Portuguese **Distributed Computing** Infrastructure

Enhanced Implementation of Cloud Services (Mi5.2)

Execution report

State:	**FINAL**
Dissemination:	**PUBLIC**
Authors:	**Catarina Ortigão, Mario David**
URL:	**http://www.incd.pt**

Date	Author	Comments
09-Jan-2019	Catarina Ortigão	Doc initial draft
24-Jul-2019	Jorge Gomes	Endpoints and dashboard.
24-Oct-2019	Mário David	Final version

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

Introduction

INCD will provide services whose features and capacity will be continuously reviewed and improved to address the research growing needs. The evolution will be directed by feedback obtained through the pilots. The development will be performed via the INCD specialized centres.

In the 1st year the focus will be on the Lisbon centre, and in planning the North and Centre region centres which will be deployed in the 2nd and 3nd years.

The services will be mainly built on open source software, standards, and best practices aiming at providing interoperable solutions. The development will follow a path of planning, development, and continuous enhancement. The processing and storage capacity associated to the services will be gradually increased over the three years.

Sub-task 1 - Cloud IaaS

Supporting a wide range of needs requires a flexible platform. OpenStack is an open-source highly modular cloud computing software with major international adoption. It is widely used by the scientific and academic communities to provide IaaS services. The IaaS level will provide virtual computing environments directly controlled by the end-users, and a basis for other high-level services.

Sub-task 2 - Online Storage

The data storage services will provide various types of storage (objects, blocks, and files). These services will be integrated with the IaaS cloud. Block and object storage will be based on Ceph an open-source flexible data storage system with good integration with OpenStack.

For file storage Lustre will be initially used until the Ceph filesystem becomes mature.

Sub-Task 3 - High-level services

Services valuable to research will be deployed on top of the IaaS cloud. These services may be generic (i.e. map reduce, data analysis, databases, PaaS, etc) or specific targeting scientific areas. Access to valuable commercial software such as optimized compilers or computational mathematics will be also considered.

Sub-Task 4 - Advanced Computing

The computing farm and grid service will be redeployed on the IaaS service using newer software. This will allow a more efficient and flexible balance of hardware capacity between services. Access to accelerators (GPGPUs), high performance networks, and servers with large memory will be made available on top of the IaaS service.

Sub-Task 5 - Data Protection

Data backup and long term storage services will be implemented. The offline data backup will take advantage of existing hardware whose capacity will be enlarged.

This report is for the project's milestone **Mi5.2 - Enhanced implementation of cloud services**, and provides evidence that this task/milestone has been done and is complete.

Openstack Cloud Infrastructure Architecture

The typical Openstack IaaS deployment is comprised of four main types of services:

1. Authentication and Authorization service: ***keystone***
2. Infrastructure services:
    a. Computing: ***nova***
    b. Storage: block volume - ***cinder***, object store - ***swift***
    c. Networking: ***neutron***

An additional service to manage the virtual machines images is called ***glance***.

The users interact with Openstack either through a Web Graphical User Interface (WebGUI) called ***Horizon Dashboard***, or through the Openstack Command Line Interface (CLI).

The INCD Cloud Infrastructure is based on Openstack Stein version. The implemented architecture is described next.

1. Two "haproxy + keepalived" nodes; expose the public services, the Horizon Dashboard and the several Openstack API endpoints: keystone, nova, neutron. The two hosts guarantee high availability and failover. The public services are exposed through SSL/HTTPS with X.509 digital certificates emitted by Digicert-Terena.
2. Three controller nodes; run the API services: keystone, nova and neutron, run the Horizon Dashboard. These nodes also run a mariadb (SQL database) with galera for active-active replication as well as a cluster of rabbit message queue service, necessary for the management of the Openstack services. This setup, guarantee high availability and failover if any of the nodes fails.

3. Two network nodes; run network (neutron) agents for several purposes, private networks and subnetworks for the Openstack projects, DHCP and Metadata services, L2 and L3 network agents. The neutron agents make use of keepalived for failover purposes of either node.
4. Two storage nodes: run cinder (storage block volumes), glance (images) and swift (object store) APIs. The haproxy and keepalived run on both nodes providing high availability and failover capability, if any of the nodes fails. The underlying storage facility is based on Ceph distributed service deployed with "replica 3" for storage disk failover.
5. Compute nodes; these are the nodes where the Virtual Machines (VMs) are instantiated. They run the nova-compute service and KVM/QEMU virtual machine service. The Neutron L2 agent also run on these nodes to provide networking to the VMs.

Operation and user interaction

The service is deployed at the following URL: https://stratus.ncg.ingrid.pt/.

For each user or group of users, a new Openstack project is created as well as the private network and sub-network, providing isolation between projects. For more advanced users, an account is created (in keystone) so they can create and manage their infrastructure resources (computing, storage and network). For less advanced users, INCD support does the creation of the necessary resources on their behalf and given them the control after.

Endpoints

The API endpoints to access the Openstack installation are shown in Figure 1.

Service	Service Endpoint
Compute	https://stratus.ncg.ingrid.pt:8774/v2.1/
Identity	https://stratus.ncg.ingrid.pt:5000/v3/
Image	https://stratus-stor.ncg.ingrid.pt:9292
Network	https://stratus.ncg.ingrid.pt:9696
Object Store	https://stratus-stor.ncg.ingrid.pt:8080/swift/v1/AUTH_44f4b4c3447740e9a39cf62a1c161ca1
Placement	https://stratus.ncg.ingrid.pt:8778
Volumev2	https://stratus-stor.ncg.ingrid.pt:8776/v2/44f4b4c3447740e9a39cf62a1c161ca1
Volumev3	https://stratus-stor.ncg.ingrid.pt:8776/v3/44f4b4c3447740e9a39cf62a1c161ca1

Figura 1 INCD Openstack installation API endpoints

Usage

At the time of writing there are 34 projects running in the INCD Openstack IaaS cloud service. The figure 2 shows the usage and quota for a given project / tenant in the INCD Openstack installation.

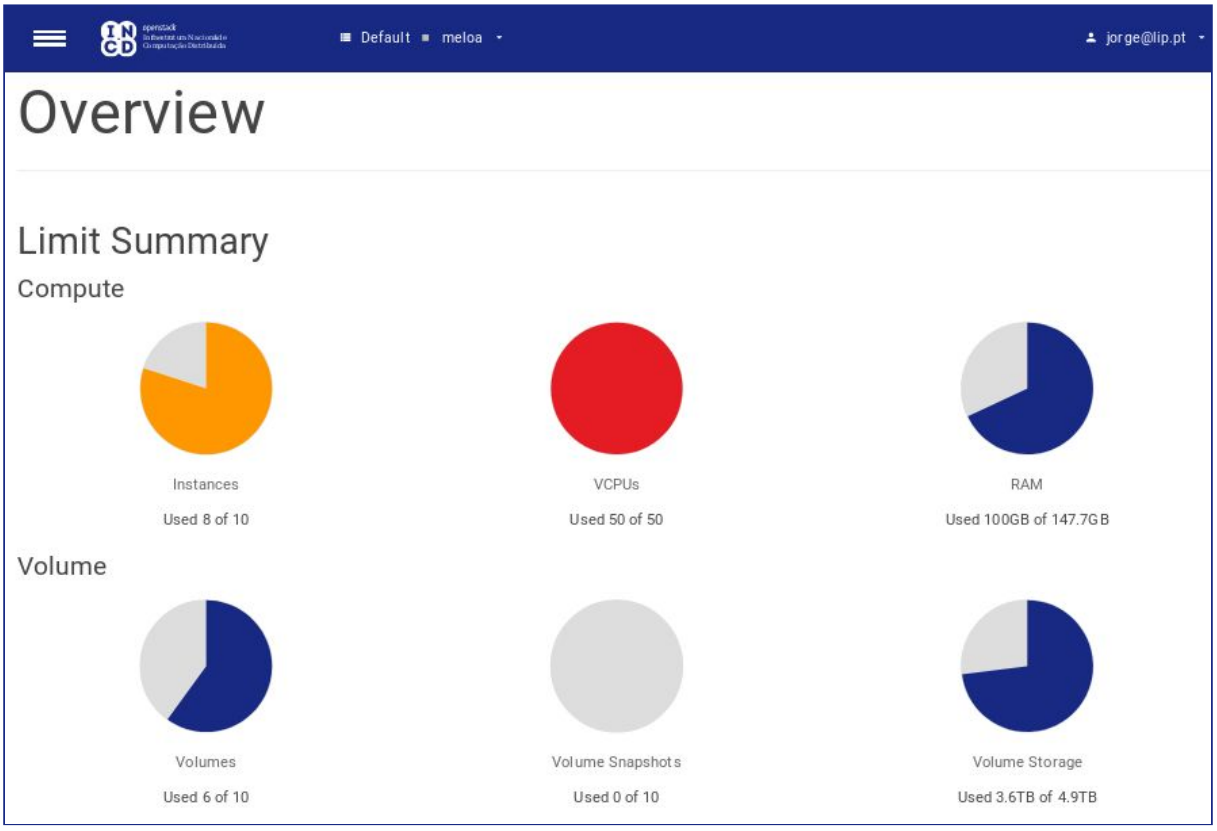

Figura 2 View of usage and quota limits for a given project