CEA: connecting co-design capabilities with the operation of large HPC and data processing facilities – a key know-how for the future of big data and extreme computing

Edouard Audit, Christophe Calvin, Jean Gonnord, Jacques-Charles Lafoucrière, Jean-Philippe Nominé

Since many years CEA has developed unique know-how for high performance numerical simulations and data management and treatment, encompassing all aspects of the value chain of simulation-based and data-intensive-based science and engineering:

- HPC technology development
- operation of large computing centres
- development and exploitation of large numerical simulation applications

HPC technology development

CEA has active R&D in all HPC topics from hardware components, system architecture, middleware, file systems and storage, to numerical algorithms, meshing and visualization. CEA has been involved in the co-design of high-end supercomputers for more than 10 years, with industrial partners such as ATOS/Bull and Intel, contributing to the successful bullx series[1], then to the recently announced SEQUANA[2].

Balanced system design between compute storage and I/O has always been a primary concern at CEA, with supercomputers having outstanding I/O capabilities such as a global parallel filesystem (Lustre) bandwidth of 500 GB/s for 1.25 Pflop/s in TERA100 in 2010 [3].

In addition to continuing a long history of LUSTRE deployment and development, we have more recently focused a significant part of our HPC systems R&D effort on virtualisation and storage, considered two key aspects of the data- and compute-capable supercomputers of the future.

Virtualization is highly used in the Cloud world. Since more than one year CEA has started to use Virtual Cluster of Virtual Machines on its production clusters. To do this and to simplify users' life, CEA has developed PCOCC, an open source tool to allocate and start virtual clusters of Linux machines. Next step is to support Windows clusters.

Object storage in general and key/value store are promising technologies for exascale. CEA is investigating how these new paradigms could be used for our I/O paths, in our storage software services, and how we could deploy related solutions in our computing centres. Participation in EU-funded projects such as Horizon 2020 SAGE and BigStorage supports and leverages these efforts with many partners in Europe[4]. CEA is also a founding member of ETP4HPC[5] and a strong supporter and player of EU extreme computing and data R&D in general.

Operation of large computing centres

CEA operates world-class supercomputers and offers related HPC expertise and services to research, industry and defence in its supercomputing complex near Paris[6] (more than 6 petaflop/s of computing power as of mid-2016, and ca. 50 petabytes of storage, with TERA1000 supercomputer for defence - 2.5 Pflop/s, CCRT/COBALT 1.5 Pflop/s for industry and GENCI's CURIE - 2 Pflop/s for research).

The proximity with such a wide and diverse portfolio of applications and their users, and the management of more than 100 TB of daily data production, make our computing centres (located on a single site) a unique place to observe the evolution of high-end simulation requirements and related data production and usage trends, to relate them to technology and system requirements and to steer virtuous relationships with technology suppliers.

Development and exploitation of large numerical simulation applications

Some quick highlights can be given on applications that are representative or forerunners of future compute and data intensive simulation, through a short selection of applications CEA is involved in.

- High Throughput Computing (HTC) is used in many different domains at CEA, such as genomics or High Energy Physics (HEP). CEA is part of WLCG[7] (Worlwide LHC Computing Grid) for the data treatment and analysis of LHC experiments. In life sciences, a mission of CEA Genomic Institute is to ensure, for the

[1] http://www.bull.com/bullx-supercomputers
[2] http://www.hpcwire.com/off-the-wire/atos-reveals-bull-sequana/
[3] Chapter 4 in Contemporary High Performance Computing: From Petascale toward Exascale - Jeffrey S. Vetter ed., 2013
[4] http://www.sage-project.eu/home.html , http://bigstorage-project.eu/
[5] www.etp4hpc.eu
[6] http://www-hpc.cea.fr/en/complexe/computing-ressources.htm
[7] http://wlcg.web.cern.ch/

benefit of the scientific community, the mass production and the processing and analysis of genomic data. To do so, a data-centric model has been setup starting from high performance sequencer with a first level of local data processing resources (300 cores and 400 TBytes) and a second level centralised in a CEA national data centre (TGCC/CCRT[8]) for global data storage and analysis (3K cores and 5 PBytes). This approach has many virtues, structuring the dialogue between users of the joint IT/HPC facility and between users and operators of the resources.

- A recent project revisited existing seismic datasets and proved that the re-organisation of data processing and subsequent 'old' data re-interpretation made possible by this approach could lead to significant scientific added value[9]. Revisiting 10 years of seismic signals, analysed in a few hours with improved data processing, it was possible to detect earthquakes in large volumes of geophysical data, so as to spot events that were previously unnoticed – by comparing the continuous flow of recordings with past earthquakes already identified. This experiment has revealed, on a 10 year timescale, the seismic activity of several volcanoes. The data processing workflow was redesigned to handle and process millions of files with a reasonable time-to-solution– which was more of a challenge than the mere 20 TB size of the dataset. This workflow has then been validated and tested on 32 000 cores of CURIE supercomputer. The technical innovation lie in distributed data improved management, together with a hierarchical dispatching and scheduling of processing tasks. This application involved two CEA divisions – HPC and Geosciences ones inside DataScale project[10].
- We can illustrate the challenges of coupling the management of huge data sets with high computing power to achieve valuable analysis with two examples.
    - CEA is part of French IPSL (Institut Pierre-Simon Laplace[11]) federation, one of the active participants in the WCRP (IPCC World Climate Research Program) through the CMIP exercises (Coupled Model Intercomparison Project[12]). Current CMIP6 involves more complex climate models with finer spatial resolution and larger ensembles. The expected increase of amount of generated data compared with the previous exercise (CMIP5) is 20- to 50-fold, yielding up to 90 PBytes of data and 215 million files. A generic library to optimize data storage and on-line processing has been developed: XIOS[13], using client/server modes with I/O delegation, parallel, asynchronous I/O, external configuration files, in-situ processing features... Initially developed for climate models, XIOS has been successfully employed for plasma fusion simulations as well. The challenge is then not only to generate these data from climate simulations or even to store them, but the different inter-comparison of models from all over the world, for which new smart workflows and dataflows must be imagined and designed.
    - The second illustration deals with life sciences and health. A French collaborative project has the main objective to provide secure solutions for therapeutic innovation using HPC technologies for the biomedical community, in order to compare numerical simulation and 3D modelling of pre-clinical models, such as brains, for the development of new drugs. The proposed solutions allow different experts to seamlessly and intuitively navigate in 3D images coming for observations and numerical simulations with access via thin clients (3D tablets without glasses). The proof of concept targets the process for the development of molecules against identified new therapeutic targets in the Alzheimer disease.

We thus believe large technological research organisations, running large computing/data centres, such as CEA have a key part to play to foster the evolution of extreme computing and data processing. Connecting co-design capabilities with the operation of large facilities, hosting a diversity of simulation-based and data-intensive-based science and engineering applications, is a key asset and driver.

[8] https://www.france-genomique.org/spip/spip.php?article39
[9] http://www-hpc.cea.fr/en/news2015.htm
[10] DataScale project was supported by a Digital Economy Development plan of the French government http://www.teratec.eu/gb/activites/projetsR_D_Datascale.html
[11] http://www.ipsl.fr/
[12] http://www.wcrp-climate.org/wgcm-cmip/about-cmip
[13] http://forge.ipsl.jussieu.fr/ioserver