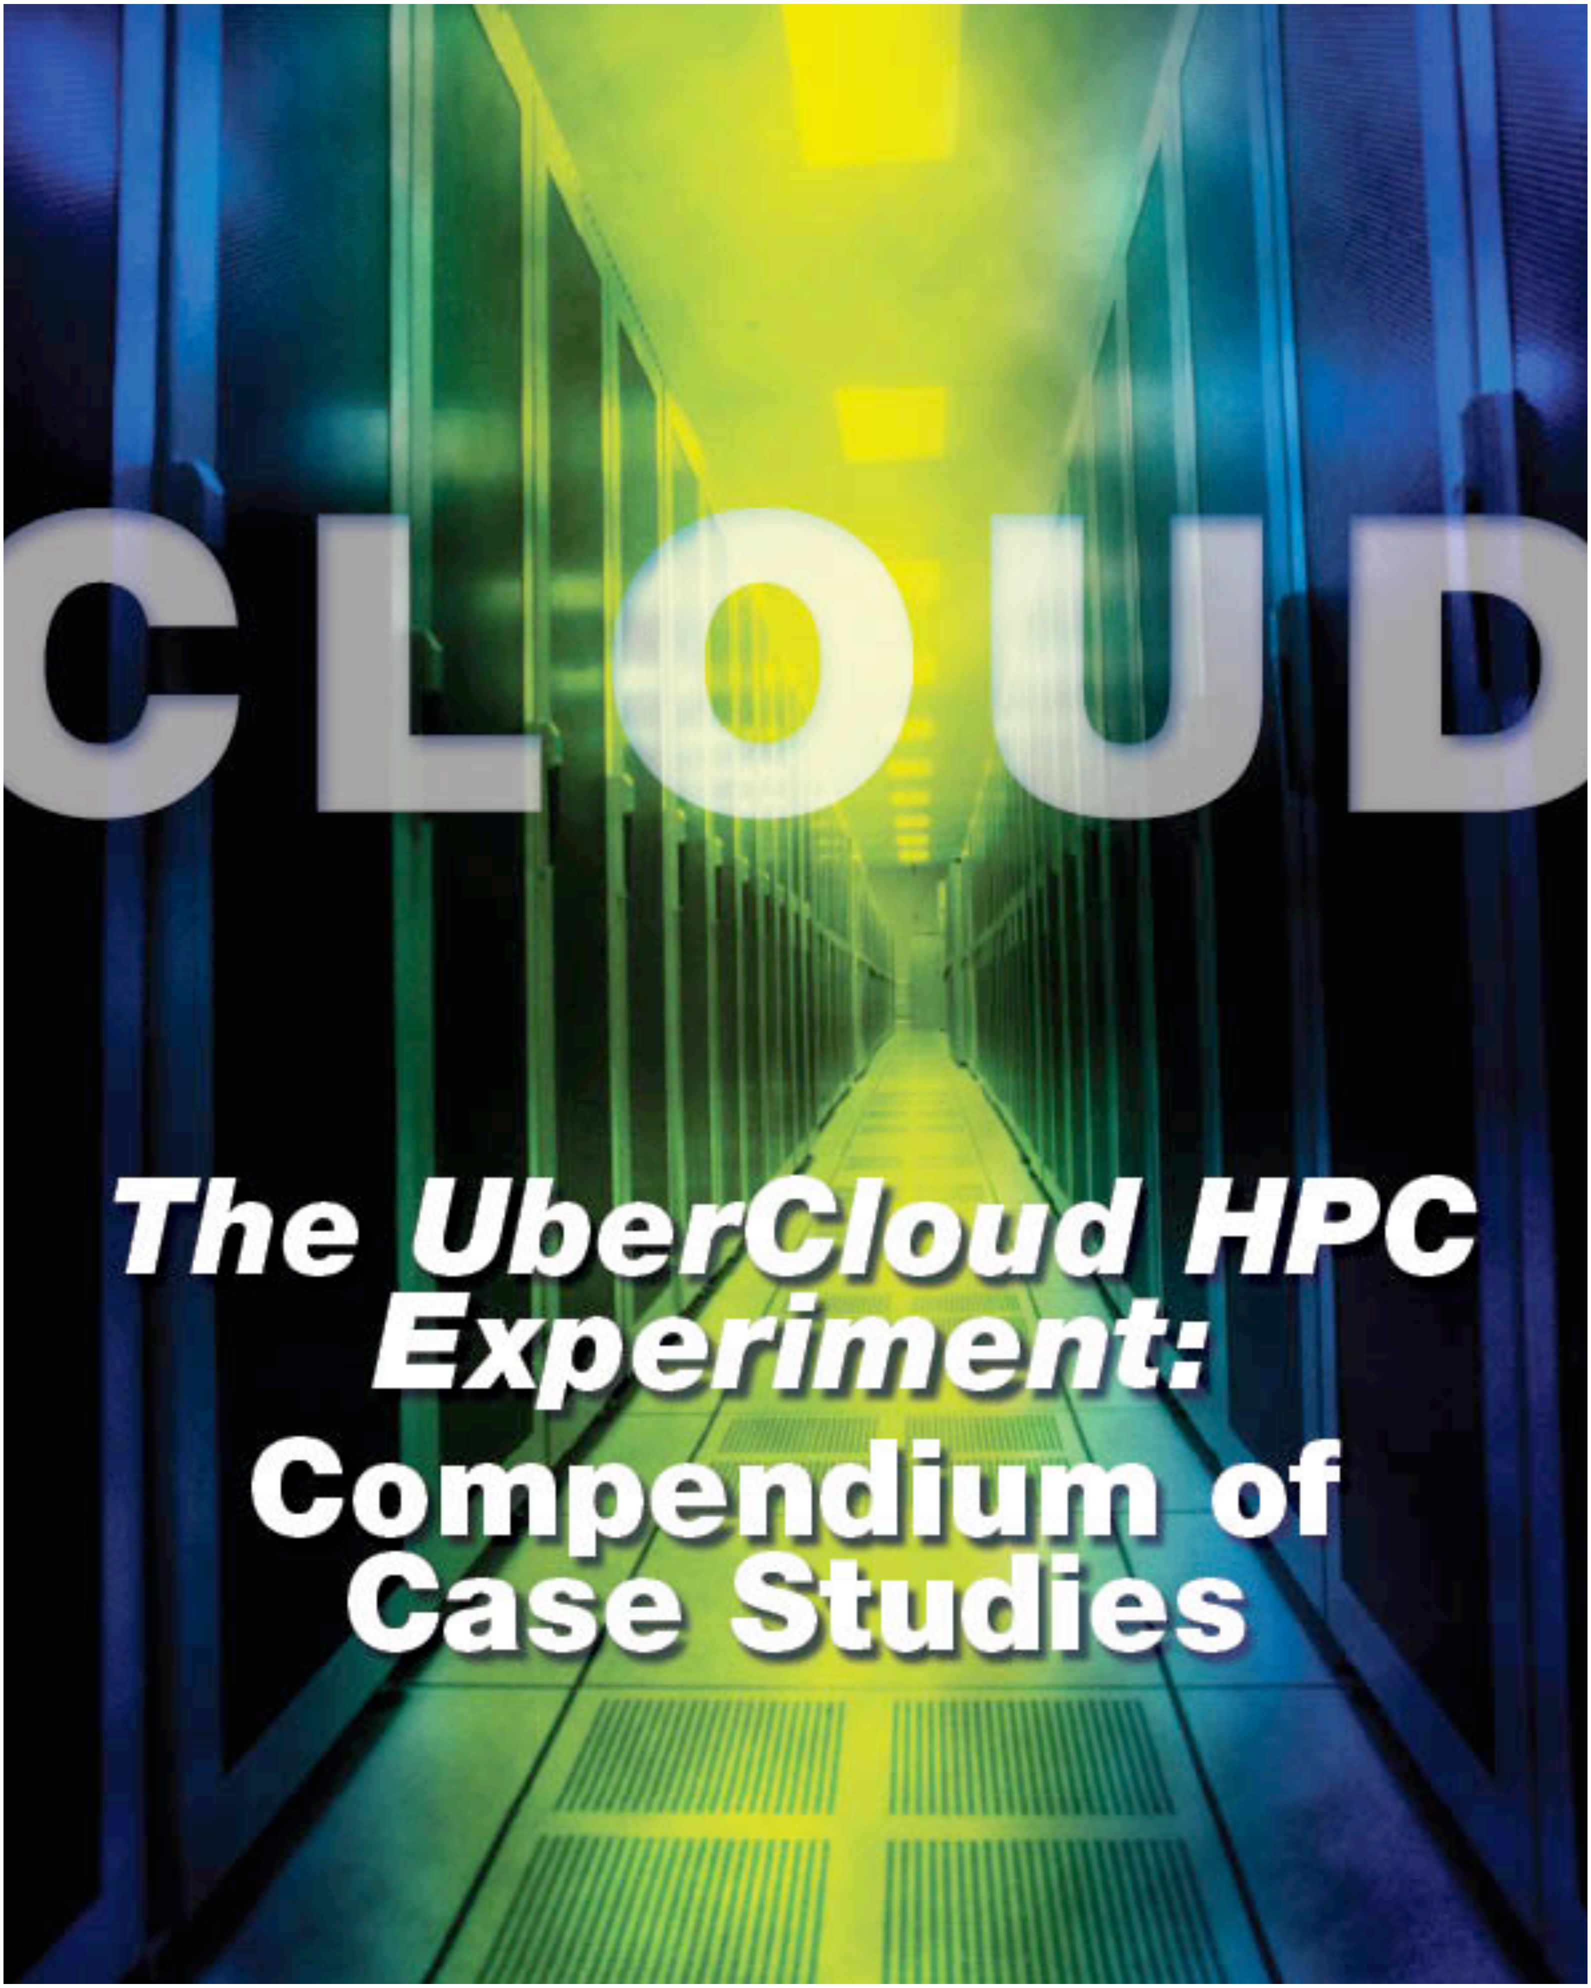
CLOUD
The UberCloud HPC Experiment:
Compendium of Case Studies

Digital manufacturing technology and convenient access to High Performance Computing (HPC) in industry R&D are essential to increase the quality of our products and the competitiveness of our companies. Progress can only be achieved by educating our engineers, especially those in the "missing middle," and making HPC easier to access and use for everyone who can benefit from this advanced technology.

The UberCloud HPC Experiment actively promotes the wider adoption of digital manufacturing technology. It is an example of a grass roots effort to foster collaboration among engineers, HPC experts, and service providers to address challenges at scale. The UberCloud HPC Experiment started in mid-2012 with the aim of exploring the end-to-end process employed by digital manufacturing engineers to access and use remote computing resources in HPC centers and in the cloud.

In the meantime, the UberCloud HPC Experiment has achieved the participation of 500 organizations and individuals from 48 countries. Over 80 teams have been involved so far. Each team consists of an industry end-user and a software provider; the organizers match them with a well-suited resource provider and an HPC expert. Together, the team members work on the end-user's application – defining the requirements, implementing the application on the remote HPC system, running and monitoring the job, getting the results back to the end-user, and writing a case study.

Intel decided to sponsor a Compendium of 25 case studies, including the one you are reading, to raise awareness in the digital manufacturing community about the benefits and best practices of using remote HPC capabilities. This document is an invaluable resource for engineers, managers and executives who believe in the strategic importance of this technology for their organizations. You can download it at:
http://tci.taborcommunications.com/UberCloud_HPC_Experiment

Very special thanks to Wolfgang Gentzsch and Burak Yenier for making the UberCloud HPC Experiment possible.

This HPC UberCloud Compendium of Case Studies has been sponsored by Intel and produced in conjunction with Tabor Communications Custom Publishing, which includes HPCwire, HPC in the Cloud, and Digital Manufacturing Report.

If you are interested in participating in this experiment, either actively as a team member or passively as an observer, please register at http://www.hpcexperiment.com

Heat Transfer Use Case

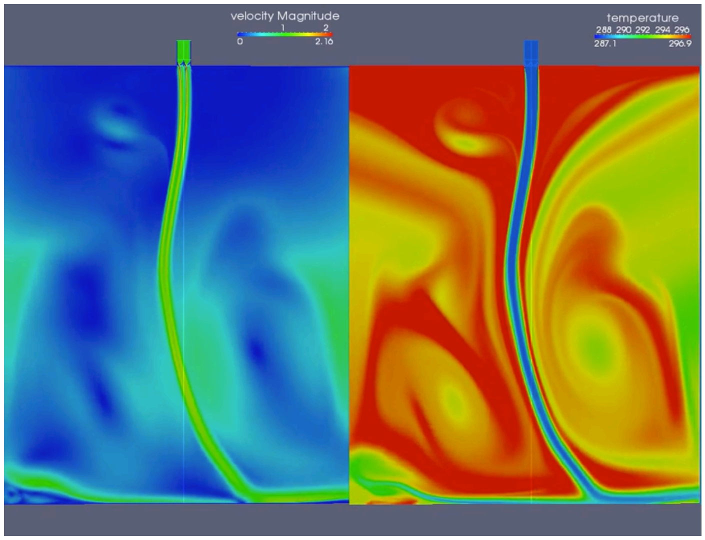

"Concerning the ease of using cloud computing resources, we concluded that this working methodology is very friendly and easy to use through the CloudBroker Platform."

MEET THE TEAM

- Lluís M. Biscarri, Biscarri Consultoria SL, Director
- Pierre Lafortune, Biscarri Consultoria SL, CAE Expert
- Wibke Sudholt, CloudBroker GmbH, CTO and Managing Partner
- Nicola Fantini, CloudBroker GmbH, CEO and Managing Partner
- Members of the CloudBroker Team for development and support
- Joël Cugnoni, researcher and developer of CAELinux
- Peter Råback, CSC — IT Center for Science, Development Manager

Organizations Involved

- *Biscarri Consultoria SL* (BCSL, http://www.biscarri.cat), based in Barcelona, is a SME engineering and consulting company specialized in CAE technology that offers simulation services and know-how transfer to industry. BCSL is focussed on the use of open-source computational mechanics software and its application to multi-physics engineering problems of the industry. The use of HPC and cloud computing hardware resources is one of BCSL main interests as well.
- *CSC — IT Center for Science Ltd.* (http://www.csc.fi) is administered by the Finish Ministry of Education and Culture. CSC provides IT support and resources for academia, research institutes and companies.

- *CAELinux* (http://www.caelinux.com) is an open source project, a ready to use Linux distribution for CAE and scientific computing. The main goal of CAELinux is to promote the use of state of the art open source software in research and engineering. The current version of CAELinux is based on Ubuntu LTS 64 bit and includes the most popular open source CAE applications such as OpenFOAM, Elmer FEM, Code-Aster, Code-Saturne, Calculix, Salome, Gmsh, Paraview and many more. CAELinux is available both as an installable LiveDVD image and as a virtual machine image on Amazon EC2.
- *CloudBroker GmbH* (http://www.cloudbroker.com) is a spin-off company of the ETH Zurich located in Zurich, Switzerland. It offers scientific and technical applications as a service in the cloud, for usage in fields such as biology, chemistry, health and engineering. Its flagship product, the CloudBroker Platform, delivers on-demand web and API access to application software on top of compute and storage resources in public or private clouds such as Amazon Web Services.

USE CASE

Background

In many engineering problems fluid dynamics is coupled with heat transfer and many other multiphysics scenarios. The simulation of such problems in real cases produces large numerical models to be solved, so that big computational power is required in order for simulation cycles to be affordable. For SME industrial companies in particular it is hard to implement this kind of technology in-house, because of its investment cost and the IT specialization needed.

There is great interest in making these technologies available to SME companies, in terms of easy-to-use HPC platforms that can be used on demand. Biscarri Consultoria SL, is committed to disseminate parallel open source simulation tools and HPC resources in the cloud.

CloudBroker is offering its platform for various multiphysics, fluid dynamics, and other engineering applications, as well as life science for small, medium and large corporations along with related services. The CloudBroker Platform is also offered as a licensed in-house solution.

Current State

Biscarri Consultoria SL is exploring the capabilities of cloud computing resources for performing highly coupled computational mechanics simulations, as an alternative to the acquisition of new computing servers to increase the computing power available.

For a small company such as BCSL, the strategy of using cloud computing resources to cover HPC needs has the benefit of not needing an IT expert to maintain in-house parallel servers thus concentrating on our efforts in our main field of competence.

To solve the needs of the end user, the following hardware and software resources existing on the provider side were employed by the team:

- Elmer (http://www.csc.fi/english/pages/elmer), an open source multi-physical simulation software mainly developed by the CSC — IT Center for Science
- CAELinux (http://www.caelinux.com), a CAE Linux distribution including the Elmer software as well as a CAELinux virtual machine image at the AWS Cloud
- CloudBroker Platform (public version under https://platform.cloudbroker.com), CloudBroker's web-based application store offering scientific and technical Software as a Service (SaaS) on top of Infrastructure as a Service (IaaS) cloud resources, already interfaced to AWS and other clouds
- Amazon Web Services (AWS, http://aws.amazon.com), in particular Amazon's IaaS cloud offerings EC2 (Elastic Compute Cloud) for compute and S3 (Simple Storage Service) for storage resources

Experiment Procedure

Technical Setup

The technical setup for the HPC Experiment was performed in several steps. These followed the principle to start with the simplest possible solution and then to grow it to fulfil more complex requirements in an agile fashion. If possible, each step was first tested and iteratively improved before the next step was taken. The main steps were:

1. All team members were given access to the public CloudBroker Platform via their own account under a shared organization created specifically for the HPC Experiment. A new AWS account was opened by CloudBroker, the AWS credit loaded onto it, and the account registered in the CloudBroker Platform exclusively for the experiment team.
2. Elmer software on the existing CAELinux AWS machine image was made available in the CloudBroker Platform for serial runs and tested with minimal test cases by CloudBroker and Joël Cugnoni. The setup was then extended to allow parallel runs using NFS and MPI.
3. Via Skype calls, screen sharing, chatting, email and contributions on Basecamp, the team members exchanged knowledge on how to work with Elmer on the CloudBroker Platform. The CloudBroker Team gave further support for its platform throughout HPC Experiment Round 2. CloudBroker and BCSL performed corresponding validation case runs to test the functionality.

4. The original CAELinux image was only available for normal, non-HPC AWS virtual machine instance types. Therefore, Joël Cugnoni provided Elmer 6.2 as optimized and non-optimized binaries for Cluster Compute instances. Also, the CloudBroker Team deployed these on the CloudBroker Platform for the AWS HPC instance types with 10GBit Ethernet network backbone, called Cluster Compute instances.
5. BCSL created a medium benchmark case, and performed scalability and performance runs with different numbers of cores and nodes of the Amazon Cluster Compute Quadruple and Eight Extra Large instance types and different I/O settings. The results were logged, analyzed and discussed within the team.
6. The CloudBroker Platform setup was improved as needed. This included, for example, a better display of the number of cores in the web UI, and the addition of artificial AWS instance types with fewer cores, as well as the ability to change the shared disk space.
7. BCSL tried to run a bigger benchmark case on the AWS instance type configuration that turned out to be preferable from the scalability runs – that is, single AWS Cluster Compute Eight Extra Large instances.

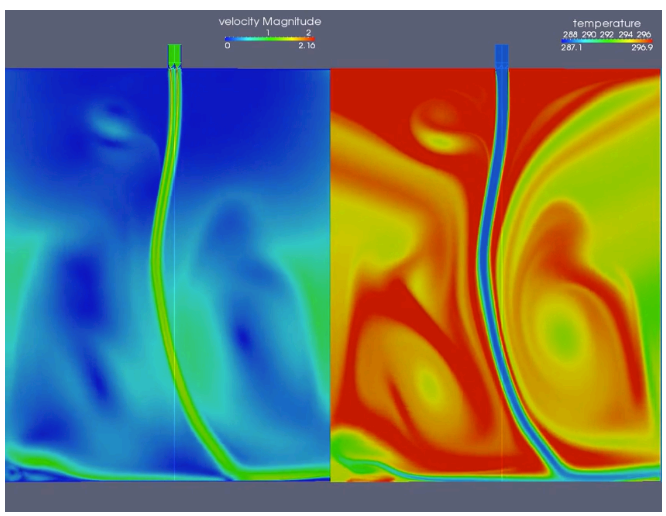

Fig. 1 - This figure shows the model employed in the scalability benchmark. The image on the right shows the temperature field, while the left image shows the velocity field at a certain time of the transient simulation.

Validation Case

First a validation case was defined to test the whole simulation procedure. This case was intentionally simple, but had the same characteristics as the more complex problems that were used for the rest of the experiment. It was an idealized 2D room with a cold air inlet on the roof (T = 23ºC, V = 1m/s), a warm section on the floor (T = 30ºC, V

= 0.01m/s) and an outlet on a lateral wall near the floor (P = 0.0Pa). The initial air temperature was 25ºC.

The mesh was created with Salome V6. It consists of 32,000 nodes and 62,000 linear triangular elements. The solution is transient. Navier-Stokes and Heat equations were solved in a strong coupled way. No turbulence model was used. Free convection effects were included. The mesh of the benchmark analysis was a much finer one of the same geometry domain, consisting of about 500,000 linear triangular elements. The warm section on the floor was removed and lateral boundaries had open condition (P = 0.0Pa).

Job Execution
The submission of jobs to be run at AWS was done through the web interface of the CloudBroker Platform. The procedure was as follows:

- A job was created on the CloudBroker Platform, specifying Job Name, Software, Instance Type and AWS Region
- Case and mesh partition files were compressed and uploaded to the CloudBroker Platform attached to the created job
- The job was submitted to the selected AWS resource
- Result files were downloaded from the CloudBroker Platform and postprocessed in a local workstation
- Scalability parameters were calculated from job output log file data

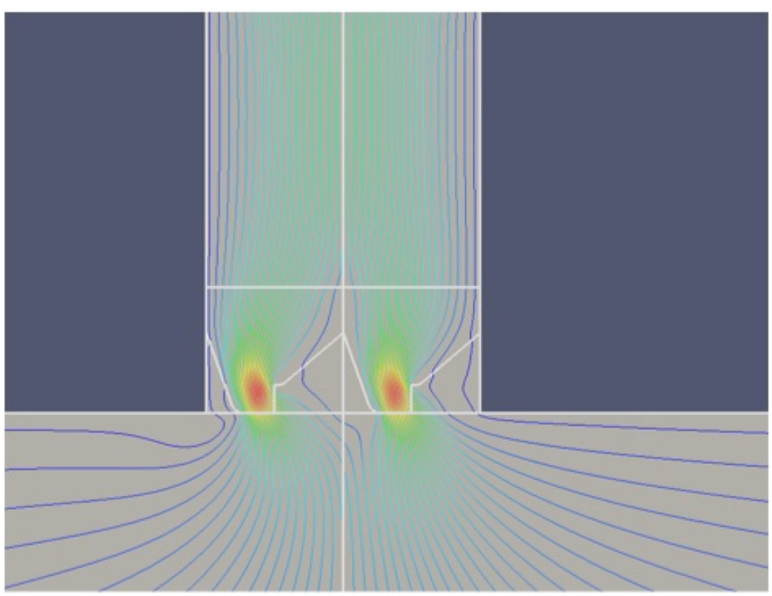

Fig. 2 – Streamline on the inlet section.

CHALLENGES

End User
The first challenge for BCSL in this project was to learn if the procedure to run Elmer jobs in a cloud computing resource such as AWS is easy enough to be a practical alternative to in-house calculation servers.

The second challenge was to determine the level of scalability of the Elmer solver running at AWS. Here we encountered good scalability when the instance employed is the only computational node. When running a job on an instances using two or more computational nodes the scalability is reduced dramatically, showing that communications between cores of different computational nodes slows down the process. AWS uses 10Gbit Ethernet as backbone network, which seems to be a limitation for this kind of simulations.

After the scalability study with the mesh of 500 Kelems was performed, a second scalability test was tried with a new mesh of about 2000 Kelems. However, jobs submitted for this study to Cluster Compute Quadruple Extra Large and Cluster Compute Eight Extra Large instances have not been successfully run yet. Further investigations are in progress to better characterize the network bottleneck issue as a function of problem size (number of elements per core) and to establish if it is related to MPI communication latency or NFS throughput of the results.

Resource Provider and Team Expert
On the technical side, most challenges were mastered by already existing features of the CloudBroker Platform or by small improvements. For this it was essential to follow the stepwise agile procedure as outlined above, partly ignoring the stiffer framework suggested by the default HPC Experiment tasks on Basecamp.

Unfortunately AWS HPC Cloud resources are limited to a 10 GBit Ethernet network. 10 Gbit Ethernet was not sufficient in terms of latency and throughput to run the experiment efficiently on more than one node in parallel. The following options are possible:

1. Run the experiment on one large node only, that is the AWS Cluster Compute Eight Extra Large instances with 16 cores
2. Run several experiment jobs independently in parallel with different parameters on the AWS Cluster Compute Eight Extra Large instances
3. Run the experiment on another cloud infrastructure which provides low latency and high throughput using technology such as Infiniband

The CloudBroker Platform allows for all the variants as described above. Variants 2 and 3 were not part of this experiment, but would be the next reasonable step to explore in a further experiment round. In the given time, it was also not possible to try out all the different I/O optimization possibilities, which could provide another route to improve scalability.

A further challenge of the HPC Experiment was to bring together the expertise from all the different involved partners. Each of them has experience on a separate set of the technical layers that were needed to be combined here (actual engineering use case, Elmer CAE algorithms, Elmer software package, CloudBroker Platform, AWS Cloud).

For example, often it is difficult to say from the onset which layer causes a certain issue, or if the issue results from the combination of layers. Here it was essential for the success of the project to stimulate and coordinate the contributions of the team members. For the future, we envision making this procedure more efficient through decoupling – for example, by the software provider directly offering an already optimized Elmer setup in the CloudBroker Platform to the end users.

Finally, a general challenge of the HPC Experiment concept is that it is a non-funded effort (apart from the AWS credit). This means that the involved partners can only provide manpower on a “best effort” basis, and paid projects during the same time usually have precedence. It is thus important that future HPC Experiment rounds take realistic business and commercialization aspects into account.

BENEFITS

Concerning the ease of using cloud computing resources, we concluded that this working methodology is very friendly and easy to use through the CloudBroker Platform.

The main benefits for BCSL regarding the use of cloud computing resources were:

- To have external HPC capabilities available to run medium sized CAE simulations
- To have the ability to perform parametric studies, in which a big number of small/medium size simulations have to be submitted
- To externalize all IT stuff necessary to have in-house calculation servers

For CloudBroker, it was a pleasure to extend its platform and services to a new set of users and to Elmer as a new software. Through the responses and results we were able to further improve our platform and to gain additional experience on the performance and scalability of AWS cloud resources, particularly for the Elmer software.

CONCLUSIONS AND RECOMMENDATIONS

The main lesson learned at Biscarri Consultoria SL arising from our participation in HPC Experiment Round 2 is that collaborative work through the Internet, using on-line resources like cloud computing hardware, Open Source software such as Elmer and CAElinux, and middleware platforms like CloudBroker, is a very interesting alternative to in-house calculation servers.

A backbone network such as 10Gbit Ethernet connecting computational nodes of a cloud computing platform seems not to be suitable for computational mechanics calculations that need to be run on more than one large AWS Cluster Compute node in parallel. The need for network bandwidth for the solution of strongly coupled equations involved in such simulations makes the use of faster network protocols such as Infiniband necessary to achieve time savings when running it in parallel on more than a single AWS Cluster Compute instance with 16 cores.

For CloudBroker, HPC Experiment Round 2 has provided another proof of its methodology, which combines its automated web application platform with remote consulting and support in an agile fashion. The CloudBroker Platform could easily work with CAELinux and the Elmer software at AWS. User requirements and test outcomes even resulted in additional improvements, which are now available to all platform users.

On the other hand, this round has shown again that there are still needs – for example, a reduction of latency and improvement of throughput (i.e., by using Infiniband instead of 10 GBit Ethernet) to be fulfilled by dynamic cloud providers such as AWS regarding highly scalable parallel HPC resources. Their cloud infrastructure is currently best suited for loosely or embarrassingly parallel jobs such as parameter sweeps, or highly coupled parallel jobs limited to single big machines. Finally, despite online tools, the effort necessary for a project involving several partners like this one should not be underestimated. CloudBroker expects though that in the future more software like Elmer can be directly offered through its platform in an already optimized way, making usage more efficient.

Case Study Authors - Lluís M. Biscarri, Pierre Lafortune, Wibke Sudholt, Nicola Fantini, Joël Cugnoni, and Peter Råback.

Thank you for your interest in the free and voluntary UberCloud HPC Experiment.

To download similar case studies go to:
http://tci.taborcommunications.com/UberCloud_HPC_Experiment

If you, or your organization would like to participate in this Experiment to explore hands-on the end-to-end process of HPC as a Service for your business then please register at:
http://www.hpcexperiment.com/why-participate

If you are interested in promoting your service/product at the UberCloud Exhibit then please register at
http://www.exhibit.hpcexperiment.com/