# Computing **LIP Computing Activities Status and Perspectives** LIP workshop 2020













## **Organisational Structure**

- Distributed Computing and Digital Infrastructures
  - LIP-Lisbon Researchers and Technical staff
  - INCD Collaborators in Lisbon and Minho
  - Evolution of the LIP-Lisbon computer centre

- Advanced Computing
  - LIP-Minho Researchers from U.Minho and P.Bragança
  - Joined LIP when the Minho centre was created



### **Advanced Computing**

- Software for High Energy Physics
  - Development and optimization of code
- ATLAS High-Level Trigger
  - Algorithms parallelization in GPUs
- Training in Scientific Computing
  - Advanced training (students)
  - Technical training (Linux, C++)
- Local Compute Infrastructure
  - LIP Minho cluster
  - CPU/GPU platform for Machine learning development
- Participating in the project POCI-01-0145-FEDER-029147
  - Understanding Big Data in High Energy Physics





### **Advanced Computing**

- António Esteves Researcher U.Minho, application of machine/deep learning techniques
- António Pina Researcher U.Minho, application performance analysis, dynamic tracing
- José Rufino Researcher P.Bragança, parallelization strategies for GPU based algorithms
- Tiago Duarte MSc Distributed Training of Deep Neural Networks
- Tiago Fernandes MSc, Study, evaluation and application of the HPX platform



### **Distributed Computing and Digital Infrastructures**

- LIP general IT services
  - Core institutional services, network services, security, and also
    LIP-Lisbon specific services
- Web development and graphics design
  - Web applications, web development, graphics design, multimedia, support to events and outreach
- Scientific computing and data processing
  - Shared compute and data services, both for LIP and other organizations (via INCD)
- R&I projects
  - Participation in european (H2020) and national projects aimed at developing competences and capabilities





### **Distributed Computing and Digital Infrastructures**

- Jorge Gomes PI, INCD coordination, projects, development, network
  - Mário David cloud, SQA, GPU computing, projects
- João Pina
  Tier-2, EGI/EOSC software management
- João Martins fabric management, compute, data storage
  - Nuno Dias desktops, notebooks, security, network services
- Hugo Gomes web development, events, multimedia
- Carlos Manuel web development, graphic design
- Samuel bernardo development, AAI, cloud, projects
- José Aparicio purchases, repairs, desktops, hardware and facilities
- Dinis Monteiro web development for INCD
- Henrique Carvalho web development, multimedia
- André Vieira INCD, fabric management, monitoring, cloud, storage
- Catarina Ortigão
  - Zacarias Benta

João Machado

- INCD, project management, communication, outreach INCD, advanced computing services and infrastructure
- INCD, advanced computing services and infrastructure



### **Projects and funding**

LIP			
EOSC-hub HR	H2020	Jan 2018	36M
EOSC-hub VA	H2020	Jan 2018	36M
DEEP-Hybrid-DataCloud	H2020	Nov 2017	30M
INCD P2020	P2020 Infrastructures Roadmap	Jul 2017	36M
EOSC-Synergy	H2020	Sep 2019	30M
BigHPC	University of Texas at Austin (UTA)	to be started	36M
INCD			
INCD P2020	P2020 Infrastructures Roadmap	Jul 2017	36M+12M
Géant Cloud Pilot	FCT contract	May 2018	12M
Protocolo RICA	FCT contract	Dec 2018	extended
EOSC-Synergy	H2020	Sep 2019	30M



### **IT infrastructure context**





BOB @ REN Riba-de-Ave HPC cluster from FCT Offered by TACC 20% is managed by INCD+LIP LIP helping in the storage

#### NCG @ LNEC

Sala-Grid Datacenter New INCD equipment Old INGRID equipment LIP specific equipment

> LIP users LIP Tier-2 INCD users



U. Lisboa @ reitoria Mail, web, VMs Home directories

LIP @ 3Is Tape library Local network+services





U. Coimbra Old Datacenter Physics Dep. Small LIP facility Point-of-Presence

ISEC - Laced Minerva Cluster 20x Dell R720 Infiniband Miguel Couceiro



### **IT network infrastructure**



- LIP IT services
- Worldwide LHC Computing Grid
- National Distributed Computing Infrastructure



# Web and design

#### 1. Web design and development

- LIP website and INTRANET continuous development, improvement, and new user applications
- Groups/Projects/Events web pages
- Web content managers/editors
  - <u>NIKO</u> Content manager + Editor
  - <u>InlineEdit</u> content editor edit pages directly/real-time in the website (no content management system needed)
- **UI starter kits** selection/preparation for websites and applications

#### Near future

- **portalIT** IT management portal
- ID LIP identity service, users password change
- **portalIT / LIPDB simple-API** access LIPDB data easily from webservices
- **Travel and Acquisitions Platform** development of new application according to LIP requirements

#### 2. Internal applications for LIP users

- Indico events management
- Event registration internal Admin for LIP secretariat
- Travel and Acquisitions Platform
- LIP-Users wiki
- New USER/EVENT online registration
- Projects Timesheets
- IDPASC website and Calls
- Room Booking (for Lisboa facilities)
- Email vacation setup (for Lisboa users)
- Outreach Moodle
- Support on RCTSaai applications - Colibri, Videocast, Filesender, etc..
- Databases design and management
- **cLIP** internal newsletter (with LIP-ECO)
- LIP Technical Report generator
- LIP-ECO Events webpages and registration (Masterclasses, Esc. Prof. CERN)
- New Wiki system for Computing group and INCD
- LIP Seminars web streaming using RCTSaai (VideoCast)

#### 3. Graphic Design

- LIP Annual Reports
- LIPNews bulletin
- LIP PowerPoint presentations and templates
- RollUps for LIP and for INCD
- DI4R conference signs and communication materials
- LIP events design support (badges, poster, etc..)



## **Authentication and Authorization**

LIP FIAM Federated Identity and Access Management.

Active development

- Registry
- Identity Management System
- APIs to connect federated applications
- Collaboration
  tools

Later

• Users portal

Computing



## WLCG Tier-2 accounting





# WLCG Tier-2 activities and issues



- Tier-2 specific activities
  - Integration of new INCD compute nodes and storage
  - Upgrade and reorganization of Lustre storage system
  - Improvement of gridftp, XRootD and webdav gateways
  - Increased pledge ~25% per experiment (+ 800 HS06, + 60 TB)
- Issues
  - The increase of capacity is largely insufficient in face of needs
  - A considerable part of the capacity is from older systems



# WLCG HL-LHC challenge







• the main challenge is storage;

Computing

but computing requirements grow 20-50x





# **INCD** usage





## Indicators Jul-2017 - Jan 2020



Groups by Organization			
Universidade Aberta	3	IST – Univ Lisboa	4
Universidade do Algarve	3	MARE - Univ Lisboa	1
CCMAR	3	FCUL – Univ Lisboa	2
CBMR	1	ISEG	1
CICECO	2	CE3C – Univ. Lisboa	1
Univ de Coimbra	3	Universidade do Minho	2
BIOISI – Univ. Lisboa	3	Univ Nova de Lisboa	6
CIBIO-InBIO	1	Fac Medicina UL	1
IGC	2	REQUIMTE – Univ do porto	4
IMM	1	IPMA	1
INESC-ID	3	CIBIT	1
LIP	17	ISEC	4
LNEC	7	ISEL	2
IDL	1	ISCTE	4
ISA – Univ. Lisboa	2	UCIBIO	5
ISPA	1	INL	1
ITQB	3	Inst Ricardo Jorge	1
FCSH	2	Cicant - Univ Lusófona	1
FCT-FCCN	1		

Computing

Metrics		
User Groups	101	
Projects supported	58	
MsC thesis supported	16	
PhD thesis supported	15	
Conference presentations	81	
Articles direct	150	



### **Infrastructure** activities

- Hardware added in 2019
  - Compute nodes
    - 25x compute nodes
    - 8x larger nodes > capacity
  - Storage Nodes
    - 6x Lustre
    - 6x Ceph
  - 10x Service nodes
  - 3x Network 25/100 GbE switches
  - 1x Infiniband switch
  - 3x racks, PDUs & mgt switches
  - 2x Nvidia GPUs
- datacenter improvements

omputi

- New PDUs and recabling
- Reorganization of servers
- Comprehensive inventory

- Service improvements/developments
  - Lustre upgrade
  - Ceph upgrade
  - Openstack upgrade
  - New equipment database Netbox
  - Monitoring and accounting
    - Prometheus
    - Nagios
    - Grafana/Graphite/Carbon
    - ELK
    - Influx DB
    - XDMoD
  - Containers
    - Portainer
    - Kubernetes
  - Documentation
    - Bookstack
    - Complete reorganization
  - Source code and packages
    - New gitlab deployments
    - Jenkins CI/CD
    - Nexus OSS

### **High Performance Computing**

- BOB High Performance Computing cluster
  - TACC offered to Portugal (FCT) part of the Stampede
  - Installed at the REN datacenter in Riba de Ave
  - INCD is managing 20% of the capacity 2560 cores
  - INCD is helping in the deployment of the Lustre filesystem
  - Limited by network connectivity and cooling
- Advanced Computing Network (RNCA/RICA)
  - National Advanced Computing Strategy (INCoDe.2030)
  - New type of FCT projects just for computing capacity
  - Pilot call of FCT computing projects possibly in April
  - INCD will be one of the resource providers
- EuroHPC

- EuroHPC will establish National Competence Centres
- FCT is partner in the competence centres H2020 proposal
- LIP is a third party in the project proposal (2M€ for PT)







- European Open Science Cloud (EOSC)
- EOSC-hub



# **European Grid Infrastructure**

### EGI.eu:

- 24 members both countries and FIROs
- Created by national infrastructures to coordinate international computing activities
- EIROS: CERN, EMBL ٠
- Portugal is member through FCT •
- Associated infrastructures in Asia, Austrália, Africa, North and South America

### **Distributed computing infrastructure:**

- Multidisciplinar
- Grid, <u>Cloud</u>, <u>Data</u>, HTC, HPC, platforms, ...
- > 47 participating countries
- > 230 user communities
- > 360 resource centers
- > 1.000.000 CPUs

- > 740 PB disk + tape
- > 4.4 billion CPU hours in 2018



**Enabling new discoveries** 1700 publications open access per year





Rainer Weiss



**Gravitational Waves** LIGO

Barry C. Barish Prize share: 1/4

Kip S. Thorne Prize share: 1/4



## LIP in EGI

- LIP is partner of all EGI flagship projects
  - InSpire (2010-2014), Engage(2015-2017), EOSC-hub (2018-2020)
- Roles as resource provider:
  - Liaison with the Portuguese and Iberian grid infrastructure
  - Infrastructure coordination together with, policy, training, dissemination and relations with user communities
- Core activities:
  - Global Coordination of the middleware rollout process
  - Quality criteria definition and validation for the middleware
  - Coordination/management vulnerabilities and intrusions both in the region and at European level
  - Regional support for users and sites for all Iberian region
    - Council member representing FCT since 2010

# Computing

### **IBERGRID**

### IBERGRID

- > 12 centres
- > 24 K CPU cores
- > 30 PB

### Portugal

- INCD-A Lisboa (NCG)
- INCD-B Norte (RDA)
- INCD-C Coimbra (UC)
- LIP
- U. Porto
- ISEC



- IBERGRID federation of Iberian digital infrastructures in Portugal and Spain
- Common participation in the European Grid Infrastructure (EGI) and EOSC activities
- Support to the participation in common R&I ATLAS, CMS, LifeWatch, INSTRUCT, AUGER,





### Project Overview

20 digital research infrastructures, EGI, EUDAT and INDIGO-DataCloud

*services, software and data for advanced, data-driven research & innovation*  EOSC-hub: Integrating and managing services for the European Open Science Cloud

### H2020 Grant Agreement ID 777536

Tot budget: €33,287,542 EU Budget € 30 000 000 100 Partners, 76 beneficiaries 3830 PMs, 106 FTEs +150 staff involved

Coordinator STICHTING EGI

> **Consortium** 100 partners 53 countries

36 Months Jan 2018 – Dec 2020

### **EOSC-hub** The 'federating core' and the wider EOSC

	EOSC Service Portfolio	Thematic services: e.g. RI services, thematic e-Infra services	Common s e.g. D manage analy visualis	services: Data ment, tics, ation	Data : e.g. ESFRI data sets, certified data sets, data catalogue listings	EOSC Data Portfolio	
	EOSC	Shared resources: Generic data processing services, public good data EOSC			EOSC		
Federating Core	Hub portf Key Access E Services, e. Marketplac	<b>olio:</b> nabling g. AAI, ce etc.	Co fr Rules c Servico Syst	ompliance amework: of Participation, e Management tem, policies	Core		

# LIP in EOSC-hub

- Main activities:
  - Task 4.6 coordination under federation operations:
    - Middleware quality assurance and rollout process
  - Thematic Service: OPENCoastS
    - On-demand Operational Coastal Circulation Forecast Service for the North Atlantic
    - Partners: LNEC, LIP, CNRS (University of La Rochelle -LIttoral, ENvironment and Societies (LIENSs) and UNICAN
  - udocker: support and maintenance



### OPENCoastS

### Coastal circulation on-demand forecast

on-demand circulation forecast systems for selected coastal areas

- 48h or 72h forecasts
- water levels, vertically averaged velocities and wave parameters
- HPC simulations
- EGI service providers LIP+IFCA

#### Integration in EOSC

- Submit jobs to EGI over DIRAC
- EGI HTC/HPC compute service
- Use EGI Fedcloud to deploy the Opencoasts Services
- Data transfer between cloud and GRID using StoRM + WebDAV
- Use EUDAT for data provenance

Computing





OSC-hub

# Software management



- EGI UMD and CMD software quality assurance (IBERGRID)
  - Validate software for services to be production-ready. Reduce the probability that the new software releases are deployed in production and cause problems to the services.
  - Quality Criteria Definition
    - Software Quality Criteria specification specific for each software
  - Quality Criteria Verification and maintain a verification testbed
    - Nearly full automateded testbed based on Jenkins
  - Staged Rollout
    - Procedure through which newly verified software releases are first deployed and tested by Early Adopter sites

### • EOSC Technical Specification for Software Quality Assurance

 Technical Specifications for all software development processes under EOSC starting from the definition of requirements, coding, release, testing and integration



# Software management

- GitHub: software repositories (ansible roles used by jenkins)
- Jenkins: open source automation server used to provide automatic validation process. Triggered by Request Tracker
- **Request Tracker:** internal track of the full process







- DEEP-Hybrid-DataCloud
- EOSC-synergy
- udocker



# **DEEP-Hybrid-DataCloud**



- Deep Learning / Machine Learning as a Service.
  - Use of specialized hardware:
    - GPUs, Low latency interconnects
  - Platform/Framework deployment: On demand
    - DL stack: Keras, Tensorflow, Theano ...
    - Jupyterhub
    - Data streaming/processing: Apache Spark/Flink
    - Virtualized environments: Openstack (VMs: QEMU/KVM)
  - Docker Orchestration: Mesos/Marathon, (*Kubernetes*)
  - Use of udocker when in HTC and HPC clusters.
  - LIP roles:

MDI

openstac

UDOCKER

**Tensor**Flow

- Work Package leader: SW release, management and support, testbeds
- Task leader: Bare-metal like performance to access GPUs and Infiniband







https://deep-hybrid-datacloud.eu/





### **DEEP-Hybrid-DataCloud**



- DEEP Open Catalogue: https://marketplace.deep-hybrid-datacloud.eu/ •
- Documentation: <a href="http://docs.deep-hybrid-datacloud.eu/en/latest/">http://docs.deep-hybrid-datacloud.eu/en/latest/</a>



Docs » Welcome to DEEP-Hybrid-DataCloud's documentation!

C Edit on GitHub

#### Welcome to DEEP-Hybrid-DataCloud's documentation!

#### **User documentation**

If you are a user (current or potential) you should start here.

- User documentation
- Quickstart Guide

#### **Technical documentation**

If you are searching for technical notes on various notes.

- Technical documentation

  - · Prepare the agent (slave) node
  - Testing Chronos patch for GPU support
  - Testing GPU support in Marathon
  - Running tensorflow docker container
- - Enabling open-id connect authentication

### **DEEP-Hybrid-DataCloud**



https://deepaas.deep-hybrid-datacloud.eu/swagger/index.html?url=https://deepaas.deep-hybrid-datacloud.eu/api/v1/web/deepaas/deep-oc/mods/swagger.json



• DEEP Architecture





# **EOSC-synergy**

 Expand EOSC Capacity and Capabilities by leveraging Investments and Existing Know-How & Resources of National Digital infrastructures, under the coordination of Ibergrid

Into EOSC



Home EOSC Synergy - News & Events Partners Documents



- Integrate/Federate National computing and storage resources: Cloud HTC, HPC
- Integrate National data repositories: towards FAIR data
- Integrate/Federate National Thematic Services: Scientific portals and platforms
- Service Quality Assurance criteria and <u>SQA</u>aaS (as a <u>Service</u>)
- Develop and provide training materials, tutorials and corresponding services and infrastructure.



# **EOSC-synergy**

- <u>Software/Service</u> Quality <u>Assurance</u> <u>as</u> <u>a</u> <u>Service</u>
  - <u>SQAaaS</u>

- Define a set of quality criteria.
- Implement criteria in a <u>C</u>ontinuous
  <u>Integration framework</u> (CI Jenkins pipeline)
- User/Developer can "Compose" the pipeline
- Provide the service and resources to execute the CI pipeline
- IF SQA criteria passes  $\rightarrow$  issue badge  $\rightarrow$  create report
- Can be used in scientific applications and thematic services.
- For software: SQA baseline document:
  - <u>https://github.com/indigo-dc/sqa-baseline</u>
  - <u>http://hdl.handle.net/10261/160086</u>



### udocker

EOSC-hub Hybrid DataCloud



- In DEEP is used to
  - docker compatible containers
  - enable ML in HPC environments
- Added support for
  - newer Linux kernels and distributions
  - OCI images
  - save containers in docker format
  - installation with pypy
  - execution in Linux ARM

### • Improvements in

Computir

- download and handling of manifests
- pathname translations in Fn modes
- handling of mountpoints
- GPU support across distributions
- and many more
- Porting to Python 3 complete

- In EOSC-hub is used to
  - execute containers in HTC
- Integrated in
  - CWL, bioconda, SCAR, etc
  - Several user communities
    - Fusion
    - Molecular dynamics
    - NMR
    - Genome & Life Sciences
    - Physics
    - Engineering
    - Etc
- EXA-LAT: Post-Brexit project with UK Lattice Field Theory
  - LFT algorithm scaling towards exascale

Search or jump to	Pull requests issues Marketplace	Explore 🗳 + - 🚍 -
⊑ indigo-dc / <b>udocker</b>		O Unwatch ▼      29      ★ Star      646 <sup>9</sup> / <sub>2</sub> Fork      72
Code () Issues 32 () Pull requests 2	Actions Projects 0 III Wiki	Security Insights   Settings
A basic user tool to execute simple docker conta	ainers in batch or interactive systems w	thout root privileges
docker containers batch user emulation	proot runc fakechroot indigo grid	i hpc root-privileges chroot docker-containers
deep-hybrid-datacloud eosc-hub Manage topics	<u>s://github.com/in</u>	<u>digo-dc/udocker</u>
6 545 commits  9 5 branches  1	package 🔿 11 releases 🛷 1 e	environment 👪 15 contributors 🎄 Apache-2.0
Branch: master - New pull request	Creat	e new file Upload files Find file Clone or download +
jorge-lip Update README.md		✓ Latest commit зазссе7 13 days ago



Poster at ISC 2019

### Thank you !



### **Events**

- **Digital Infrastructures for Research 2018** 
  - ISCTE Lisbon
  - EOSC-hub, GÉANT, OpenAIRE and PRACE
  - More than 400 participants
- IBERGRID 2018
  - ISCTE Lisbon
  - Co-located with DI4R
- IBERGRID 2019

- CESGA Santiago de Compostela
  - More than 100 participants











### **INCD integrated approach**







### Present

Computing







EuroHPC

LHC and HL-LHC

ATLAS Preliminary

Resource needs

(+15%/year)

2018

(2017 Computing model) — Flat budget model

2020

2022

**Data Repositories** 

Services for Science

**Further EOSC activities** 

2024

2026

2028

# WLCG Tier-2 country comparison



Computir