

Analysis of Elastic Cloud Solutions in an HPC Environment

- ▶ Author: Johannes Coym
- ▶ Type: Master's Thesis
- ▶ Date: 2021-10-23
- ▶ Reviewers: Jun.-Prof. Dr. Michael Kuhn, Prof. Dr. Thomas Ludwig
- ▶ Supervisors: Jannek Squar, Jun.-Prof. Dr. Michael Kuhn
- ▶ Download: PDF

Cloud Services, like AWS, Azure, and Google Cloud, are a growing market, and all of them also started providing their cloud services for HPC use cases in the last years. Amazon even developed their own fabric adapter for AWS, claiming greatly improved load distribution compared to existing solutions. This thesis aims to evaluate the current performance and profitability of cloud services for HPC applications. Specific focus will be on the rentability of running several specific or even all applications in the cloud. For this cloud side, AWS, Azure and Google Cloud will be used for profitability analysis. These cloud providers will have to compare to an on-premise HPC cluster for different job configurations. On the side of the on-premise HPC cluster, the cluster usage of a whole year of one cluster will be analysed to gain the required data. Additionally, the costs of running those variants will be compared on a typical lifespan of an HPC cluster with all of its acquisition costs and running costs. These cost factors and requirements will then be taken into a cost function to assist a cluster owner's decision when HPC cloud systems provide a better value than running an independent cluster. The choice between an on-premise HPC system and the cloud will also not just be looked at as an explicit or, as there is also the possibility of owning a smaller cluster and outsourcing some parts to the cloud. The cost function for the comparison against the cloud can then provide a way to outsource specific jobs to the cloud, reducing the total cost of the cluster and potentially even optimising the remaining jobs for the job scheduler.