Contribution ID: **3** Type: **Presentation**

From bare metal to the cloud: the endless journey of a scientific data center

Wednesday, 8 May 2019 11:10 (10 minutes)

Nowadays it is accepted that Cloud computing is a disruptive paradigm that has been rapidly adopted by industry and government sectors due to its unique features such as reduced costs, elastic scalability, self-service provisioning, etc. This cloud advent is not something only restricted to the IT industry, and also research and education have embraced it. However, the road to the adoption of the Cloud computing model is not something that sprouted suddenly and it is worth looking at how cloud datacenters have adopted and migrated into a Cloud model [1].

Starting on 2005, virtualization (when paravirtualization and hardware assisted virtualization blasted-off) has become a commonly adopted solution within modern data centers due to its widely discussed advantages over the usage of traditional machines [2]. Virtualization is one of the key technologies that paved the road for the Cloud computing advent, as virtual machines (VMs) provided the abstraction needed for Cloud resources to be provisioned. From that initial seed, nowadays Cloud Computing providers are able to provide not only virtual machines, but also physical machines, virtual networks, etc.

Taking into account this context, the natural evolution for a virtualized datacenter is to move towards a private cloud scenario, where the virtualized resources are not managed anymore manually or with in-house developed tools by the infrastructure administrators, but rather this management is leveraged to a Cloud Management Framework (CMF) such as OpenNebula, OpenStack or CloudStack. Taking this step forward means that the datacenter has evolved into a private cloud deployment, where system administrators, even if they are aware of the underlying infrastructure, can now operate its infrastructure as if they were managing their resources in any cloud provider. Moreover, managing a datacenter in a cloud-like mode opens the door to a new world of possibilities, as it makes possible to offer the infrastructure resources to a wider public.

In this session we will present how the scientific datacenter at the Instituto de Física de Cantabria, a research center in Spain has evolved from a virtualized infrastructure into the Cloud model. We will discuss how our datacenter has evolved and the challenges and opportunities that we have faced.

Type of abstract

Presentation

References

1: Kate Keahey, Manish Parashar, Enabling on-demand science via cloud computing , 20142: Ranadive, Adit; Kesavan, Mukil; Gavrilovska, Ada; Schwan, Karsten, Performance implications of virtualizing multicore cluster machines, 2008

Primary author: Dr LOPEZ GARCIA, Alvaro (CSIC)

Co-authors: AGUILAR, Fernando (CSIC); CABRILLO, Iban (CSIC); NUNEZ-VEGA, Miguel (CSIC); ORVIZ, Pablo (CSIC); PALACIO, Palacio (IFCA)

Presenter: Dr LOPEZ GARCIA, Alvaro (CSIC)

Session Classification: Cloud: State of the Federation and Future Opportunities