



## HPC FOR MACHINE LEARNING: CARME

1 *Simplified scheme of the most important system components and their connections*

Machine learning has an increasingly higher priority in both scientific and industrial enterprises. This is evident from the investment in new, above all, GPU-based hardware – from simple desktop computers to high performance computing clusters. Computing clusters are used in Data Analysis (DA) and highly complex Machine Learning (ML) systems to process and simulate very large amounts of data – to include even the human brain.

Machine learning in HPC clusters presents certain challenges. The procurement of the individual hardware components is the least of these challenges. The biggest questions arise subsequent to that acquisition:

- How to manage existing resources?
- How to make an application scalable to several GPUs?
- How to solve the challenge of data storage and continuous upload to the program?
- How to train users to effectively use the hardware?

The answers to these questions begin with our open-source software stack Carme. The basic concept is to combine the world of machine learning and data analysis with the world of HPC systems. We achieve this using established ML and DA tools with HPC back ends. Specifically, we use a variety of HPC and ML technologies. Some of these technologies are developed in this department, for example, the highly reliable parallel file system BeeGFS for fast data links.

Carme combines the worlds of machine learning and HPC clusters. ML is a steady and fast growing field of technology. This new agility challenges data centers to provide very different applications for single users. It is not enough to have one user interface for the user; rather there must also be a guarantee of a seamless integration of this surface in existing and emerging clusters. To make clusters attractive to ML and DA users, an intuitive software environment must be provided to the clusters. Interactive management of the cluster is essential in the development of ML applications. Users must have the opportunity to use tools they are familiar with on a complex HPC cluster, making it easier for them to migrate to and use the cluster.

