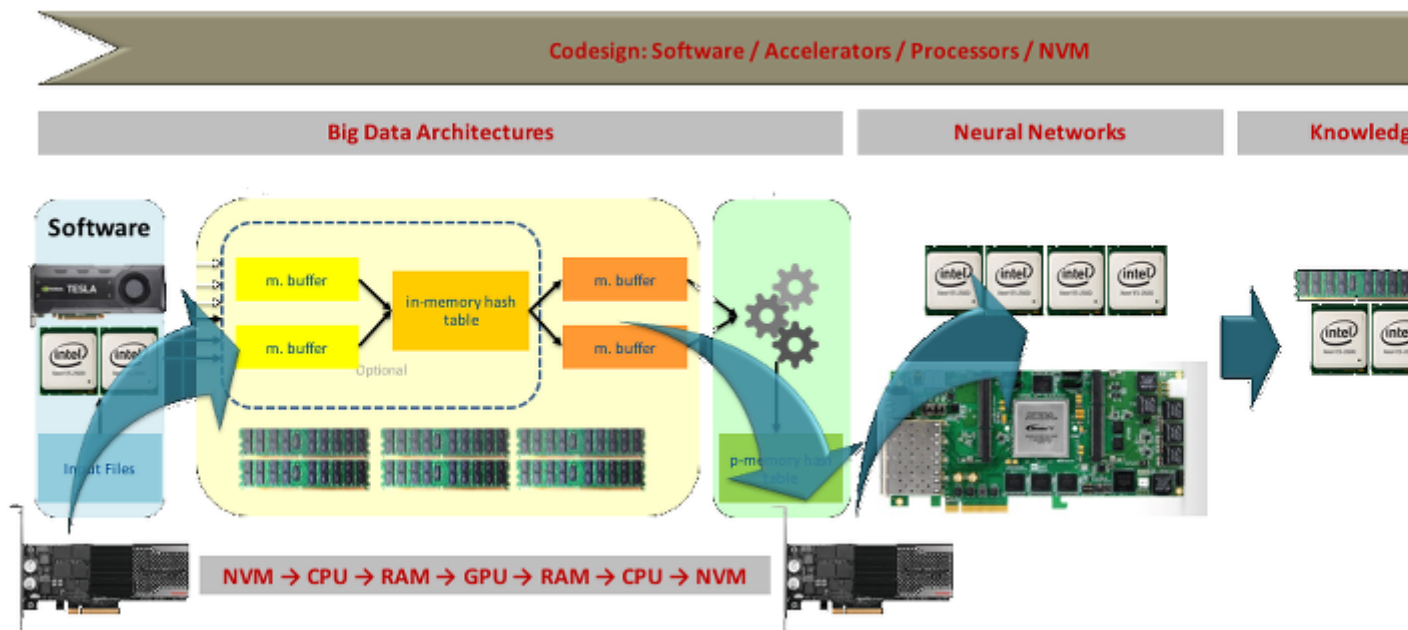# Data-Centric Architectures

This research line aims to develop new data-centric architectures that leverage emerging technologies (accelerators, NVMe) to accelerate workloads, including the development of new interfaces to access the devices as well as new programming paradigms (active storage, KV stores).

Summary

Supercomputers are critical infrastructures to address Grand challenges in the field of Bioinformatics, Physics, and Earth Sciences. A common belief in the supercomputing field is that the future bottleneck for many data intensive scientific advances will be on the technological approach used currently used to process data. This situation is clearly visible observing at the sustained growth rate of the GenBank genome database from NCBI1, or the experiences in the Atlas Project in the process to discover the Higgs boson at the Large Hadron Collider (LHC). In both cases, the scientific grand challenge resembles more a data mining problem than a classical CPU-intensive supercomputing project. But they need, in both cases, the ultra-high capacity of supercomputers to address the problems that each project aims.

The response to this challenge passes by the combination of high-speed networks and high-performance and high-capacity memories. Over the last years an intersection of different technologies is gaining momentum performance-wise, developing a research field known as Active Storage Technologies: RDMA-enabled network technologies and persistent memories interconnected to CPUs to provide storage capabilities attached to compute power. The goal of combining these technologies is to provide low latency and high

bandwidth all-to-all in a network topology, and therefore fast access to all data in a distributed dataset. High-speed networks require lightweight protocol stacks and CPU offloading to move data between nodes at high speeds what is achieved using RDMA-enabled networks such as Infiniband or iWARP. At the same time, to achieve high bandwidth to store data, fast memories are used instead of slow rotational disks. Such memory can be accessed in different ways: through conventional disk interfaces, what is the case of Solid State Disks (SSD); through conventional PCIe buses, what is the case of PCIe Flash boards; or directly through the memory buses in the processors, what is the case of the future Phase Change Memories (PCM) or any generic Storage Class Memory (SCM) technology. As a result, the set of technologies used in next generation supercomputers is extremely complex and has no precedent in the past.

Objectives

- Explore novel architectures for workload acceleration in scale-up solutions (appliances)
- Explore the integration of different accelerator technologies (FPGA, GPU) in compute nodes
- Explore the use of NVM technologies for memory extension
- Develop new interfaces for NVM devices

Barcelona Supercomputing Center - Centro Nacional de Supercomputación