Portuguese **Distributed Computing** Infrastructure

Centre and Northern Nodes Deployed (Mi6.2)

Execution Report

State:	**FINAL**
Dissemination:	**PUBLIC**
Authors:	**Jorge Gomes**
URL:	**http://www.incd.pt**

Date	Author	Comments
07-05-2019	Jorge Gomes	Initial version
26-06-2019	Jorge Gomes	Technical deployment information
30-07-2019	Jorge Gomes	Final version
11-09-2019	Jorge Gomes	Update

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

Portuguese **Distributed Computing** Infrastructure

Executive Summary

INCD aims to support a wide range of scientific domains and projects with very different requirements. Supporting them across all stages of simulation, processing, analysis and data sharing requires a flexible infrastructure on top of which new added value services and specific virtual research environments can be implemented and delivered.

The activity 6 (deployment) includes: the deployment of the services designed in the activity 5 (development), as well as the installation of equipment to be purchased and/or committed by the INCD users.

This milestone refers to the deployment status of the INCD operations centres in the North (Minho) and centre (Coimbra).

During the second year the services to be deployed at these centres was further detailed. Several factors contributed to make the deployment more difficult than initially foreseen. First, cash flow issues caused by the need to pay purchased equipment prior to request the corresponding funding from FCT. Second, the deployment of the BOB supercomputer in Minho, and the proposal for a Petascale supercomputer in the EuroHPC context, changed the housing requirements and priorities. This contributed to make the housing situation less clear, and made INCD postpone some of the investments planned for the north and centre regions.

In the spring of 2019, the BOB supercomputer was installed at the REN datacenter in Riba de Ave (Minho). A partnership with INCD was established and the deployment of the INCD-Minho operations centre started. The focus of the INCD-Minho operations centre is:

- High Throughput Computing: provided by a batch farm.
- High Performance Computing: provided by a batch farm.
- Cloud Computing: to be provided by an OpenStack IaaS cloud.

Regarding the INCD-Coimbra operations centre, its development and deployment has started. However the purchase of new equipment has been postponed to take place in a fourth project year through a project extension. This centre will be focused on long term data storage, both for data archival, data protection and recovery.

Portuguese **Distributed Computing** Infrastructure

The INCD-Minho deployment

The general architecture of the INCD batch services for High Performance Computing and High Throughput Computing were defined in the milestone Mi6.1. The deployment of the corresponding services for the INCD-Minho operations centre were performed maximising synergies in the context of the cooperation for the installation of the BOB supercomputer.

HPC and HTC services

The BOB supercomputer is composed of 20 racks of the former Stampede supercomputer offered by TACC to Portugal in the context of the cooperation agreements FCT - UT Austin. This equipment has been complemented by storage acquired by FCT. INCD cooperated in this process by performing the installation of the storage solution currently based on a Lustre parallel filesystem. INCD is also managing this storage system which is shared by both the Minho Advanced Computer Centre (MACC) and also INCD.

The INCD-Minho operations centre provides HPC and HTC services that currently exploit 20% of the BOB supercomputer. This is further complemented by equipment purchased by INCD including servers for data storage and servers for metadata equipped with large memory and NVMEs, all incorporated in the shared Lustre installation. Furthermore INCD has purchased additional racks, storage servers, storage gateways and network switching equipment to complement the local installation. A virtualization platform with dedicated hardware was also installed to facilitate the management and provisioning of additional services.

Overall at the time of the writing of this milestone, the INCD-Minho operations centre in Riba de Ave has a computing capacity of 2560 CPU cores interconnected by Infiniband, complemented by 1.5 Petabyte of Lustre data storage. It offers both HPC and HTC processing services that are being exploited by INCD users.

The batch scheduling system installed is Slurm. This deployment, offered to INCD the opportunity to gain experience with the Slurm batch system using a new separate installation. It is foreseen that the experience with Slurm obtained at INCD-Minho will be later applied to migrate the INCD-Lisbon batch services from SoGE to Slurm, thus making the HTC and HPC services at both operations centres homogeneous.

The HPC service supports the execution of parallel processing applications. Multiple parallel environments have been configured in the batch system and made available to the end-users these include MVAPICH, OpenMPI, etc. INCD also purchase the Intel Parallel Studio XE suite for INCD-Minho and is supporting the Intel compilers and MPI libraries.

The CVMFS read-only filesystem was also implemented, enabling transparent sharing of software (applications and libraries) between the INCD-Lisbon and INCD-Minho operations centres. Thus facilitating the execution of applications in both centres and maximising synergies in the applications support.

Cloud Computing service

The deployment of an INCD cloud computing installation in Minho is now being prepared. The purchase of the equipment is foreseen to take place between the third and fourth years. This cloud computing installation will follow the same architecture of the INCD IaaS cloud deployed at the INCD-Lisbon operations centre. It will complement the overall INCD cloud capacity and further improve the resiliency with two redundant cloud regions.

The team supporting INCD-Minho

The INCD operations centre in Minho is backed by two engineers contracted by the INCD association under the 01/SAICT/2016 project. These contracts were initially foreseen for INCD-Coimbra, however they were moved to INCD-Minho to further exploit the opportunities created by the advanced computing activities related to the BOB supercomputer, RNCA and EuroHPC. The team is based at the LIP-Minho facilities in the University of Minho, and is complemented by LIP researchers and engineers from LIP-Minho and LIP-Lisbon.

The INCD-Coimbra deployment

The INCD-Coimbra operations centre is leveraging the LIP-Coimbra facilities in the department of physics of the University of Coimbra. It is using a datacenter at the physics department where INCD is already present with a small installation, composed of four racks with computing and storage equipment already connected to the INCD private L2 network operated by RCTS. The INCD-Coimbra site provides resources that support development and testing of the INCD services.

In addition the INCD-Coimbra deployment lead to a collaboration agreement with the Instituto Superior de Engenharia de Coimbra (ISEC), under which INCD is exploiting computing capacity from the ISEC HPC cluster in Coimbra and collaborates in the management of this computing facility. The

ISEC computing cluster is now integrated in the INCD services portfolio. This partnership decreased the pressure of having newer computing resources at INCD-Coimbra.

The INCD-Coimbra operations centre will be further improved to constitute a storage node integrated in the INCD infrastructure. These activities will leverage the INCD participation in the H2020 project EOSC-Synergy scheduled to start in September 2019, where INCD aims to contribute in the area of FAIR data repositories. The following services are under development in sub-task 5 of the project activity 5.

Further Plans

The execution of the INCD 01/SAICT/2016 project is delayed especially in what concerns the acquisition of equipment and related activities. INCD intends to request an extension to complete the implementation of the infrastructure. The plan for third and fourth years is:

INCD-Minho:

- Implementation of a IaaS cloud computing service;
- Integration of the cloud computing service into the IBERGRID and EGI distributed computing infrastructures.

INCD-Coimbra:

- Implementation of a data protection and data backup on-demand service;
- Implementation of a long-term data storage service;
- Implementation of a public data archive service;

Appendix I – hardware purchased and/or installed

During the second project year, the following hardware was purchased to reinforce the INCD-Minho operations centre.

- Two 19” racks
  - APC NetShelter SX 48U AR 3307
  - PDUs APC 8886
- Two Network management switches
  - 48x gigabit 1000Base-T ports
- Two storage nodes for Lustre with:
  - Two Intel Xeon Silver 4110 having a total of 16 CPU cores
  - 64 GB of RAM per storage node
  - Two SSD with 480GB
  - Dual 25GbE SFP+
  - 24x 8TB SATA3 disks
  - Mellanox Infiniband ConnectX-4
- Three servers for Lustre metadata with:
  - Two Intel Xeon Silver 4110 having a total of 16 CPU cores
  - 192 GB of RAM
  - Two SSD with 240GB
  - 2x NVMe 1TB
  - Dual 10GbE SFP+
  - Mellanox Infiniband ConnectX-4
- Two servers for supporting services (I/O and network gateways) with:
  - Two Intel Xeon Silver 4110 having a total of 16 CPU cores
  - 64 GB of RAM
  - Two SSD with 240GB
  - Dual 10GbE SFP+
  - Mellanox Infiniband ConnectX-4
- One Layer 3 Ethernet managed network switches with:
  - 48x 10GbE SFP+ ports
  - 6x 40/100GbE QSFP+ ports
- Complemented by 12 servers for virtualization hypervisors (offered by CERN):
  - Two Intel E5540 having a total of 8 CPU cores
  - 24 GB of RAM per node
  - Two or Three 250 GB disks

The hardware was installed and tested between June and July 2019. The equipment is now fully operational and integrated in the HPC, HTC, cloud and data services.