle (very good but requires root on client side)
- See how-to
 - <https://wiki.incd.pt/books/how-to-access>



Access - data

- Via SSH
 - scp, sftp, sshfs
- Webdav in Lisbon
 - Using grid credentials
 - Using OAuth2/OIDC tokens Macaroons
- Storage Elements in Lisbon
 - StoRM
 - XRootD



Software

- Using conventional Environment Modules
 - dynamically loading software packages
 - we add regularly new compilers and libraries
 - we support Intel compilers and libraries
 - we provide software compiled for Intel and AMD
 - as usual use **module avail**

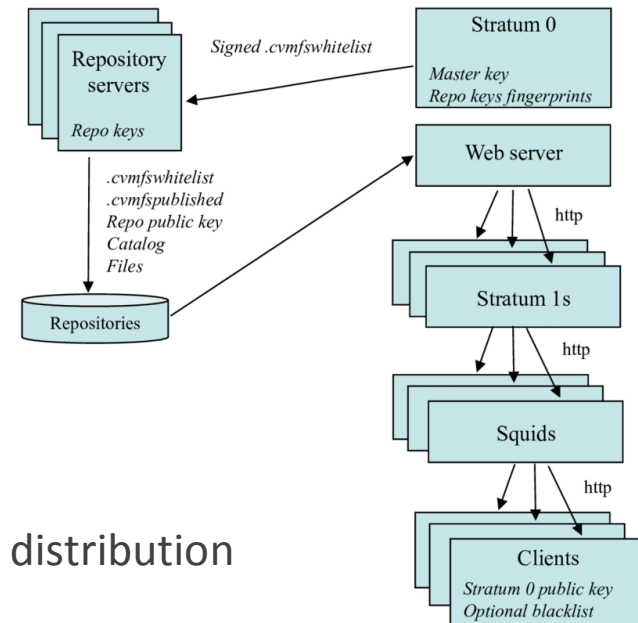
Maybe not so conventionally ...

- our software is on CVMFS
- modules is setup to use the software on CVMFS
- Shared group areas for software also supported



Software distribution - CVMFS

- Wide area network read-only POSIX filesystem
 - Developed by CERN
 - Meant for worldwide software distribution
 - Scalable, reliable and low-maintenance
 - Aggressive caching and reduction of latency
- Architecture:
 - Hierarchy of https caches
 - FUSE module for Linux
 - Data integrity by cryptographic hashes
- At INCD:
 - Used for software, small RO files and containers distribution
 - Used across the sites (Lisbon and Minho)
 - Stratum 0 and Stratum 1 CVMFS servers





udocker



- Running Linux containers in user space
 - Does not require privileges for installation
 - Does not require privileges for use
 - Does not require namespaces
 - Does not require recompilation
 - Works both on older and newer kernels
- Suitable for
 - HTC, HPC and accelerated computing
 - Grid computing and computing clusters
 - Batch and interactive systems
- Empower users to execute applications in containers

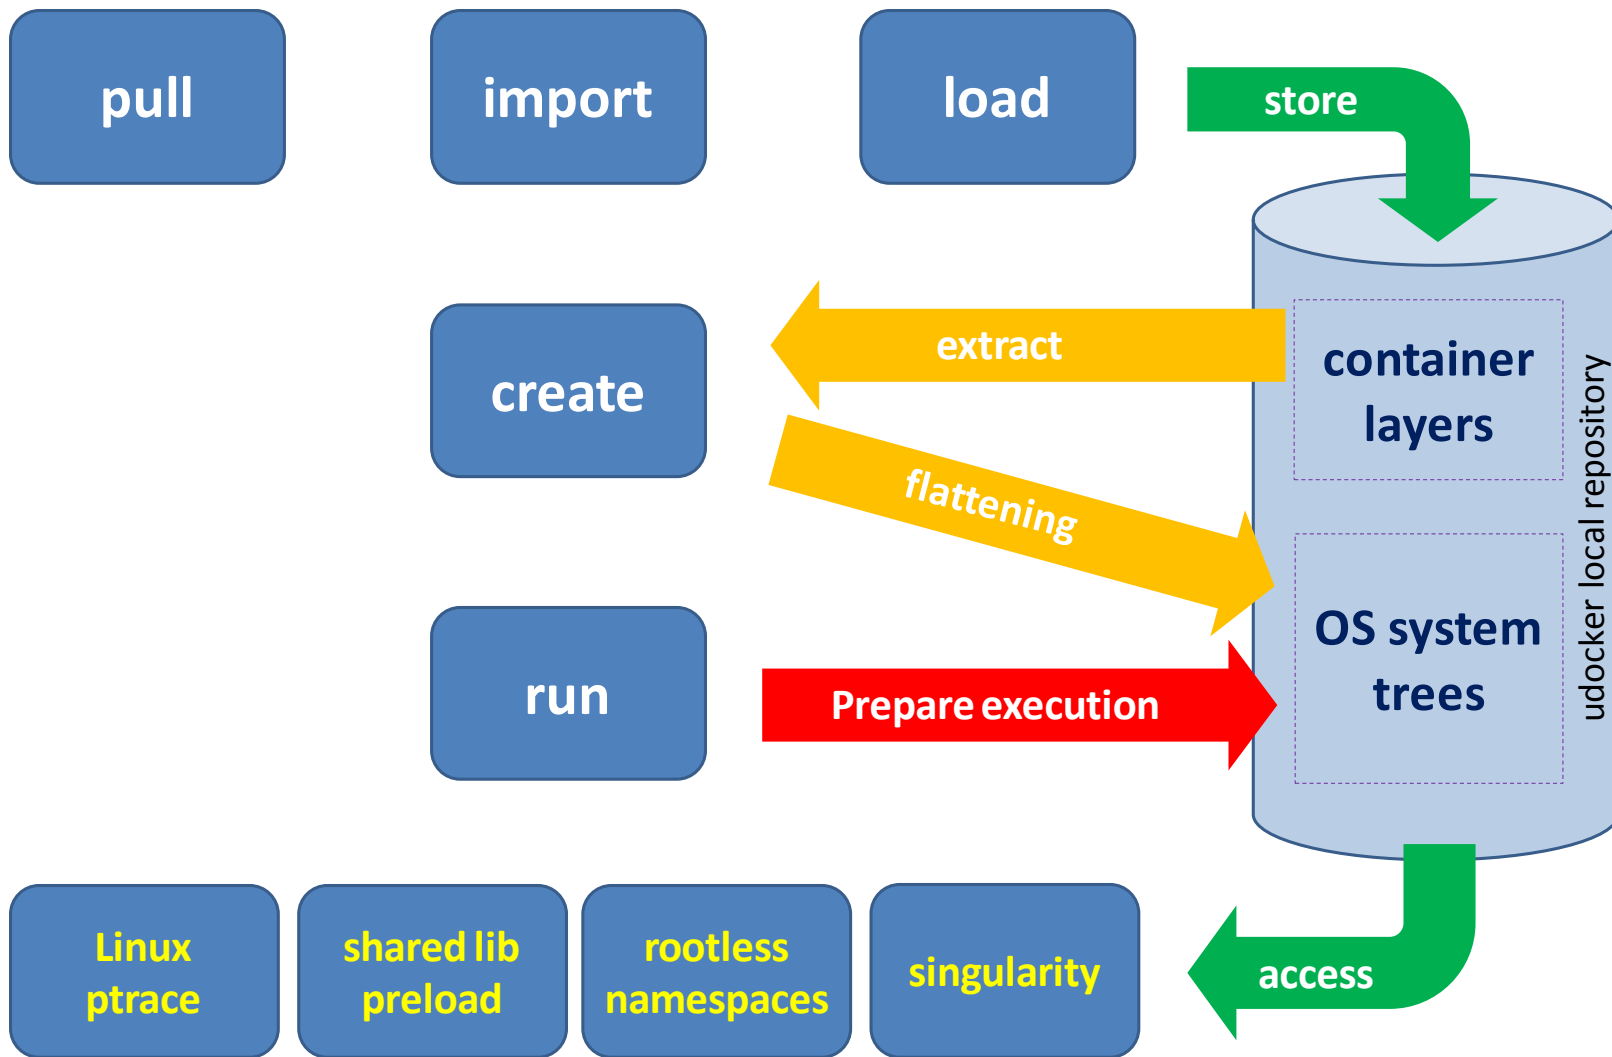


EOSC-hub





udocker





udocker



- Supports multiple execution techniques
 - Selectable via execution modes

Mode	Base	Description
P1	PRoot	PTRACE accelerated (with SECCOMP filtering) ← DEFAULT
P2	PRoot	PTRACE non-accelerated (without SECCOMP filtering)
R1	runC / Crun	rootless unprivileged using user namespaces
F1	Fakechroot	with loader as argument and LD_LIBRARY_PATH
F2	Fakechroot	with modified loader, loader as argument and LD_LIBRARY_PATH
F3	Fakechroot	modified loader and ELF headers of binaries + libs changed
F4	Fakechroot	modified loader and ELF headers dynamically changed
S1	Singularity	where locally installed using chroot or user namespaces





udocker




Container:

- Latest GPU version of Tensorflow (from Docker Hub).
- Train a model to recognize handwritten digits (the MNIST data set).



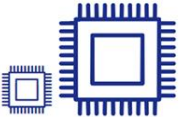
<https://github.com/tensorflow/models.git>

EXECUTION TIME





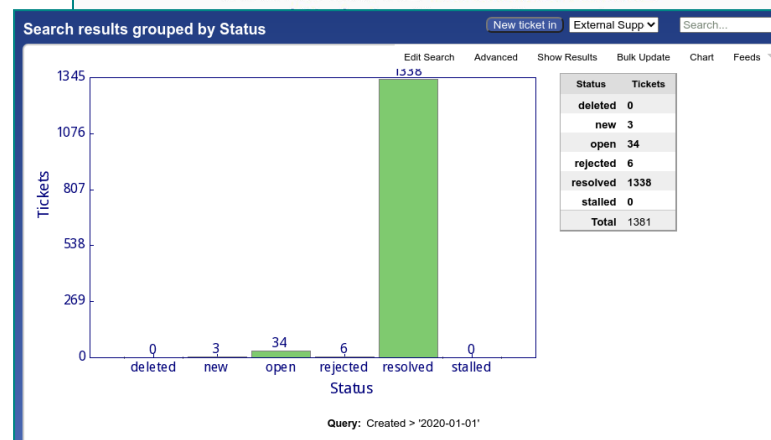
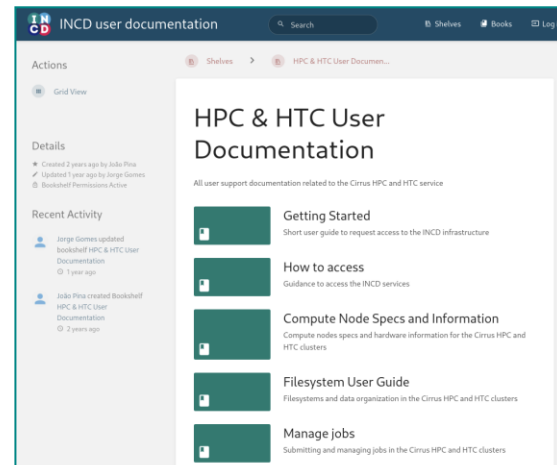
Services and how to use them

Service type	Access	Software Environment
 <p>Cloud Computing</p>	<ul style="list-style-type: none">• Through web dashboard• Through command line and APIs• Through orchestrators and APIs• Federated Cloud (EGI, IBERGRID)	<ul style="list-style-type: none">• Pre-defined images• Bring their user images• KVM, LXD, Docker, udocker, etc
 <p>HTC Computing</p>	<ul style="list-style-type: none">• Login SSH and command line• Integration with workflows• Grid (EGI, IBERGRID, WLCG)	<ul style="list-style-type: none">• CVMFS and modules• Shared software space• udocker, singularity
 <p>HPC Computing</p>	<ul style="list-style-type: none">• Login SSH and command line• Integration with workflows• Grid (EGI, IBERGRID)	<ul style="list-style-type: none">• CVMFS and modules• Shared software space• udocker, singularity• Low latency interconnect



Support

- Web site
 - www.incd.pt
 - Access request
 - Usage and Access Policies
- Documentation
 - BookStack
 - <https://wiki.incd.pt>
 - Sections on
 - HPC & HTC
 - Cloud
- Helpdesk & Contact
 - helpdesk@incd.pt
 - Request Tracker (RT)



OPENCoastS.pt

- **On-demand circulation forecast systems for European coastal areas.**
- Generates daily forecasts of water levels, vertically averaged velocities and wave parameters over the region of interest for 48 or 72 hours
- Numerical **MPI** simulations of all physical processes as-a-service.

- Forcings are provided by:

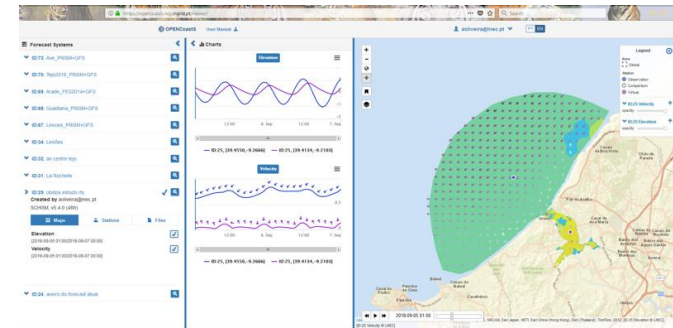
- [NOAA / GFS](#)
- [MétéoFrance](#)
- [PRISM2017](#)
- [FES2014](#).

- Tide gauge data:

- [EMODnet Physics](#)

- LNEC service jointly with LIP & INCD

- Also part of EGI-ACE
- <https://opencoasts.ncg.ingrid.pt>





www.incd.pt

helpdesk@incd.pt

Thanks

A screenshot of the INCD website homepage. The top navigation bar is dark blue with white text for 'THE INFRASTRUCTURE', 'SERVICES', 'NEWS', 'USE CASES', and 'TRAINING', along with a globe icon and 'EN'. The main header features the INCD logo and the text 'Infraestrutura Nacional de Computação Distribuída'. Below this is a large banner with the text 'Computação para a Ciência e para o Ensino'. The main content area states 'The ideal solution for Researchers, Teachers and Students with computing needs.' At the bottom, there are three service highlights: 'Tailor made solutions' (with a person icon), 'Close and permanent monitoring' (with a handshake icon), and 'Open access to the community' (with a padlock icon).