

Lifting scientific computing to the cloud

Erich Birngruber, Petar Forai, Petar Jager, Ümit Seren

Vienna BioCenter

The Vienna BioCenter (VBC) is building its next scientific computing infrastructure on top of OpenStack. We want to share our experience of deploying a scientific private cloud on a heterogeneous set of hardware with the Austrian HPC community. We will discuss the deployment and configuration of different CPU, Memory and GPU accelerators configurations for various application paradigms. Since the traditional batch scheduler workload is not going away, we will give an insight why we chose to virtualise compute nodes and how we achieved close to bare metal performance on these. We want to start a discussion about issues not addressed yet, share the pitfalls we encountered on the way and we welcome exchange with others in the community going the same route.

The Vienna BioCenter hosts 3 institutes involved in life sciences research (GMI, IMBA, IMP). As the 3 existing HPC installations on campus are reaching their end-of-life, a unique opportunity presented itself to put efforts together and replace these installations with a single unified concept. Additionally the High Energy Physics Institute of the Academy of Sciences (HEPHY) joined this effort, pooling personnel and hardware resources.

At the VBC a diverse range of instruments and appliances need to be integrated into scientific workflows. While some workflows remain based on traditional work load managers, this is not always a good fit for modern ways of software deployment such as containers or Big Data Frameworks as Apache Spark and Hadoop. However, as we expect a shift in these requirements towards a more cloud-native methodology, we looked for a solution that provides as much flexibility in deployment as possible. OpenStack has a proven track record for public and private cloud offerings.

The Wellcome Sanger Institute [1], one of the leading institutes in life sciences research is successfully running OpenStack installation for 2 years, offering private cloud services for scientific customers. CERN is running one of the largest OpenStack deployments [2], hosting a significant part of the scientific workload. The OpenStack deployments at these institutes show that private clouds fill a demand of the scientific community that needs to be addressed.

Designed as a cloud framework, OpenStack consists of decoupled services that make up a fully fledged cloud deployment similar to commercial public cloud offerings. One of its core features is the deployment of virtual and bare metal resources. These capabilities, not found in typical HPC cluster deployment tools, allow for a flexible configuration and use of resources. Furthermore shifting the emphasis between batch scheduler workloads and cloud native software deployments like containers and virtual machines can only be achieved with a system, that allows dynamic re-provisioning. Particularly for small institutes with limited financial resources it is important to make efficient use of datacenter resources, therefore it is key to achieve best possible integration with these existing resources. OpenStack with its decoupled service model allows deep integration with core datacenter services like storage and networking through vendor provided drivers.

References

- [1] Wellcome Sanger Institute, 2018 HPCWire award for "Best Use of HPC in the Cloud" at SC '18
- [2] Bell, T and Bompastor, B and Bukowiec, S and Castro Leon, J et al. (2015). Scaling the CERN OpenStack cloud. Journal of Physics: Conference Series. 664. 022003. doi:10.1088/1742-6596/664/2/022003.
- [3] Erich Birngruber, and mit Seren, AHPC17: Past Present and Future of scientific computing in life sciences