ff

FORTISSIMO

D10.13

First Tranche of Success Stories

Workpackage:	10	Dissemination and Training	
Author(s):	Francis Wray		scapos
Authorized by	Guy Lonsdale		scapos
Reviewer	Tomi Ilijaš		ARCTUR
Reviewer	Gregor Pipan		XLAB
Reviewer	Daniel Vladušič		XLAB
Dissemination Level	PUBLIC		

Date	Author	Comments	Version	Status
2014-12-05	Francis Wray	Initial draft	V0.0	Draft
2014-12-13	Francis Wray	Further draft using complete set of success story templates and feedback from the QA reviewers	V0.1	Draft
2014-12-17	Francis Wray	Revisions following clarifications from experiments and further input from the QA reviewers	V0.2	Draft
2014-1219	Francis Wray	Version accepted by QA reviewers	V1.0	Final

Executive Summary

The Fortissimo project is concerned with the development of a marketplace for Cloud-based-High Performance Computing (HPC) simulations, particularly targeting SMEs in the engineering and manufacturing sectors. The development of this marketplace is driven by experiments which both capture its requirements and which demonstrate the potential of industrial simulations running on an HPC Cloud. This document summarises and collates the results of an initial tranche of 20 experiments. In total, three such tranches involving ~50 experiments are anticipated over the course of the Fortissimo project.

The objective the first tranche has been to demonstrate the potential of Cloud-based HPC across a range of applications, mostly from the manufacturing and engineering sectors. In all cases, Cloud-based HPC has been used to simulate some aspect of an engineering or manufacturing process. The contents of this document will be taken forward to develop a set of material, promoting the Fortissimo Marketplace and the benefits of HPC-based simulation, using a range of appropriate media including flyers, videos and web pages.

20 experiments were conducted using HPC systems available at 8 different centres across Europe in Sweden, the UK, The Netherlands, France, Germany, Slovenia, Italy and Spain. Most experiments comprised an end-user, a computer centre and either a commercial software package or an open-source software package and a domain expert from the computer centre or from a further company such as a consultancy. Some of the end-users were completely new to HPC-based simulations and others had a level of exposure varying from moderate to high.

Most of the experiments were able to demonstrate a quantifiable financial benefit from their results and most demonstrated additional benefits such as quicker times to market. In many cases it was confirmed that Cloud-based HPC has significant financial advantages over owning and maintaining an in-house HPC system and that the high entry costs of purchasing an HPC system can be overcome by the pay-per-use paradigm inherent in the Cloud approach. It should be noted that the benefits of the pay-per-use paradigm apply also to the provision of HPC and domain expertise offered by relevant experts. The following benefits were demonstrated by the experiments:

Improvements to Industrial Processes

- A reduction in the times for the simulation of designs by a factor of 10 is possible;
- HPC can lead to a reduction in development and manufacturing costs;
- Improved process control is possible when HPC enables better simulations.

Cost-benefit Analysis

- It is 10 times cheaper to use a Cloud-based system than own and maintain an HPC system in some relevant cases;
- Significant financial benefits result in some cases with a return on investment of 10 times, that is, for a simulation costing X, a financial return of 10X can be realised;
- The break-even point where it makes financial sense to use a Cloud-based system rather than in-house hardware has been demonstrated. This point is around 50% of the peak system capability;
- Significant savings in the overall design process have been demonstrated.

Already significant benefits have been demonstrated from the use of HPC simulation. 30 further experiments are now planned. These will add considerably to the weight of material already developed and will demonstrate not only the value of HPC, but also the benefits the Fortissimo Marketplace will bring to SMEs through the wider availability of and support for HPC simulation.

Table of Contents

1 Introduction

The Fortissimo project is concerned with the development of a marketplace for Cloud-based-HPC simulations, particularly targeting SMEs in the engineering and manufacturing sectors. The development of this marketplace is driven by experiments which both capture the requirements that marketplace should offer and which demonstrate the potential for simulations of industrial processes running on an HPC Cloud. This document focuses on demonstrating the potential of that marketplace. It comprises deliverable D10.13 of the Fortissimo project, First Tranche of Success Stories. In total, three such Tranches of success stores are anticipated, over the course of the Fortissimo project, based on the outcomes of ~50 experiments.

To support the development of these Success Stories, Work Package 10, in conjunction with the Fortissimo Project Management, developed a template to enable each of the experiments to report its outcomes in a uniform way. This document is concerned with the summary of these templates, the homogeneous presentation of the experiments and a discussion of their outcomes. It collates and summarises success stories from each of the first tranche of 20 Fortissimo experiments.

The objective of all these experiments has been to demonstrate the potential of Cloud-based High Performance Computing (HPC) across a range of applications, mostly from the manufacturing and engineering sectors. In all cases, Cloud-based HPC has been used to simulate some aspect of an engineering or manufacturing process. This document constitutes the deliverable D10.13. Its objective is to capture and summarise information from each of these experiments to demonstrate the benefits of simulation using Cloud-based HPC resources, particularly to SMEs in the manufacturing and engineering sector. The contents of this deliverable will be taken forward within work package 10 of Fortissimo to present well documented business cases for the use of HPC-based simulation. This will result in a set of material, promoting the Fortissimo Marketplace and the use of HPC-based simulation, using a range of appropriate media including flyers, videos and web pages.

This document will serve as a guide and reference for its target audience of SMEs in the manufacturing and engineering sectors to the benefits of HPC-Cloud-based simulation and will support the partners in the Fortissimo project in the exploitation of its outcomes, particularly involving that target audience.

This document is the first of three, each covering a successive tranche of experiments. It comprises four main sections:

- Section 1 comprises this introduction;
- Section 2 comprises edited summaries of the context and business benefits of each of the experiments;
- Section 3 comprises a discussion of the overall outcomes of the experiments taken together, lessons learned and conclusions drawn; and
- Appendix A which contains the completed, unedited Success Story Templates for each of the experiments. Some templates also contain some graphics which may be used in the promotional material.

D10.13 is a public deliverable. However, the contents of Appendix A, the completed Success Story Templates are work in progress and will be superseded by the promotional material: flyers videos and web pages. These Success Story Templates will be presented in the deliverable D10.13 for review, but will be suppressed in the public version of the deliverable. However, the public deliverable will contain the promotional material or links to where it can be viewed.

2 Summary Success Stories

This section comprises a description of each of the 20 experiments carried out in the Fortissimo project using remote HPC services. The objective of each of the projects has been to demonstrate the benefits of using Cloud-based HPC with a particular focus on SMEs in the manufacturing and engineering sectors. It is anticipated that, in these sectors, the pay-per use approach to HPC simulation for hardware, software and domain specific expertise will remove barriers to its use by SMEs with resultant economic impact. This section serves to demonstrate ways in which that impact might be realised.

Each subsection of section 2 presents the outcomes of the experiments in a homogenous way, despite the wide-ranging areas covered by the experiments. To support that homogeneity, the following headings have been used for each experiment namely:

- Overview of the Experiment;
- The Industrial Setting;
- Improvement to the Industrial Process;
- A Cost-benefit Analysis;
- Other Business Benefits;
- The Need for an HPC-Cloud.

In section 3, a discussion of the experiments, the lessons learned and conclusions brings together several themes running across the experiments.

2.1 Experiment 401: Aircraft Design

2.1.1 Overview of the Experiment

The objective of this experiment was to evaluate the benefits of using Cloud-based HPC in the design of light aircraft. The flow of air over the surface of aircraft has been simulated in order to improve its design. This experiment enabled the end-user, Pipistrel, an aeroplane designer and manufacturer, to design its aircraft in a more optimal and cost-efficient way. For this experiment, a standard software package running on an HPC system was used to simulate the flow of air over the aircraft.

2.1.2 The Industrial Setting

There are two ways how to examine a behaviour of the flow of air over an arbitrary body: the first one is a wind tunnel test of a physical, usually scaled, body and the second one is to simulate the flow of air in a computer. Although both options have pros and cons, a standard practice for large companies in the aerospace industry is to use both in a consecutive manner. Computational simulations are usually used throughout the whole design phase, whereas wind tunnel tests are used only at certain phases of the design, since they are much more expensive from the cost and time point of view.

For an SME it is virtually impossible to use wind tunnel tests during the design phase of a new aircraft, because this option is simply too expensive. The only option an SME has is to simulate the flow of air as accurately as possible user a high-performance computer. To replace wind-tunnel tests satisfactorily, aerodynamic models with a high degree of fidelity to the real world need to be deployed. Such models require significant compute cycles and memory.

2.1.3 Improvement to the Industrial Process

The use of Cloud-based HPC allowed Pipistrel to run simulations of higher fidelity than was possible with its in-house systems. These simulations closely modelled real-world behaviour and gave accurate information on how the airplane behaved in flight. To simulate the flow with the required degree of accuracy, a large computer model was used for the experiments. In doing this, Pipistrel learned how to run, handle and post-process such big cases on a cloud-based HPC system. A typical large model would run in approximately 2 to 3 days on the HPC system. Such a problem would either be too big for the in-house systems to would take too long to run for an effective design process (around 20 to 30 days). The use of HPC therefore enabled Pipistrel to obtain results of much more complex simulations in a reasonable time.

2.1.4 A Cost-Benefit Analysis

Pipistrel needs to simulate the flow of air over the body of an aircraft only occasionally during the design process. It estimates that it is 10 times cheaper to use Cloud-based HPC simulations than have a suitably powerful in-house system which is only used for part of the time. The indicative annual costs of using Cloud-based HPC simulations are approximately €30k, which shows that this saving is considerable.

2.1.5 Other Business Benefits

This experiment allowed Pipistrel to use HPC for the first time and to learn about its capabilities. Pipistrel ran more demanding, higher fidelity simulations. It gained considerable experience in the use of HPC-based simulation. This experience will help Pipistrel to estimate the time and the cost of such simulations better. This will help them to decide if the use of HPC is justified or not in future projects.

Pipistrel learned that the use of HPC will be very valuable during a design phase of future aircraft. HPC can be used to run much more demanding simulations that improve the fidelity of results. The time needed for such simulations running on an HPC system is roughly the same as the coarser simulations currently run on Pipistrel's in-house cluster. The higher-resolution simulations give more and better data that can be incorporated into each design phase. This accelerates the design phase and reduces the number of the design cycles.

2.1.6 The Need for an HPC-Cloud

In this experiment simulations requiring high computer power and large amounts of memory were run. On the HPC system it was clearly demonstrated that these large simulations scaled well. In other words, by doubling the size of the computer used, the time to a solution was halved. HPC systems are characterised by much higher-speed internal communication bandwidth than standard cluster offerings. From the experiments carried out, it was clear that an HPC system was necessary for the simulations and that a standard cluster would not have the necessary internal communication bandwidth to run these simulations efficiently.

2.2 Experiment 402: Steel Casting

2.2.1 Overview of the Experiment

The pouring of liquid steel from the ladle to the tundish is the most critical stage of the continuous casting process. In order to prevent surface oxidation, the molten metal in the ladle is covered by a layer of protective slag, which may be partly transferred into the tundish during the final stage of ladle emptying. Slag transfer into the casting lines may cause the breakage of the solid skin of the solidifying cast product, an event with serious economic consequences which is also hazardous to the operatives. Quantitative knowledge of the flow of steel in the ladle emptying mechanism is required in order to develop an effective slag detection system, based on vibration analysis.

The objective of this experiment is to simulate the flow of steel from the ladle to the tundish using Cloud-based HPC resources so that Ergolines can use the results of the simulation to develop such an effective slag detection system.

2.2.2 The Industrial Setting

In the field of continuous casting there is an increasing industrial demand for the development of new technologies for preventing slag transfer from the ladle to the tundish. Such an event may cause the breakage of the solid skin of the solidifying cast products, which results in the hazardous dispersion of liquid steel within the industrial plant. Ladle-slag monitoring is currently performed by operatives on an empirical basis. Given the safety and economic implications of such an event, there is a significant demand for an effective, automated system for ladle-slag control. While passing through the ladle shroud, liquid slag induces characteristic vibrations which can be measured. In order to develop an effective detection system, it is necessary to correlate the vibrational signal to the fluid flow. This experiment is therefore concerned with simulating the flow of steel using a simulation package running on a Cloud-based HPC system. If successful, the outcome of such a simulation would be an innovative automated system for ladle-slag control with significant economic and safety implications.

2.2.3 Improvement to the Industrial Process

Dedicated numerical simulations based on Cloud-based HPC resources has enabled Ergolines' researchers to perform a thorough analysis of the different ladle emptying mechanisms. As a result, it has been possible to establish a correlation between the shroud vibrational signal and the fluid dynamic of the system. Based on this, it is now possible to develop a novel technology for the detection of ladle slag based on vibrational analysis. By preventing slag from passing from ladle to tundish, this innovative system would reduce the risk of breakouts, thus addressing key safety issues and contributing to an increase steel plants' productivity. In order to prevent slag transfer into the tundish, the pouring of liquid steel must be interrupted before the ladle is completely empty.

Excellent agreement between numerical predictions and experimental results demonstrated the effectiveness of the simulations, which enabled a key understanding of the correlation between the measured vibrational signal and the fluid dynamics of the system. These very relevant results constitute the basis for the future development of a dedicated system for ladle slag detection, an innovative technology featuring a high commercial potential in the field of continuous casting.

Due to the complexity of the physical phenomena involved, dedicated numerical simulations imply an exceptional computational load, thus requiring the use of HPC. Participation in the experiment provided Ergolines with the necessary HPC resources needed to address this scientific and industrial challenge. Dedicated numerical simulations have been performed which allowed for accurate modelling of ladle emptying. The fluid dynamics of liquid steel and slag has been studied by considering both a single-phase and a double-phase configuration. Realization of a dedicated laboratory scale model allowed for successive validation of the numerical results.

2.2.4 A Cost-Benefit Analysis

Given the complexity of the phenomenon to be simulated, a very fine discretization in terms of geometry and time is needed in order to obtain accurate results. Such a fine discretization involves a significant computational load and therefore requires adequate computational capabilities. As the company does not possess the necessary computational infrastructure, the possibility to use cloud-based HPC resources proved fundamental in addressing this specific industrial and scientific challenge. In fact, participation in the Fortissimo experiment allowed Ergolines to exploit supercomputing resources and reduced computational times without having to cover the costs of owning and managing a dedicated infrastructure.

The ability to detect the slag while it is passing through the shroud would enable the steel plant to delay the closing of the ladle, with an estimated saving of 1% of its production costs. In order to produce 1 million tons of steel per year, a medium size factory must sustain steel costs of up to 300 M€. Envisioned annual savings for a medium size factory are around 3 M€ for low-medium quality steel, rising significantly for higher-quality cast products. Fast return on investment is thus guaranteed, together with full coverage of the overall simulation costs. In fact, if no access to HPC resources and services had been granted, the overall simulation costs, comprehensive of hardware and personnel, would have been ~ €130,000.

2.2.5 Other Business Benefits

Besides providing significant business benefits for the company, the use of an HPC cloud in the present experiment will also bring relevant economic advantages to the steel casting industry. In fact, HPC-based fluid dynamic simulations enabled Ergolines to acquire the key scientific competencies needed for the development of an innovative system for ladle slag control, thus significantly reducing the time to market and improving product design. By reducing the risk of breakouts, this novel technology would in turn bring relevant economic benefits to the continuous casting industry, while contributing to increase steel plants productivity: the ability to detect the slag while it is passing through the shroud would in fact enable the steel plant to advance the closing of the ladle, saving up to 1 ton of steel for each casting sequence, with an estimated saving of 1% of the production. In order to produce 1 million tons of steel per year, a medium size factory must sustain steel costs up to 300 M€. Envisioned annual savings for a medium size factory are of the order of 3 M€ for low-medium quality steel and may rise significantly for higher quality cast products.

Dedicated HPC-based simulations followed by experimental validation provided Ergolines with key insight into the physics of the system. The results obtained will constitute the basis for the development of an innovative slag monitoring technology, which would significantly contribute to increase both occupational safety and steel plants productivity.

2.2.6 The Need for an HPC-Cloud

Simulating complex physical phenomena, such as the fluid dynamics of ladle-emptying mechanisms, requires a computational load that cannot be satisfied by the company's current infrastructure. The possibility of using cloud HPC resources offers the unique advantage of exploiting high-performance computational capabilities together with domain specific expertise without having to sustain the managing costs of standard cluster offerings.

The scalability of the simulation has been analysed, namely the improvement of calculation times resulting from parallelising the simulation on an increasing number of cores. Although only a limited speedup and scalability were noted on an HPC system, (only a factor of 8), a standard cluster offering would have shown no speedup because of the limitations in its internal network. Although a factor of 8 is disappointing from a technical point of view, it nevertheless represents a significant commercial benefit because of the improvement it brings to the design process.

2.3 Experiment 403: Design of High-Pressure Vessels

2.3.1 Overview of the Experiment

This experiment simulates the properties of pressure vessels made from composite materials using an open-source package and a Cloud-based HPC system. Such vessels are used in the storage of natural gas as the fuel for automotive vehicles. Composite materials have the potential to replace widely used steel and aluminium, often with better performance. Composite components can lead to weight savings of between 60% to 80% by replacing steel components and between 20% to 50% by replacing aluminium components.

The objective of this experiment is to improve the design of high-pressure vessels used for the storage of natural gas in vehicles using HPC-Cloud-based simulations.

2.3.2 The Industrial Setting

Composites are multiphase materials and offer the possibility for the modification of their properties or the addition of new functionality by a proper choice and combination of different phases. The biggest advantage of modern composite materials is that they are many times lighter than, and yet as strong or stronger than widely used metals. By choosing an appropriate combination of matrix and reinforcement material, a new material can be made that exactly meets the requirements of a particular application. Composites also provide design flexibility because many of them can be moulded into complex shapes. Mikrosam manufactures winding machines which are used in the manufacture of high-pressure vessels. It is involved in the sale of winding machines, the design of pressure vessels and the setting of its machines to produce such designs.

Filament winding has become a popular construction technique in a wide variety of industries for creating composite structures with high stiffness-to-weight ratios. Filament-wound composite pressure vessels offer significant weight savings over conventional all-metal pressure vessels for the containment of high-pressure gases and fluids.

This experiment will simulate the use of composites in the construction of pressure vessels particularly with regard to the storage of natural gas as a vehicle fuel. In the design of such vessels it is mandatory that certain standards are followed. The simulations presented take full account of such standards.

2.3.3 Improvement to the Industrial Process

Composite pressure vessels should take full advantage of the extremely high tensile strength and high elastic modulus of the fibres from which they are made. One approach to model the behaviour of a composite material is micro-mechanics, where the interaction of constituent materials is examined in detail to determine the behaviour of a heterogeneous composite material. The use of simulations using the properties of a range of materials has decreased design and testing costs by 10% because less material is wasted.

The experiments have shown that using HPC-based simulation can reduce the time for composite design by about 30%. The use of HPC resources allows more simulations to be made and reduces the amount of actual 'physical' testing, which has to be done to design vessels using composite materials.

2.3.4 A Cost-Benefit Analysis

For each filament-winding machine used in the production of high-pressure cylinders, different combinations of materials and winding angles need to be taken into account. Using Cloud-based-HPC simulation enables a final design to be tested with a reduction in the number of expensive physical prototypes needed.

Without using simulation, the price for production to Mikrosam of one physical prototype is €250, and one filament winding machine requires on average 10 units to be tested to validate both its functionality and the design of the final product, the total material cost is €2,500 and the labour cost is €1,459. With Cloud-based HPC simulation there is a saving of 10% on material and 20% on labour. Taking into account that the computing costs for the simulations are €100, we arrive at a net saving of €442 for the design of each pressure vessel and the settings of the winding machine. This represents a saving of over 10% in the overall design process. For a winding machine manufacturer such as Mikrosam, developing and designing a range of pressure vessels, the overall cost of the design process can be very significant.

2.3.5 Other Business Benefits

Using optimization in composite design in order to simulate the micro-macro mechanical properties of a filament-wound high-pressure vessel is an innovative idea which will result in improved services and together with increased competitiveness. Improvement of the time needed for the design of composite laminates may add 10% to 20% to the positive evaluation of Mikrosam's potential customers for winding machines.

2.3.6 The Need for an HPC-Cloud

This experiment has demonstrated the benefits of Cloud-based-HPC simulation to Mikrosam. An essential feature has been the support it has had from its HPC provider and from an HPC expert. This has been central to the success of this experiment. Such support would not have been available in a standard cloud offering.

2.4 Experiment 404: Simulation of Tightening Flanges

2.4.1 Overview of the Experiment

Flanges are joints between metallic parts which are used by many manufacturing companies in various sectors such as oil and gas, chemical and pharmaceutical. This experiment addresses flanged joints that are used for high pressure and temperature gases or liquids. To seal the flanges a gasket is placed inside a groove which is located on both flanges. During the process of closing and tightening the joint, the gasket and the flanges may suffer deformations and high levels of stress which, if higher than the materials elastic limit, may damage the flange. This means that it cannot be reused, once it has been opened (i.e., for maintenance) or may result in sealing problems, leakages or a decrease in the joint's operational life. Even worse, if the closing process is done inadequately, workers at the plant could suffer serious or fatal accidents accompanied by damage to the surrounding community and the environment.

The experiment develops a model for simulating the tightening process and a front-end application to control the parametric simulations, which will permit improvements to the design of the tightening process.

2.4.2 The Industrial Setting

Texas Controls is a Spanish SME that offers large industrial facilities, technological products and services oriented towards instrumentation, tightening and sealing solutions. In the field of tightening and sealing processes, today Texas Controls is the sole provider of engineering solutions. As such, one of its priorities is to provide the highest-quality solutions.

These solutions are especially important to customers of the gas and oil industry. In these sectors, probably more than in any other, it is imperative to avoid leaks in pipes and pressure vessels or reactors that could occur under extreme pressure and temperature conditions. The costs of preventing leakage are much lower than those generated by its consequences. However, many variables must be considered in order to effectively prevent hazards: the sealing process involves plastic interactions between elements so an expected deformation may occur in certain areas. The response and the service life of each and all of the components will be determined by the extreme and varying pressure and temperature conditions. This is certainly a complex and fascinating engineering problem to be solved. Therefore, it is crucial to study and predict the behaviour of flanged joints and to understand how the elastic interactions between elements are produced.

2.4.3 Improvement to the Industrial Process

Industrial clients often require immediate responses for maintenance or critical incidents. This is because the infrastructures affected by tightening and sealing technologies regularly play a major role in providing energy to large populations and to industry. Thus, to run all the parametric cases simultaneously in order to get the urgently required results as soon as possible is therefore the best possible solution.

With the simulation model, developed in this experiment, Texas Controls will be able to conduct analysis before the tightening process begins. Thus, time and money will be saved at the same time as damage to components will be avoided, and greater safety for facilities, workers, the community and the environment will be achieved.

It should be noted that currently the simulations being attempted for the challenging text case do not yet converge. This appears to be a model problem because two different software packages both fail to produce a convergent result. Technical work is continuing to understand this problem and to realise the significant potential benefits of this simulation.

2.4.4 A Cost-Benefit Analysis

A key benefit of using the Fortissimo HPC Cloud will be the reduction of operational time in the maintenance of huge and critical oil and gas heat exchangers or reactors. Another fundamental benefit will be avoiding damages in flanges during tightening processes.

Based on previous field work experience (of a time past when simulations were not used), a non-optimised tightening process on a 24 stud bolt flange consumed 108 man-hours while a current optimised tightening process consumed 72 man-hours. The optimisation allows the technicians to reduce the number of tightening passes from 6 to 4, in all this represents a 33% time saving per flange. This obviously represents considerable savings in labour costs. Such savings provide a competitive edge for Texas in international tenders.

With the simulation model, developed in the experiment, Texas Controls hope to be able to conduct analysis before the tightening process begins. Thus, time and money will be saved at the same time that damage to components will be avoided, and greater safety for facilities, workers, the community and the environment will be achieved.

2.4.5 Other Business Benefits

The most important short-term benefits are qualitative rather than quantitative, but in the mid-term quantifiable economic impact is anticipated

Relevant improvements are the reduction of the uncertainties in the process of tightening and the increasing safety of the final result of the process. This is because it can be guaranteed that in any step of the process neither the flange nor the gasket is overstressed. These factors results in reduced maintenance costs of facilities in the long-term. The cost of a leak is much higher than the cost of avoiding it.

2.4.6 The Need for an HPC-Cloud

Flange tightening is by no means a regularly scheduled activity. Peaks and troughs in workloads are common. Tightening procedures in the same installation can involve several different tightening processes for flanges of different types. In many cases it is necessary to perform many simulations in a short period of time. However there could be fairly long periods in which no simulations may be required. HPC cloud offers fast response and the number of resources required in each case.

Industrial clients for our solutions often require immediate response times for maintenance and critical incidents since the infrastructures affected by tightening and sealing technologies regularly play a major role in providing energy to large populations and industries. Thus, to run all the parametric cases simultaneously in order to get the urgently required results as soon as possible is therefore the best possible solution.

From a financial standpoint, it is indeed unthinkable for a small engineering SME such as Texas Controls to make the regular investments needed to maintain an operational and updated HPC cluster to provide the necessary simulating infrastructure to respond to its unusually high yet sporadic demands.

2.5 Experiment 405: Design of High-Voltage Cables

2.5.1 Overview of the Experiment

Cables for high-voltage power transmission are subject to electromagnetic losses during operation, when a load electric current flows along the conductor. It is important to model the behaviour of such cables under load to improve designs. The objective of this experiment is simulate the behaviour of high-voltage transmission cables as supplied by Prysmian, world leader in the energy and telecom cables, and systems industry, in order to improve their designs. This experiment will use a combination of open-source and proprietary software packages to perform the necessary simulations.

2.5.2 The Industrial Setting

The losses in a cable for high-voltage power transmission arise from several sources all of which need to be modelled. These sources include the resistive heating, dielectric losses inside the insulation layer and induced currents in the metal screening and armoured layers. These losses are dissipated as heat.

During steady state the normal operating conductor temperature is not allowed to exceed a prescribed limit, so as not to increase the ageing of electric insulation and so reduce cable operating life. Most of the losses are directly affected by the load current. This raises the problem of determining the rating of a cable system that is the intensity of the electric current that the system can sustain on a continuous basis without overheating. This involves:

- Correct computation of all loss contributions inside the cable;
- Precise modelling of heat transfer phenomena with the surrounding environment.

Up to now, modelling and calculation activities have relied on formulas and procedures mostly based on analytical or semi empirical models that have the advantage of being easily implementable into spread sheets or self-made routines. However, in many cases these calculations are affected by a high level of approximation, which may result in excessive overdesign of the systems. The increasing need to optimise the performance of cables together with challenging installation and operation conditions demands more sophisticated methodologies. Cloud-based-HPC simulations now offer designers powerful tools to meet the stringent requirements for future high-voltage cable transmission systems.

2.5.3 Improvement to the Industrial Process

The use of Cloud-based HPC simulation has enabled Prysmian to calculate more accurately the losses within a cable and to model the transfer of heat from the surface of cables to the surrounding environment. This simulation has involved the use of both proprietary and open-sources computer codes running on an HPC system. The use of this advanced modelling offers the opportunity to address in detail thermal transport phenomena that play a critical role in the design and operation of high-voltage cable transmission systems.

2.5.4 A Cost-Benefit Analysis

Tests performed during this experiments have now shown that the use of HPC-based simulation using a Cloud and external expertise will result in considerable savings with respect of acquiring internal resources, over a three year period, as described later in more detail.

In this cost-benefit analysis, there are three costs to be considered. These are the cost of computing, the cost of qualified personnel to set-up and perform the simulations and the cost of software licences. To perform the simulations in house, these costs were estimated to be respectively €300k, €80k and €30k over the three year period, totalling €380k. Based on the outcome of this experiment and its benchmarking tests, to perform these experiments using a Cloud-based HPC system from a computer centre with high-quality technical support and appropriate licences would cost 30% less, a very significant saving.

2.5.5 Other Business Benefits

The other benefits of Cloud-based-HPC simulations can be summarised as follows:

- Increased speed, as a result of obtaining simulation results within hours instead of weeks.
- Increased flexibility, occurring mainly through services utilization on demand and payment out of specific project budgets.
- Improved quality, as a result of the production of an enhanced solution based on higher resolution models.
- Faster design processes, driven by optimization made possible by increased CFD simulation throughput
- Increased competitive advantage on the basis of the above mentioned benefits. This means the potential edge for winning projects with typical budgets of around €10-100 million;
- Lower entrance barrier for company wishing to use Cloud-based HPC simulation.

2.5.6 The Need for an HPC-Cloud

The Cloud-based-HPC approach has four key advantages over the use of a standard cluster:

- The availability of expertise both in the technical area of the simulation and in the optimisation and running of such codes;
- The availability of computing nodes with very large RAM and high-bandwidth interconnect which are not available on standard cluster offerings such as Amazon EC2; and
- The availability of remote visualisation techniques, again not available on standard cluster offerings.

2.6 Experiment 406: HPC Tools for Urban Planning

2.6.1 Overview of the Experiment

IES develops world-leading suites for the simulation of buildings and their environments. The purpose of this experiment is to demonstrate the use of a tool depending on HPC simulations because of the scale of the simulations being addressed. Previously this tool had been used for small-scale simulations running on a workstation. In this experiment the workstation will be used to visualise the outcomes of the simulations running on a Cloud-based HPC system.

The main objective will be to enable Virtual Environment (VE) desktop installations and web-based interfaces to access the calculation resources hosted on an HPC-cloud infrastructure.

2.6.2 The Industrial Setting

IES have developed a planning tool for cities, which will enable relevant stakeholders to assess, for example, the energy efficiency of a city, quality of living and more. The tool can be used at any stage of a city's life, and can be used in cities only beginning the journey towards a 'Smart' City or those that are well on their way towards sustainable advancement and integration with 'Smart' Technologies. This interactive decision support tool can be used by City Architects/Planners, Urban Planning Consultancy Firms, the Public and Business Community, ICT Renewable Energy Technology system suppliers and potential investors in a City.

This tool relies heavily on the availability of HPC cloud-based simulation because of the very large amounts of data associated with multiple buildings and their interaction within the urban context.

2.6.3 Improvement to the Industrial Process

Planning of cities requires access to a huge mass of varied data. The use of Cloud-based HPC can enable detailed simulations enabled by this data to be displayed on a desktop using the VE. Furthermore these simulations can be performed in a fraction of the time that a workstation would need thus enabling these detailed simulations to become a feasible, commercial reality, making the process easier, simpler and more precise for all involved. Furthermore, the ability to perform very fast simulations creates the possibility to analyse "what-if" scenarios thereby greatly increasing the usefulness of the simulations.

2.6.4 A Cost-Benefit Analysis

Prior to launching a service, building model simulations of various sizes ranging from small to very large were tested. Typical speedups (comparing the workstation to the HPC system) were between 5 and 10 times. The major benefit is that simulations taking unrealistically long compute times of days or weeks on the workstation could run in a few hours or days using the Cloud-based HPC system with clear commercial benefits. Indeed, based on the successful experiment, IES is now offering an HPC-based service to its customers. This service embodies a pay-as-you-go approach which is underpinned by HPC-systems available from the EPCC computer centre at the University of Edinburgh.

2.6.5 Other Business Benefits

Additional business benefits of using IES's VE system backed up by a Cloud-based-HPC system in Edinburgh include:

- Faster roll-out of new features to end-users;

- End-users always have access to the latest software version;
- Support teams are able to assist end-users better since support staff are immediately able to access the user models on the server;
- A similar credits payment system can be used to control end-user payment for support.

2.6.6 The Need for an HPC-Cloud

The VE of IES is concerned with time-related simulations. The nature of these integrated dynamic simulations requires that nodes of the HPC-system are closely coupled for optimal performance. The IES VE Suncast engine has also been optimised to use MPI on HPC systems such as those available in Edinburgh. The savings and benefits achieved thus far in the Fortissimo experiment are not possible without access to the larger number of cores per node available on the EPCC HPC platform. This large number of closely connected nodes are not available on standard Amazon EC2 offerings.

2.7 Experiment 407: Airflow Over a Vehicle

2.7.1 Overview of the Experiment

This experiment addresses the provision of cloud-based supercomputing to enable more in-depth Computational Fluid Dynamics (CFD) simulations used within the automotive industry. The open-source package, OpenFOAM, hosted on a Fortissimo platform has been deployed to tackle some representative external aerodynamic problems widely used in automotive design.

2.7.2 The Industrial Setting

The automotive industry heavily relies on CFD for its design needs. Traditionally, commercial CFD software has been used. Given the growing complexity of these simulations, HPC is often the only realistic way to produce results within an acceptable time frame. The commercial software licensing models, which often charge on a per-core or per-CPU basis, can become very costly, making them beyond the reach of SMEs. For this reason, SMEs are increasingly moving to the alternative open-source solutions.

OpenFOAM is a well established, open-source general-purpose CFD package. It is known to have capabilities comparable to its commercial counterparts. It is widely adopted in both academic and industrial settings. And it is well supported by both dedicated commercial offerings and freely available community help. These factors make OpenFOAM a credible alternative to those commercial software packages. The simulations reported here have used OpenFOAM on an HPC system.

2.7.3 Improvement to the Industrial Process

The combination of an open-source software package and Cloud-based access to HPC systems enables both large and small companies to have access to significant computing resources on a pay-per-use basis. This is very attractive to many companies. They do not need to have a very significant outlay to buy a large machine which is only used for a fraction of the time, yet which is powerful enough to perform simulations quickly in a burst mode. The economics of this are clear and unless an HPC system is loaded at more then ~50% of its capacity then it is more cost-effective to buy cycles rather than buy and support a hardware system. Furthermore it is advantageous for many small companies to move computing costs from Capital Expenditure to Operational Expenditure with consequent smoothing of cash flow.

2.7.4 A Cost-Benefit Analysis

To demonstrate the cost-effectiveness of Cloud-based HPC, running a representative job for 60 hours on 512 cores would cost €3,900 using a pay-per-use on-demand approach. Running the same job on a comparable machine in-house system used to 100% of its capacity would be ~€1,500, but if it were only used 30% of its capacity, the cost of the calculation would be ~€5,000. This clearly demonstrates the benefits of the Cloud-based pay-per-use approach.

2.7.5 Other Business Benefits

There are quantitative benefits of using a well managed Cloud-based system which include data storage, secure mechanisms for data sharing among users, and additional tools for data analysis and visualisation.

2.7.6 The Need for an HPC-Cloud

Simulations on an HPC platform can achieve very good scalability, therefore making more efficient use of computing resources, in particular if the cloud-based service is charged on a pay-per-use basis. Furthermore, this high scalability means that very ambitious computations can be undertaken. Such computations may not be feasible on a standard cluster offering because its limited internal bandwidth may limit its scalability.

2.8 Experiment 408: Optimisation of Aeroplane Wiring

2.8.1 Overview of the Experiment

This experiment uses a documented industrial use-case provided by KE-works and the design optimization company Noesis Solutions in a simulation aimed at reducing the weight and cost of wiring systems in complex products such as aircraft. The objective of the experiment is to demonstrate that Cloud-based-HPC simulation can provide a cost-effective way of running such applications.

2.8.2 The Industrial Setting

There are many auxiliary companies serving the aeronautic industry, some of them SMEs. An important problem in this industrial sector is the optimisation of wiring. This can be addressed using computer models, but applications to do this are very computationally intensive. SMEs generally do not have the available resources to buy and maintain, in house, the necessary large computer systems. Neither do they have the required expertise to use such systems.

The availability of Cloud-based-HPC systems with associated technical expertise is a potential solution for SMEs. Computationally intensive simulations can be run on a pay-per use basis, which is much cheaper than owning and maintaining a large system, big enough to satisfy peak demands. Similarly external expertise can be made available also on a pay-per use basis.

2.8.3 Improvement to the Industrial Process

The traditional trial-and-error approach of European SMEs to the design the wiring layout in aeroplanes lies in the past. Computer-based optimisers now do this, reducing costs and improving results. However, such computations are extremely computationally intensive. Bringing together supercomputing resources, domain specific expertise and licensed simulation packages in a one-stop shop can overcome this barrier and enable SMEs to compete without the need for very large capital investments in high-performance computers. Through the use of Cloud-based HPC, SMEs can compete with large enterprises, utilising thousands of cores to reduce costs and times to solution.

This experiment showed that a wiring optimisation application would run 20 times faster when using 30 nodes in parallel, than on a single node. Such a result is of significant industrial relevance.

2.8.4 A Cost-Benefit Analysis

Comparing the current non-optimal process with the automated and optimized process in a cloud HPC environment shows 2.5% reduction in cost and weight of the wiring system. This benefit is a recurring benefit for hundreds of ship sets that are produced across several decades.

The cost of computation on a Cloud-based HPC system is about EUR 660 for a single experiment run, In contrast, the cost of an in-house HPC system is about 61 kEUR / year. Because of the limited number of runs per year, having an in-house HPC facility is much more expensive than using cloud HPC. Furthermore, a cloud-based HPC system is much more flexible in terms of the number of processors which can be applied to a particular optimisation if more computational power is required. This demonstrates the feasibility and cost-effectiveness of using Cloud-based HPC for engineering simulations. It also demonstrates that SMEs are able to compete with larger organisations in the use of HPC, because the cost barrier has been eliminated.

The use of a concept such as the Fortissimo one-stop shop which could bring together supercomputing resources, domain specific expertise and licensed simulation packages could overcome these barriers. This would enable European SMEs to compete globally through better design and quicker times to market.

2.8.5 Other Business Benefits

The use of Cloud-based HPC simulations results in faster response times to Requests for Proposals, high-quality estimates and the offer of lower-cost and lower-weight wire harnesses due to optimized routing.

2.8.6 The Need for an HPC-Cloud

The Cloud-based HPC platform combines high-performance hardware with a support team, which is able to maximize the efficiency of resource usage.

This experiment would not have been possible without the help of such a support team.

Furthermore, fast interconnection between nodes has remarkably reduced the simulation time. This offering is not available on standard commercial clusters.

2.9 Experiment 409: Simulating Automobile Components

2.9.1 Overview of the Experiment

BULL is a company which supplies HPC systems and related services. SDI is a market leader in developing software to simulate the mechanical behaviour of automobile components. Scilab Enterprises develops and publishes the open-source software, Scilab, for numerical computation. The objective of this experiment is to demonstrate a cloud-based version of Scilab which can be combined with SDI software to offer a cloud-based simulation service.

2.9.2 The Industrial Setting

The automobile industry is a highly competitive and secretive business where all parties need to have the confidence that their proprietary information is secure. This particularly applies to simulations involving new vehicles. Two different scenarios were tested in this experiment. The first is based on a brake manufacturer and a vehicle modeler working together, but remotely on a cloud-based infrastructure. The second addresses a main contractor wishing to compare the performance between 2 brake manufacturers. Each use case illustrates well different levels of confidentiality between main and sub-contractors.

2.9.3 Improvement to the Industrial Process

The possibility of launching simulations via a Cloud-based HPC system will not only help to increase the numbers of scenarios which can be evaluated, but it will also reduce overall compute times. For an end-user, the benefit of using a Cloud-based HPC system could be either an increase of productivity or a commercial advantage compared to company using a classic client-server model. Further advantages can be realised by avoiding the need to maintain a computing cluster in-house with associated hardware costs and need to provide related in-house expertise.

2.9.4 A Cost-Benefit Analysis

The results of the experiment are very promising. The speed-up obtained scales well with the number of processors deployed. That is if N processors are used the simulation runs nearly N times faster. This paves the way to new optimization techniques where parameters are chosen automatically in a great number of simulations. This contrasts with the current approach at SDI where an engineer carefully selects only those parameters deemed important. On the basis of the experiments carried out, SDI expects that a simulation service, relevant to the needs of its customer base, would run on a Cloud-based HPC system in only 30% of the current time needed at a cost of only €100 per month.

2.9.5 Other Business Benefits

One of the goals of the experiment was to demonstrate that data could be handled securely. Sharing a simulation with business partners in a unified manner without exposing the know-how was satisfactorily demonstrated during the experiment.

2.9.6 The Need for an HPC-Cloud

Currently, Scilab has limitations when running on large numbers of cores in an HPC system because it has not yet been optimised for such systems. However, a version of Scilab, optimised for HPC systems with high-performance interconnect, is under development. Scilab 6 will be able to take full advantage of such HPC systems to give a performance and to enable simulations not realisable on a standard cluster.

2.10 Experiment 410: Prediction of Air Quality

2.10.1 Overview of the Experiment

European regulations on air-quality protection require more and more the testing and evaluation of adaptation and reduction scenarios. From the traditional evaluation of 2 to 3 scenarios, consulting companies and regional air-quality agencies are now facing the need to evaluate tens of scenarios. This requires a large increase in their capacity for computing beyond what they can easily manage with in-house resources.

The objective of this experiment is to demonstrate the use of Cloud-based-HPC services to investigate air-quality at the scale of cities. The experiment uses the ADMS-Urban software from CERC running on the Extreme Factory HPC offering from BULL.

2.10.2 The Industrial Setting

Not only is there a need to increase the number of scenarios to be evaluated as described in the previous section, but it is also required to quantify uncertainties associated with air-quality simulations. This further increases the number of scenarios to be modelled. The demonstration of a cost-effective service based on Cloud-based-HPC simulation would represent a breakthrough and solution for the end-user faced with the need to evaluate this increasing number of scenarios.

2.10.3 Improvement to the Industrial Process

The possibility of running urban air-quality simulations using Cloud-based HPC will help to increase the numbers of scenarios which can be feasibly simulated. Using an HPC system will also reduce the computational time needed for such simulations. The outcome of using such a system includes shorter times for simulations with resultant competitive advantage and the cost reduction of using pay-per use resources rather than owning and maintaining an expensive in-house system which may only be used on a part-time basis.

2.10.4 A Cost-Benefit Analysis

An analysis can be made of the cost-benefits of using a Cloud-based HPC system as follows:

- If there is no constraint on the time to produce results for a typical urban air-quality study (i.e. simulations can take 3 months or more), the cost to exploit an HPC cloud system is higher than the costs of to acquiring and maintaining an internal server. In contrast, if the time to perform a simulation is short (1 month or less), then the exploitation of a Cloud-based HPC system (using pay on demand) is economically interesting (from a factor 2 to 10 cheaper). This can result in an increased margin while offering the client a reduction of time to complete a study. Alternatively the HPC system can support additional services such as the evaluation of uncertainties or the testing of more scenarios.
- When the average percentage of use of an internal server falls below 40% (regardless of the number of cores), a pay-on-demand cloud service becomes economically viable compared with the costs to acquire and maintain that server. If the average percentage use of internal servers is low, then the benefits of using a Cloud-based system can be significant. It should be noted that an important part of the cost for internal server is the cost of IT staff which may be a significant burden for an SME.

2.10.5 Other Business Benefits

The partners in this experiment can now offer a software solution for evaluating air-quality at the scale of cities as a pay-per-use Cloud-based service rather than requiring a customer to purchase an annual licence and install and run the software locally. This enables an attractive pricing option for customers needing only an infrequent use of the software and its underlying model.

2.10.6 The Need for an HPC-Cloud

The ADMS-Urban software could be run on less sophisticated architectures such as those offered by standard clusters, but with less efficiency, less security, and a perceived cost which does not take into account the time and efforts to setup an operational HPC cluster on an Amazon-type infrastructure.

2.11 Experiment 411: Simulation of Risers in the Oil Industry

2.11.1 Overview of the Experiment

DeepLines® provides oil and gas engineering companies with an integrated tool to design and optimize their risers, flowlines and moorings systems for shallow to deep-water offshore projects. DeepLines is used daily by engineering companies in the oil and gas sector as well as the marine renewable energy sector. The objective of this experiment is to determine the benefits of running DeepLines on a Cloud-based-HPC system.

2.11.2 The Industrial Setting

The interest of having a version of DeepLines running on HPC cloud is to enable both Principia and its DeepLines customers to use these optimizations to obtain results more quickly, but also to do larger case studies to represent more accurately the offshore environment seen by the system being designed over two to three decades.

2.11.3 Improvement to the Industrial Process

The main objective of this experiment is to perform simulations that are not accessible right now. Due to computer limitations, the current practice in the offshore industry is to simplify the meteocean data to limit the number of simulations necessary for the design whilst still being conservative. Due to the risk involved, all simplifications are overly conservative and can lead to a design which are not optimized. This is a general trend in the offshore industry where systems are to stay in place for 20 to 40 years. Over the past decade, the number of cases to run has dramatically increased both in sheer number (now above 1000), but also in term of simulation length (from 500s to 10800s).

2.11.4 A Cost-Benefit Analysis

Currently, the cost advantage is limited since half of the necessary calculations (post-treatment) are not available on an HPC system. However, once the post-treatments are available, the time will be saved across the whole calculation chain, thereby reducing delays and the costs of studies, despite the additional costs for licences and the computer cycles.

For the first HPC model some changes were made to Deeplines involving the development of a post-processing tool. For simple case studies test, the compute time reduced from 60 days on the customer's system to 3 days on the HPC system with significant benefits.

2.11.5 Other Business Benefits

No other business benefits have been identified at this stage.

2.11.6 The Need for an HPC-Cloud

The use of HPC systems is more advantageous compared to a standard cluster such as Amazon for this experiment. The main reason is the capability to tailor the configuration of CPU, RAM and especially disk type to the specific needs of the computation.

2.12 Experiment 412: Reduced Vehicular Emissions of CO_2

2.12.1 Overview of the Experiment

This experiment deals with the optimization of a vehicle's architecture in order to minimize CO_2 emissions. This involves running large-scale parameter studies involving the use of vehicular simulation software running on a Cloud-based HPC system. The objective is to demonstrate to SMEs in the automotive industry how perform large-scale simulations without the need for a large primary investment in computational hardware and permanent software licences. In outline, the experiment comprises:

- Calculating vehicle CO2 emissions using the simulation software tool AVL CRUISE;

- Analyzing simulation results using the post-processing software tool AVL CONCERTO;
- Selecting an optimal vehicle architecture by running large number of vehicle model variants in a distributed cloud environment
- Preparing data on a desktop environment for computations on a remote system;
- Enabling large-scale computations at the high performance computer centre, HLRS;
- Addressing the end-user through a secure, user-friendly web-based front-end platform, featuring model/data sharing and the export of results.

2.12.2 The Industrial Setting

Many vehicle system optimization tasks in the automotive industry require the simulation of a large number of variants of a particular vehicle model. In the automotive sector, such simulations are often performed by SMES. AVL supplies software which performs such simulations and visualises their output. In order to undertake large-scale simulations in the automotive industry, SME service providers could either:

- run simulations of different vehicle variants sequentially on one machine with one software license (clearly this increases the time taken to get results); or
- acquire a larger number of CPUs and multiple software licenses (clearly this costs more)

Many SMEs which are suppliers to the automotive industry and which are dealing with mid-scale, single-focus and short-term projects, cannot afford owning and maintaining large hardware and software infrastructures. A new approach is needed to offer simulation resources to those SMEs in an economically competitive way. This is a pay-per use approach where software licences and computing resources can be made available of a pay-per-use basis.

2.12.3 Improvement to the Industrial Process

In contrast to the traditional approach, this experiment utilises Cloud-base HPC computing which allows many simulation tasks to run simultaneously, following to a customisable schedule and results to be obtained in a shorter time.

This simulation service together with a tool for the analysis of results can be offered to SME end-users within a user-friendly, web-based front-end environment, enabling them to access and perform large-scale parameter studies as needed, when needed, whilst preserving control over progress and results.

2.12.4 A Cost-Benefit Analysis

The most clear cost benefit to SMEs of using Cloud-based HPC resources is the possibility to lease a powerful computer system for single projects instead of acquiring and maintaining computational resources which would be underutilized most of the time and which would probably not be sufficient when really needed. Clearly, the pay-per-use approach both the computing cycles and to software licences is of significant benefit to such organisations. This leads to a better management of computational and development resources, opening a whole new set of applications such as extensive parameter variations.

The monetary benefit varies for different use cases. A good analogy would be the monetary benefit of buying an airplane ticket to fly from A to B, instead of buying and maintaining a private airplane which will be used just a couple of times a year. Roughly, for example, if an end-user needs 400 concurrent (annual) SW licences ,to run an application on 400 cores, and an appropriate computing infrastructure to run a large scale parameter study, an investment of millions of Euros would be required. Taking all costs into account, a short-term project running millions of simulations on 400 cloud-based CPU cores for a couple of weeks, several times a year would bring significantly more than 50% savings, when compared to the cost ownership of the necessary hardware and software licences. In particular the saving alone in computing costs for 400 cores running for 10 weeks would be around €50k to €100k.

2.12.5 Other Business Benefits

Different business models (short-term lease, pay-per-use e.g. per CPU-time, per simulation time usage or per simulation run) can make cloud-based computing attractive to SMEs and may lead to new applications and cost effective offers for simulation of a range of products and services.

2.12.6 The Need for an HPC-Cloud

In contrast to standard computing clusters, Cloud-based HPC is more suitable for this type of task because it offers more standardized utilization of resources in terms of hardware and software, queuing systems, status checks error handling etc. This enables users to focus on the engineering tasks rather than the computer science issues.

2.13 Experiment 413: Multiphysics Simulations

2.13.1 Overview of the Experiment

The objective of this experiment is to provide multiphysics simulation services using Cloud-based HPC where computer cycles, HPC and domain expertise and software licences are available on a pay-per-use basis. The sector addressed is the electromotive where equipment such as transformers, motors and actuators need to be simulated and designed.

2.13.2 The Industrial Setting

A fast growing number of engineers running manufacturing and production applications know that they should enhance their current simulations with multiphysics features and tools to make them more realistic. However, there are barriers to this. Those engineers will only start software evaluations if they can be sure to get a valuable feedback from their tests after a very short time that is a few day. This is a good opportunity for the use of Cloud-based HPC systems where the pay-per-use approach means that quick results can be obtained without the need for significant investment in hardware and software.

2.13.3 Improvement to the Industrial Process

The improvement to industrial processes will come through the availability of a wide range of simulation tools covering several industrial sectors involved in engineering and design. These simulation tools will be available on a pay-per-use basis which make them readily accessible by SMEs which would otherwise not be able to afford the necessary hardware and software investment to use such tools.

2.13.4 A Cost-Benefit Analysis

A simple cost-benefit analysis shows that Cloud-based HPC simulations using the pay-per-use approach give better value on the following basis.

- If the usage of an HPC system is less than ~50% of its capacity it is cheaper to use a Cloud-based approach than buy and maintain a system;

- If less than ~1500 hours of compute time is used per year, it is cheaper to use an HPC-Cloud on a pay-per-use basis than buy yearly software licences. This, of course, varies from vendor to vendor.
- If HPC or domain expertise is needed, then it is cheaper to get this resource from an HPC Centre rather than employ a specialist member of staff, unless the use of HPC is very intensive.

2.13.5 Other Business Benefits

The major advantage for the multiphysics newcomer will be the general possibility to test a new simulation environment through a cloud service instead of installing new and complex software packages locally just for an evaluation run.

2.13.6 The Need for an HPC-Cloud

The benefits of having both HPC and sector specific advice are important to SMEs which may not have that expertise in house. Such support is generally not available in a standard cloud environment. Furthermore several of the simulation packages need to run on an HPC system with high internal bandwidth. Such a feature is not available in standard cloud offerings.

2.14 Experiment 414: Calculating Properties of Hazardous Materials

2.14.1 Overview of the Experiment

The objective of this experiment is to simulate the thermophysical properties of substances using Cloud-based HPC systems. Such calculations have significant advantages compared to physical testing and may even calculate properties which it is not possible to assess in the laboratory.

2.14.2 The Industrial Setting

Nowadays, powerful predictive methods, using computer-based simulations, exist that calculate the thermophysical properties of compounds. These are the basis for the design and optimization of chemical engineering processes. Typically, the chemical industry measures the required data experimentally. However, if they are needed for hazardous substances (explosive, toxic or mutagenic), the associated costs of physical testing may be prohibitive. In such cases, computer-based simulations are very attractive alternatives. To carry out such simulations, a large number of model runs are necessary requiring a very powerful computer.

2.14.3 Improvement to the Industrial Process

The main improvement to the industrial process is the increased flexibility the Cloud-based simulations bring and to the significant reduction in costs there use brings. Also the use of HPC-based simulation can take place much more quickly than physical measurements with the inherent advantages of faster time to market.

2.14.4 A Cost-Benefit Analysis

There are significant cost-benefits in the use of HPC-based simulations.

Experimental pure component densities cost around 2,700 € per substance for a very limited temperature and pressure range, when bought from an external supplier. Compared to that, 60 molecular simulations carried out in the entire fluid region up to arbitrary high pressures will cost around €1,600 yielding not just the density but every static thermodynamic property simultaneously.

For mixtures of compounds, the difference in costs is more extreme. 60 experimentally measured gas solubility data points of a binary mixture can cost up to €50,000, while the cost for the molecular simulation usually increases by a factor of two compared to a pure component (€3,200). It should be noted that the prices for the experimental data gathered by physical measurements only apply to moderate conditions and non-hazardous substances. Such measurements at high temperatures or pressures can be much more expensive or even impossible to conduct.

Lonza is a company which manufactures chemical apparatuses and needs to know these thermophysical properties. The use of simulation can bring massive savings to Lonza's production process. For example, for the design of a distillation column costing €1.5 million, the following cost calculation can be made:

Physical Experiment:

Staff costs: 100 person days at €1,000 = €100,000

Computer-based simulation

Cost of cycles = €6,400

Staff costs: 4 person days at €550 = €2,200

Subsequent physical measurement for verification: Staff costs 5 person days @1,000 = €5,000

Total: €13,600

***Overall saving:* €86,400**

2.14.5 Other Business Benefits

A clear benefit of the HPC-based simulations is that they can be performed much more quickly than physical experiments with the advantage of faster times to market.

2.14.6 The Need for an HPC-Cloud

The internal details of the code used in the simulations require the use of an HPC system with high-bandwidth connectivity in order to perform the calculations efficiently. This would not be possible using a standard cluster offering.

2.15 Experiment 415: Design of Moulds for casting

2.15.1 Overview of the Experiment

IMR produces foundry equipment for brass alloys and bronze. This equipment comprises moulds and casting channels. This equipment used in 80% of the cases for the production of taps and valves and in 20% for other high performance artefacts. The moulds mounted on the installation are filled by injecting the liquid metal at temperatures of about 1000 ° C.

The objective of this experiment is to simulate the flow of metal in a mould using open-source software in order to improve the design of moulds and also the design process.

2.15.2 The Industrial Setting

The success of the entire system depends on the quality of the castings (taps and valves) obtained. To avoid cracks and defects, it is therefore essential to design the moulds and the casting channels to obtain a laminar flow in the liquid metal and a constant cooling gradient throughout the final piece.

Currently, the design of moulds and filling channels is entrusted to the experience of mould makers in collaboration with experts from the foundry. Generally, the exchange of experience has obtained good results, but it is often necessary to modify the mould several times and repeat tests before committing to production. The application of HPC-based simulation avoids the need to make expensive, physical redesigns of the mould.

2.15.3 Improvement to the Industrial Process

SMEs operating in this sector use the traditional method of test-error-correction-new test. Simulation is not generally used in this sector because of the high cost and the difficulty for an SME to use of the software. Furthermore, processing times with normal processors are very long, about 4 to 6 hours, for each change of parameters. The possibility of using a HPC with open-source software allows a considerable saving of time and investment. This experiment has demonstrated that an SME with limited knowledge of HPC can get good results from the simulation so reducing the time to design and test the moulds.

The impact on the market and increased competitiveness will bring great benefits in terms of reduced lead times, reduced development costs and a growth in the market for IMR.

2.15.4 A Cost-Benefit Analysis

The cost of designing a mould using the traditional trial and error is approximately €41k. The use of Cloud-based HPC in this process produces a saving of around €8k because fewer iterations and less testing are involved. The cost of the computing cycles needed to perform these simulations is ~€2k. This results in an overall saving of €6k per mould. Over a year, IMR will design 8 sets of moulds representing an overall saving of €48k. Against this the need to train IMR staff to use the system and the software is a once-off cost of around €3k.

The use of a Cloud-based-HPC system rather than the purchase and maintenance of an in-house system can represent a considerable saving. In particular when the utilisation of an in-house system falls below 50% of its peak capacity, then it is cheaper to use a Cloud. The use of open-source software represents a further saving.

2.15.5 Other Business Benefits

A principal benefit of HPC simulation is the reduction in time to market and a reduction in the time needed to design a mould with consequent cost savings. Furthermore the use of such simulation may open up new markets in similar areas.

2.15.6 The Need for an HPC-Cloud

In order to speed up the calculations significantly it was necessary to use a Cloud-based HPC system with good internal connectivity. A standard cluster would not be able to realise the level of speed-up shown by the HPC system.

2.16 Experiment 416: Structural Crash-Test Simulations

2.16.1 Overview of the Experiment

The objective of this experiment is to demonstrate the use of Cloud-based-HPC services to investigate the behaviour of automotive components in a structural crash-test simulation using a commercial software package. The software used will be the Virtual Performance Solution Software Suite from ESI running on the Extreme Factory system from BULL.

2.16.2 The Industrial Setting

Currently, market requirements related to the reduction of vehicle weight and the need to cut costs are driving the industry to greater innovation and to introduce new design processes, new materials and new manufacturing processes. The challenge for the supply chain, usually

SMEs, is to handle conflicting requirements and bring revolutionary changes to vehicles, while at the same time cutting development costs and times drastically.

The Green Vehicle initiative is a major driver for the automotive industry to answer environmental requirements and regulations. This challenge calls for disruptive innovation in vehicle design and material. Steel remains the main cost effective material for Body In White (a stage in automotive design in which a car body's sheet metal components have been welded together) mass production for the foreseen future. This challenge is a good business opportunity for new industry leaders such as GESTAMP to emerge, who have developed a new grade of steel called Press Hardening together with associated manufacturing processes. This is one of the major innovations introduced successfully to the market to address the reduction in the weight of vehicles with recognised advantages.

2.16.3 Improvement to the Industrial Process

The design of vehicles involves trade-offs between conflicting requirements such as weight reduction, performance requirements and manufacturing constraints. To answer this challenge, Virtual Prototyping was adopted with a holistic view in which several full vehicle simulation models all based on one unique central Body In White subsystem were used. This is called the Single Core Model approach. This disruptive approach enables a collaborative decision making process to be supported during the development phase.

This innovative development methodology calls for extensive computer resources. Cloud Computing seems to be the ad-hoc solution to this and is a major enabler. It offers the necessary flexibility to access HPC resources when they are needed at an affordable cost.

This approach should reduce the development period thus shortening time to market and to improve quality whilst retaining competitive pricing and so result in increased competitiveness.

2.16.4 A Cost-Benefit Analysis

A simple cost-benefit calculation shows that the use of the Single Core model and Cloud-based HPC solutions can result in an overall reduction in personnel effort by a factor of over 3. The compute costs are nearly the same for both the previous method using four independent applications and the Single Core Model. In particular, for a typical test run of the Single Core Model, 140 staff hours can be saved from the time needed to set up the model. This has benefits both in terms of cost and in terms of shortened development times.

2.16.5 Other Business Benefits

The use of HPC cloud services will also be a huge benefit for multi-site companies. HPC enables data to be accesses and shared, simulations to be run and results to be analysed in all the relevant locations through a conventional web browser. In summary, Cloud-based HPC applications will enable more efficient project handling by reducing data-transfer times and by improving the communication between the different stakeholders of a project.

2.16.6 The Need for an HPC-Cloud

Extreme factory is a pure HPC cloud offering, with optimized scalable architecture, dedicated management software and fast InfiniBand interconnect. It is also built from scratch with security and ease-of-use in mind. The Virtual Performance Suite software could be run on less sophisticated architectures, but with less efficiency, less security, and a perceived cost which does not take into account the time and efforts to setup an operational HPC cluster on an Amazon-type infrastructure.

2.17 Experiment 417: Optimised Design of a Supercar

2.17.1 Overview of the Experiment

This experiment brings together CINECA, an HPC centre, ICON, the owner of a Computational Fluid Dynamics (CFD) simulation code, NTUA, a research organisation specialising in computation and KOENIGSEGG, a niche, high-performance sports-car manufacturer. In this experiment, KOENIGSEGG has evaluated the use of Cloud-based-HPC simulations using ICON's software in the design of its sports cars.

2.17.2 The Industrial Setting

CFD is widely used in industry in a wide range of applications involving the motion of fluids. CFD can bring significant cost savings to product development. Currently, it is the fastest growing segment in Computer Aided Engineering.

Within the automotive industry, one of the major uses of CFD is to compute and optimize the drag and down force components for different car configurations (pitch/yaw angles, additional aero devices, overtaking, etc.). These then lead to a reduction in fuel consumption and improvements in the overall aerodynamic performance. Moreover, in the Formula One industry, intensive application of CFD is carried out in order to reduce the cost of wind tunnel testing. Models used to replicate real-life cars with high accuracy, consisting of every geometric detail such as rotating wheels and integrated components (including heat exchangers, fans and condensers) can be very large. Fully time-dependent simulations of such models are very complex and the use of an HPC solution, with the support of suitable domain-specific expertise, can make a significant difference. Past experiences have shown that the full aerodynamic design of a hyper car can be entirely conducted using CFD with minimal road and wind tunnel testing. However in a production environment tight deadlines must be met, placing an emphasis on the use of significant HPC resources.

2.17.3 Improvement to the Industrial Process

The use of ICON's simulation software on a Cloud-based-HPC system enables KOENIGSEGG to reduce or even avoid wind tunnel testing in some circumstances. Accessing powerful computing resources remotely also reduces hardware expenses and maintenance costs. Before the start of the Fortissimo Project, KOENIGSEGG had only limited computer resources available to them and little or no experience in HPC-based CFD. In 2013, 100% of the aerodynamic development of the KOENIGSEGG One:1 has been conducted using CFD. In less than eight months, hundreds of simulations to test various configurations have been carried out. The results were an impressive 250% increase in down-force with only a 15% increase in drag at 250km/h and with a 50% higher down-force at 440km/h, the vehicle's maximum speed.

2.17.4 A Cost-Benefit Analysis

Tests have shown that the use of HPC-based simulation supported by external expertise led to a return on investment in less than three months of production of a new car configuration. Significant costs can be saved and transferred to other critical parts of the development and production process.

The benefits obtainable by the use of the Fortissimo HPC-Cloud have been quantified as a 5% saving in operational costs, a 30% saving in design costs, a reduction of 50% in wind tunnel and physical testing, a 60% saving in prototyping costs, and a 30% shortening of the time to market. Furthermore, savings in development were about €90K per year on the design process, corresponding to a 1.5% reduction in overall development costs.

2.17.5 Other Business Benefits

ICON has already demonstrated the maturity of automatic optimisation technologies when combined with open-source CFD solvers, can offer significant cost benefits over competing manual CFD optimisation methods. Benefits from Cloud-based-HPC-cloud simulation are:

- Increased turnaround, as a result of obtaining simulation results within hours instead of weeks.
- Increased flexibility, occurring mainly through on-demand services and payment out of specific project budgets.
- Reduced cost, encompassing pay-per-use, reduced costs of depreciation and more effective deployment of labour.
- Improved quality, through higher resolution models.
- Optimised design enabled by faster simulation.
- A reduced barrier to the use of CFD by SMEs.

2.17.6 The Need for an HPC-Cloud

The benefits of CFD simulations are clear. By moving such simulations onto a Cloud-based-HPC environment, processing times can be significantly reduced. Moreover, the anticipated royalties and utilisation costs of such services can be drastically reduced, enabling the simulation code provider to reach to new markets, while SMEs that were unable to "purchase" such services can be given the opportunity to do so, moving from applications consuming significant effort and resources to more flexible SaaS solutions.

2.18 Experiment 418: Design of Centrifugal Pumps

2.18.1 Overview of the Experiment

This experiment simulates fluid flow in a centrifugal pump to optimise its design. Several runs of a commercial code will be made in order to investigate the range of design parameters of the pump. The objective is to find the optimum set of parameters for the particular design goals of the pump. This experiment has involved CINECA, a high-performance computer centre, and Enginsoft, a consultancy active in the area of pump design.

2.18.2 The Industrial Setting

Centrifugal pumps are widely used in many industrial applications, from oil and gas to water treatment, automotive and home appliances. Such devices may be required to operate over a wide range of flows. The prediction of operating characteristic curves is essential for a designer. Unfortunately, theoretical methods and experimental tests are unable to determine the source of poor performances sufficiently well. Simulation of such flows using an HPC system is becoming an important and common tool for pump designers. Many flows can be simulated much faster and cheaper than by means of physical experiments. Most importantly, the complex internal flows in water pump impellers can be predicted well. Consequently, Computational Fluid Dynamics (CFD) is now an established industrial design tool, helping to reduce design times and improve processes throughout the engineering world. Unfortunately to run such simulations it is necessary to use a high-performance computer which presents a significant cost barrier, particularly for SMEs. The availability of a Cloud-based HPC system together with expert support would help overcome this barrier.

2.18.3 Improvement to the Industrial Process

The numerical simulation of centrifugal pumps is not easy due to some difficulties including complex geometry, unsteadiness, turbulence, secondary flows, separation and boundary layers. These aspects require a high-fidelity CFD model, or, in other words, very fine computational grids and transient analysis. This approach is quite prohibitive for a typical SME. The aim of the Fortissimo HPC Cloud solution is to provide an attractive solution in terms of the business and technological perspective for SMEs which are unable to perform the necessary simulations using their own resources. The simulations carried out have already shown significant benefits for SMEs in the design of centrifugal pumps.

2.18.4 A Cost-Benefit Analysis

Tests have shown that the use of HPC-based simulation using a Cloud and external expertise will result in a return on investment in less than six months for an end-user. This means that a project to design a single pump with optimised performance can be run in a timeframe of six months. This contrasts with the usual 2 to 3 years timeframe involving the hiring of a junior engineer with a long learning curve. The availability of such simulations can give both pump manufacturers and associated consultancy companies a significant commercial advantage.

Enginsoft have made the following calculation. Comparing in-house costs of owning, maintaining, operating and providing expert support for an HPC system against those of a pay-per-use Cloud-based-HPC system with expert support, over a seven-year period, the Cloud-based solution is cheaper with the break even point coming at the end of that period. These costs are based on Enginsoft running around four Cloud-based projects per year each costing €40K. The overall cost is then €880K for the period. However, the in-house solution involves a high-level of risk because staff may move on. Certainly over the greater part of the seven-year period the Cloud-based-HPC solution is more cost effective.

HPC-based simulation has been used to optimise the design of centrifugal pumps. A prototype service is available to Enginsoft and its customers for pump design. This includes computing resources that were not previously available to the company. Due to these improvements, considering the average number of project lost in the last couple of years because Enginsoft could not be compliant to the requests of the end-user in terms of time-to-result, Enginsoft expects to increase its market share by at least 1% with a profit of about €100K. Considering the additional expenses for Enginsoft to use a cloud service instead of the internal resources, this represents a positive return on investment in only six months.

2.18.5 Other Business Benefits

The provision of a Cloud-based HPC solution would offer significant innovation into the marketplace. In particular the innovation would be derived from the possibility for the end-user to tackle simulations at an unprecedented level of refinement. This offer benefits both to Enginsoft as a Consultant firm allowing the opening to a market segment hitherto precluded, and to SMEs end-users which have now an instrument allowing a better, more refined, design with a reduced time to market (up to 50%).

In summary, customer value and benefits from such HPC-cloud CFD simulation could be summarised as follows:

- Increased speed, as a result of obtaining simulation results within hours instead of weeks;
- Increased flexibility, occurring mainly through utilization of services on demand and payment out of specific project budgets;
- Reduced cost, encompassing pay on demand capabilities, reduced depreciation costs, and more effective labour deployment;
- Improved quality, as a result of the production of an enhanced solution based on higher resolution models;

- Faster design processes, driven by optimization made possible by increased throughput of simulations;
- Increased competitive advantage on the basis of the above mentioned benefits;
- Lower entrance barrier for SME's approaching CFD

2.18.6 The Need for an HPC-Cloud

The numerical simulation of centrifugal pumps is not easy due to complex geometries, unsteadiness, turbulence, secondary flows, separation and boundary layers. These aspects require a high-fidelity CFD model, that is, very fine computational grids and transient analysis.

This translates into intensive calculations, that can be addressed in timescales suitable to a modern industry lifecycle management only with the careful combination of a software with very powerful hardware with high-connectivity between the processors. This is not available in standard clusters and so an HPC-system is essential for this type of computations.

2.19 Experiment 419: Simulation of Drifting Snow

2.19.1 Overview of the Experiment

This experiment comprises simulations of eolian snow transport in a civil engineering context using a Computational Fluid Dynamics (CFD) package. Such simulations allow the assessment of snow loads on the rooftops of buildings due to structures and machinery. The CFD software to be used is snowFoam, based on the CFD library OpenFOAM. Modifications will be made to snowFoam in an attempt to improve the performance of simulations for this particular application.

2.19.2 The Industrial Setting

Roof collapses due to accumulated and drifting snow are responsible for losses in the order of hundreds of millions of Euros as well as bodily injuries and loss of lives every year in Northern Europe in particular, and the Northern hemisphere in general. The simulations to be undertaken in this experiment can help identify rooftop areas that are prone to increased snow loads and provide the corresponding snow loads. Such analysis can help civil engineers design rooftop structures in perspective and avoid potentially catastrophic events. Additionally, simulations of drifting snow can help design protective barriers and fences for roads and highways that help prevent or mitigate drifting snow and reduce its disruptive impact on transportation.

Drifting snow loads are not properly addressed by present Building Codes. Moreover, the state-of-the-art in drifting snow simulations is highly empirical and limited to very specific situations. The leading technology used by Binkz in its snowFoam software does not suffer from such limitations, but is too computationally costly to use on desktop computers and even advanced scientific workstations. Consequently these simulations need to be run on an HPC system.

2.19.3 Improvement to the Industrial Process

The use of snowFoam on an HPC system allows engineers to design their rooftops for all applicable snow types and roof configurations. The downstream savings as compared to rooftop collapse and the potential for bodily injury, loss of lives, downtime of industrial operations and insurance and compensation costs is orders of magnitude larger than the analysis and computational cost that is typically €10,000-€50,000. Moreover, the analysis of drifting snow can be completed within a few weeks to a couple of months. These reasons make the use of snowFoam highly advantageous for engineering firms which are the principal users of accurate simulations of drifting snow.

2.19.4 A Cost-Benefit Analysis

Although the simulation of drifting snow is very computationally intensive, it can still be cost-effective. The downstream savings as compared to roof collapse and the potential for bodily injury, loss of lives, downtime of industrial operations as well as insurance and compensation costs is orders of magnitude larger than the analysis and computational cost that is typically 10,000-50,000 Euros.

Typical CFD consulting projects, such as computing wind loads and evaluating structural integrity, are provided with costs in the order of €15, out of which €5 - €10k is spent on computational resources, and typically take up to four weeks of overall throughput time. For the snow drift CFD consulting service, the cost and throughput time will have to be similar to this, in order to be sufficiently attractive for civil engineering companies.

Typically, a typical simulation of drifting snow requires 50k computing hours. With a price of 0.20 euro per CPU hour, the resource costs of €10k falls within the limits for these consulting projects. The scalability of the HPC cloud allows the simulations to be run in two weeks, well within the expected throughput time. It should be noted that the Cloud-based HPC approach makes significant computing resources available without the need for significant capital outlay which would not be acceptable in this market.

2.19.5 Other Business Benefits

As described above, the avoidance of roof collapse has significant benefits to a wide community.

2.19.6 The Need for an HPC-Cloud

To run these simulations an HPC system with high-bandwidth interconnect is essential. A standard cluster offering would not have the necessary bandwidth to run the simulations effectively.

2.20 Experiment 420: Computational Chemistry

2.20.1 Overview of the Experiment

This experiment addresses the modelling of large-scale molecular systems relevant to the chemical-engineering sector. The objective is to demonstrate the cost-effectiveness of simulating the behaviour of new compounds through the use of Cloud-based-HPC simulation.

2.20.2 The Industrial Setting

There is a clear industrial demand for faster, predictive computational chemistry simulations. Modelling is a proven powerful tool in chemistry research, providing key information for the design of new chemicals and materials. The software for modelling large-scale molecular systems has applications in sectors such as electronics (including semiconductors and other electric elements), organic chemistry and other areas of chemistry (fluids and polymers for foods, paints, home and personal care formulations, dyes, adhesives, petrochemistry, etc.) and materials design (such as alloys and ceramics for aerospace and automotive industry).

The Albemarle Corporation, a partner in this experiment, is a perfect example of a research and innovation-driven company with very significant research budgets. Albemarle's research into refinery catalysts is an important part of its green chemistry efforts, focused primarily on the development of more effective catalysts and related additives to produce clean fuels in a commercially viable way. Catalysis research relies on the kind of insight that only computational simulations can bring.

2.20.3 Improvement to the Industrial Process

The materials-by-design approach outlined above requires powerful computational resources and significant expertise to use them. This places a burden on innovative companies seeking to increase their competitiveness and reduce times to market by means of computational simulations. In particular, SMEs have smaller budgets and fewer IT staff than larger companies, and the increasing complexity of modelling platforms often outstrips the growth of their IT resources. In the materials modelling sector there exists a plethora of alternative methods and software implementations. Factors such as licensing costs, setup and maintenance make it harder for industrial modellers to use the best tool for the specific problem. Cloud-based HPC services, with their reduced costs and overheads, constitute a unique opportunity to provide dynamic European SMEs with powerful, easy-to-use and cost-efficient modelling tools, helping them to innovate and compete on the global scale.

This experiment involves modelling the behaviour of molybdenum sulphide clusters, as used in the catalyst-enhanced HydroDesulphurisation processes. Desulphurisation of feedstocks is an important process in refineries, preventing SO_x emission from transport fuels that are derived from these feedstocks, and is an important component of modern clean-fuel technologies. The cloud-based setup to remotely access simulation software on SURFsara's hardware is already allowing simulations which would have been computationally prohibitive with in-house resources available to Albemarle.

2.20.4 A Cost-Benefit Analysis

The benefits associated with using an HPC cloud are apparent, with clearly reduced hardware and software costs. A cloud-based HPC setup allows industrial end-users easily to scale up simulations and reach molecular system sizes that would have demanded far too long computation times with their in-house resources. In a traditional, non-cloud setup, this would have required an important investment in new hardware and software licences, as well as the manpower required to set up and maintain the software. For example, a low-end computing cluster can easily have a total cost of ownership (hardware, maintenance, power etc.) of around €60,000 per year and a commercial licence for a typical materials modelling package can well exceed €80,000 per year. On the other hand, commercial HPC providers charge around 2,500 EUR/year for the usage of 8 cores, clearly a more cost-efficient option. In general, end-users of cloud services avoid hardware licences and hardware maintenance and upgrade costs, and have reduced up-front software licence costs.

In this experiment, Albemarle's area of expertise is in chemical engineering, not hardware and software deployment and maintenance. The company can cut significant costs by outsourcing the latter to HPC specialists. The following table illustrates the yearly costs associated with a typical computing environment for materials modelling as required by Albemarle's simulation team (~128 CPU cores), comparing two different setups:

- An in-house computing system that includes hardware and software licences costs (which have been depreciated over 5 years), as well as the staff required to manage it.
- A dedicated HPC provider, in a business model where the end-user has 16 fully dedicated 8-core machines, and the usage of the software is charged as an additional 100% on top of the HPC costs. The price corresponds to a yearly contract. It assumes use of 25% of the total resources of the HPC provider's system and includes an important discount with respect to pay-per-use options.

Annual Cost Comparison between in-house resources and Cloud-based HPC		
	In-house System	**Cloud-based HPC**
Manpower	€50,000	€30,000[1]
Hardware	€16,000 [2]	
Software licenses	€16,000	€20,000
Total	€82,000	€50,000

[1] Includes manpower and cycles. The cost of cycles including supporting manpower is based on ~1 million cycles @ €0.03 per core hour ; [2] Based on 5 years' amortisation

It is apparent that hardware costs are not necessarily the most important ones, easily eclipsed by the staffing expenses required to manage a computing system (or even the software licenses). It is easy to see how this experiment has convinced Albemarle to allocate a sizeable budget for HPC cloud computing next year.

2.20.5 Other Business Benefits

However, the value of cloud computing goes beyond reducing hardware and software costs. Cloud-based HPC services allow companies to focus on their core skills to maximize growth and value. The agile and scalable character of cloud computing allows even small businesses to move in new and innovative directions to capture new markets, or to keep up with the market by growing as quickly as the market will allow. In this particular experiment, the end-user can test different modelling approaches without the full, up-front licence costs associated with each option.

3 Discussion, Lessons Learned and Conclusions

This section comprises a discussion of the experiments, lessons learned from their conduct and conclusions that can be drawn from the overall exercise.

20 experiments were conducted using HPC systems available at 8 different centres across Europe in Sweden, the UK, The Netherlands, France, Germany, Slovenia, Italy and Spain. The purpose of each experiment was to demonstrate the benefits of using HPC to simulate an industrial process. Most experiments comprised an end-user, a computer centre and either a commercial software package or an open-source software package and a domain expert from the computer centre or from a further company such as a consultancy. Some of the end-users were completely new to HPC-based simulations and others had a level of exposure varying from moderate to high.

Most of the experiments were able to demonstrate a quantifiable financial benefit from their results and most demonstrated additional benefits such as quicker times to market. In many cases it was confirmed that Cloud-based HPC has significant financial advantages over owning and maintaining an in-house HPC system and that the high entry costs of purchasing an HPC system could be overcome by the pay-per-use paradigm inherent in the Cloud approach. It should be noted that the benefits of the pay-per-use paradigm apply also to the provision of HPC and to domain expertise offered by relevant experts.

In the few cases where a quantifiable benefit could not be demonstrated, a quantitative benefit was nevertheless asserted, for example the greater convenience of using a Cloud-based HPC system rather than a standard cluster, where that convenience offered one or more business advantages. Some of these quantifiable benefits are given below under the heading "Other Business Benefits".

The following lessons were learned from the experiments and are reported in the experiments as presented above:

Improvements to Industrial Processes

- It is possible to realise a reduction in the times for the simulation of designs often by a factor of 10;
- It is possible to realise a reduction in development and manufacturing costs through the use of HPC;
- Improved process control is possible through better simulation using HPC.

Cost-benefit Analysis

- It is 10 times cheaper to use a Cloud-based system than own and maintain an HPC system in some relevant cases;
- Significant financial benefits result in some cases with a return on investment of 10 times, that is, for a simulation costing X, a financial return of 10X could be realised;
- The break-even point where it makes financial sense to use a Cloud-based system rather than in-house hardware has been demonstrated. This point is around 50% of the peak system capability;
- Significant savings in the overall design process have been demonstrated.

Other Business Benefits

- HPC enables higher-fidelity simulations leading to a quicker design process;
- The use of advanced simulation improves the reputation of companies through a quicker response to their customers;
- The experiments have enabled companies to become aware of opportunities for new products and services;

- The experiments have demonstrated the value of open-source software packages;
- The experiments have demonstrated the possibility of using very powerful systems in burst mode to perform big simulations quickly in an affordable way;
- Cloud-based HPC enables SMEs to compete with large companies and so drive an economic renewal in Europe;
- HPC offers the opportunity to evaluate many more scenarios in the design process;
- The experiments have demonstrated the value of open-source software particularly for SMEs;
- Companies participating in the experiments expect to realise an increased market share due to faster responses to customers and calls for tenders;
- HPC simulation is an enabling technology for improved services to the public sector particularly with regard to regulatory issues;
- HPC simulation offers the possibility to study compounds too hazardous for conventional assay.

This document summarises the outcome of the first 20 experiments in the Fortissimo project. Already significant benefits have been demonstrated from the use of HPC simulation. Around 30 further experiments are now planned. These will add considerable to the weight of material already developed and will demonstrate not only the value of HPC, but also the benefits the Fortissimo Marketplace will bring to SMEs through the wider availability of and support for HPC simulation. These experiments and their success stories will form the basis of a growing promotional campaign in Fortissimo’s Work Package 10, which will create awareness and stimulate the use of the Fortissimo Marketplace.

4 Appendix A: Completed Success Story Templates

4.1 *Experiment 401: Aircraft Design*

4.1.1 Overview of the experiment

The objective of the Experiment 1 (WP 401) named Cloud-based simulation of aerodynamics of light aircraft is to improve the fidelity of aerodynamic simulations in the regime where the laminar-turbulent transition flow phenomena in the boundary layer occurs. The experiment enabled the end-user Pipistrel, airplane designer and producer, to calculate the flow behaviour with a k-kl-omega turbulent model over the complete aircraft at relevant Reynolds numbers. Pipistrel is in this way able to design his aircrafts in more optimal and cost efficient way. During the experiment open source software OpenFOAM was used as a CFD tool.

4.1.2 The industrial setting

There are two ways how to examine a behaviour of an air flow passing an arbitrary body: the first one is a wind tunnel test of a physical, usually scaled, body and the second one is a computation fluid dynamic (CFD) simulation that is run on the computer. Although both options have pros and cons, a standard practice for large companies in aerospace industry is to use both in a consecutive manner. CFD simulations are usually used throughout the whole design phase, whereas the wind tunnel tests are used only at certain phases of the design, since they are much more expensive from the cost and the time point of view.

For a SME it is virtually impossible to use wind tunnel tests during the design phase of a new aircraft, because this option is simply too expensive. The SME is therefore left with CFD simulations that need to describe the flow behaviour as close to real world as possible in order to accurately estimate the forces that act on the aircraft. Reynolds averaged Navier-Stokes (RANS) simulations with fully turbulent models can give a good estimate of flow behaviour over a simplified shape at smaller angles of attack and can be run on smaller in-house computers/clusters. But if the wind tunnel tests are to be satisfactorily replaced by CFD simulations, aerodynamic models with higher complexity needs to be employed. Such improved fidelity aerodynamic simulations can be very time, computer memory and computer power consuming processes. A use of HPC services is therefore very appreciated for an aircraft designer when CFD simulations are the only tool to give the insight into aircraft in-flight behaviour.

4.1.3 The improvements in the industrial process through the use of an HPC Cloud

The use of HPC allowed Pipistrel to run simulations of higher complexity than was used to and that closer describe the real world and give better information how the airplane behaves in-flight. Pipistrel successfully used a low-Reynolds k-kl-omega turbulent model simulating an airflow past a complete aircraft at cruise speed conditions. Such turbulent model is able to estimate the laminar-turbulent transition location that significantly affects the performance. For this purpose a mesh with more than 115M cells was calculated where 13 wall layers were employed. The thinnest layer right at the wall was 1E-6m thick. Pipistrel therefore learned how to make a mesh with more than 10 times larger number of cells and more than 100 times thinner first wall layer then used to. At the same time Pipistrel learned how to run, handle and post-process such big cases on cloud-based HPC.

Pipistrel also showed that CFD simulations of a typical complexity that were run on an in-house cluster can be almost linearly speeded-up with the number of extra cores used on the HPC. A simulation with a mesh composed of 10M cells converged on the in-house cluster's 8 cores in one day whereas on HPC's 80 cores in approximately 2.5 hours. On the other hand simulations with 115M cell mesh converged on 180 cores in approximately 2-3 days. The use of HPC therefore enabled Pipistrel to obtain results of much more complex simulations in approximately the same time.

4.1.4 The cost benefits of using an HPC Cloud

Rental of HPC resources for approximately 600k CPUh is more than 10 times cheaper than buying and operating own cluster with 120 cores if an assumption is made that the cluster gets depreciated over the time period of two years.

CFD simulations done at Pipistrel are typically of a research nature, which means there isn't a constant need for HPC resources. The need depends on a case by case basis. It is therefore even more reasonable for Pipistrel to rent HPC hours when needed than to buy its own computer power.

The Fortissimo experiment in particular enabled Pipistrel to use HPC for the first time and in this way gave the opportunity to learn what are the capabilities of HPC. Pipistrel ran more computer power and computer memory demanding simulation, which improved the fidelity of aerodynamic simulations. During the experiment Pipistrel learned what is the proper workflow using a cloud-based HPC, it learned how to produce much finer mesh of better quality and it learned approximately how long does it take for the simulation of this kind to converge. The latter information will help Pipistrel in future projects to better estimate the time and the cost of such simulations ran on HPC and in this way to decide if a use of HPC is justified or not.

4.1.5 Any Other Business benefits

Pipistrel learned from the Fortissimo experiment results that HPC services will be very appreciable also during a design phase of some new future aircraft. On one hand HPC can be used to run much more demanding simulations that improve the fidelity of results. Time needed for such simulations on HPC doesn't increase substantially compared to simulations currently run on Pipistrel's in-house cluster. But they give more and better data that can be incorporated into each design phase, which accelerates the design phase and reduces the number of the design cycles. On the other hand HPC can be used also to run larger number of less demanding simulations at the same time, e.g. the complete 3D polar (dependence of lift and drag coefficients with respect to angle of attack and side slip angle) of an aircraft can be obtained in the same time as only one simulation at a chosen angle of attack and side slip angle on Pipistrel's in-house cluster. Again, much more data is in this way available to use in the design phase.

4.1.6 HPC as opposed to standard cluster offerings

In the experiment a very computer power and computer memory demanding simulations were run. Typically 60 and 180 cores have been used for mesh generation and simulation, respectively. OpenFOAM showed good scalability of simulations run on up to a few hundred cores in our case. OpenFOAM tipically shows a good scaling for cases that are sufficiently large that the problem is not dominated by inter-processor communication. We did some tests of scalability, but that was just for the experiment specific use case. The results of these tests are not suitable for general publication due to their specific nature. Around 66GB of computer memory was used for decomposition/reconstruction and for preview of results. During the experiment more than 210 meshes were constructed and more than 180 simulations were run in order to obtain proper results. These kinds of figures can be achieved only applying larger number of nodes with high-bandwidth and low latency interconnection like in HPC. Standard cluster can't match these needs.

4.1.7 Feedback to Work Package 2

The core service that is to be deployed with respect to the workflow of the 401 experiment is the ability for remote visualization capabilities.

The simulation produce a substantial amount of data, transferring this data between the End User and the Infrastructure location is very time consuming and irrelevant since the intermediate results that are obtained are relevant only for the interim adjusting of the simulation parameters. In order to do this adjustments the results have to be visualized. That is why the most appropriate approach is to have a visualization service running on the remote machine and communication with the End User's local computer.

Apart from the visualization service the 401 experiment demanded different versions of the same software due to the fact that not all of the needed functionalities were present in the newer version of the software. For the majority of the experiment the newest version of OpenFoam was used and an older version of the software was used to do a part of the simulation.

Other services that are used for the experiment is only the remote connection of the cluster and the data transfer service.

4.1.8 Feedback to Work Package 3

The technical features that have to be enabled and deployed to fit the demands of the 401 experiment are the ones described in the section above. To resume:

- visualization (we propose NX client)

- SW versioning and version selection

- SSH

- data transfer.

Furthermore once the platform is operational we propose also to implement a set of accounting features. The parameters we expect here are: price per unit of measure (preferably CPUh), number of CPUh remaining (out of the amount that has been purchased), usage backlog, storage usage and accounting.

4.1.9 Input to Work Package 8

As already described in section 15 the usage based pricing model is the most preferred one, but for it to be applicable there must be a clear and understandable method of keeping track of the available or spent resources.

4.1.10 Input to Work Package 9

The experiment has used an Open Source SW solution (OpenFoam) for the execution of the experiment. The reason for this is that other commercial alternatives of SW solutions either do not offer the desired functionalities or are simply too expensive to justify the investment into purchase and training.

As in the infrastructure usage also in the field of software licensing a flexible, usage based pricing model is the preferred option If It can be applied.

4.1.11 Any further input to Work Package 10

The training of HPC users is a fundamental activity. We took part at an OpenFoam workshop during the Fortissimo work. The experience exchanged with other users and lecturers is of help when preparing our own simulations.

4.2 Experiment 402: Steel Casting

4.2.1 Overview of the experiment

The pouring of liquid steel from ladle to tundish is a most critical stage of the continuous casting process. In order to prevent surface oxidation, the molten metal in the ladle is covered by a layer of protective slag, which may be partly transferred into the tundish during the final stage of ladle emptying. Slag transfer into the casting lines may produce the breaking of the solid skin of the solidifying cast product, an event which may cause severe economic damage to the steel plant, while being highly hazardous to the operators. Quantitative knowledge of the fluid dynamics of the ladle emptying mechanism is required in order to develop an effective slag detection system based on vibration analysis. Due to the complexity of the physical phenomena involved, dedicated numerical simulations imply an exceptional computational load, thus requiring the use of HPC. Participation in the Fortissimo Experiment 402 - "Cloud-based Simulations of Continuous Casting" provided Ergolines with the necessary HPC resources needed to address this scientific and industrial challenge. Dedicated numerical simulations have been performed which allowed for accurate modelling of ladle emptying. The fluid dynamics of liquid steel and slag has been studied by considering both a single-phase and a double-phase configuration. Realization of a dedicated laboratory scale model allowed for successive validation of the numerical results. Excellent agreement between numerical predictions and experimental results demonstrated the effectiveness of the simulations, which enabled a key understanding of the correlation between the measured vibrational signal and the fluid dynamics of the system. These very relevant results constitute the basis for the future development of a dedicated system for ladle slag detection, an innovative technology featuring a high commercial potential in the field of continuous casting.

4.2.2 The industrial setting

In the field of continuous casting there is an increasing industrial demand for the development of new technologies for preventing slag transfer from ladle to tundish. Such an event may in fact cause a breakout, i.e. the breaking of the solid skin of the solidifying cast products, which results in hazardous dispersion of liquid steel within the industrial plant. Ladle slag monitoring is currently performed by operators on an empirical basis. Given the relevance of both safety and economic implications of such an event, there is a significant demand for an effective, automated system for ladle slag control. While passing through the ladle shroud, liquid slag induces characteristic vibrations which can be measured. In order to develop an effective detection system, it is necessary to correlate the vibrational signal to the fluid dynamics of the system. Dedicated HPC-based simulations followed by experimental validation provided Ergolines with key insight into the physics of the system. The results obtained will constitute the basis for the development of an innovative slag monitoring technology, which would significantly contribute to increase both occupational safety and steel plants productivity.

4.2.3 The improvements in the industrial process through the use of an HPC Cloud

Dedicated numerical simulations based on cloud HPC resources enabled Ergolines' researchers to perform a thorough analysis of the different ladle emptying mechanisms. As a result, it has been possible to establish a correlation between the shroud vibrational signal and the fluid dynamic of the system. Based on this achievement, it will be possible to develop a novel technology for ladle slag detection based on vibrational analysis. By preventing slag from passing from ladle to tundish, this innovative system would reduce the risk of breakouts, thus addressing key occupational safety issues and contributing to increase steel plants productivity. In fact, in order to prevent slag transfer into the tundish, the pouring of liquid steel must be interrupted before the ladle is completely empty. The ability to detect the slag while it is passing through the shroud would enable the steel plant to delay the closing of the ladle, with an estimated saving of 1% of the production. In order to produce 1 million tons of steel per year, a medium size factory must sustain steel costs up to 300 M€. Envisioned annual savings for a medium size factory are of the order of 3 M€ for low-medium quality steel and may rise significantly for higher quality cast products. Fast return on investment is thus guaranteed, together with full coverage of the overall simulation costs. In fact, if no access to HPC resources and services had been granted, the overall simulation costs, comprehensive of hardware and personnel, would have been of the order of 130.000 €.

4.2.4 The cost benefits of using an HPC Cloud

Given the complexity of the phenomenon to be simulated, a very fine discretization in terms of geometry and time is needed in order to obtain accurate results. Such a fine discretization involves a significant computational load and therefore requires adequate computational capabilities. As the company does not possess the necessary computational infrastructure, the possibility to use cloud-based HPC resources proved fundamental in addressing this specific industrial and scientific challenge. In fact, participation in the Fortissimo experiment allowed Ergolines to exploit supercomputing resources and reduced computational times without having to sustain the relevant managing costs of a dedicated infrastructure.

4.2.5 Any Other Business benefits

Besides providing significant business benefits for the company, the use of an HPC cloud in the present experiment will also bring relevant economic advantages to the steel casting industry. In fact, HPC-based fluid dynamic simulations enabled Ergolines to acquire the key scientific competencies needed for the development of an innovative system for ladle slag control, thus significantly reducing the time to market and improving product design. By reducing the risk of breakouts, this novel technology would in turn bring relevant economic benefits to the continuous casting industry, while contributing to increase steel plants productivity: the ability to detect the slag while it is passing through the shroud would in fact enable the steel plant to advance the closing of the ladle, saving up to 1 ton of steel for each casting sequence, with an estimated saving of 1% of the production. In order to produce 1 million tons of steel per year, a medium size factory must sustain steel costs up to 300 M€. Envisioned annual savings for a medium size factory are of the order of 3 M€ for low-medium quality steel and may rise significantly for higher quality cast products.

4.2.6 HPC as opposed to standard cluster offerings

Simulating complex physical phenomena such as the fluid dynamics of ladle emptying mechanisms requires a relevant computational load that can not be managed by the current company's infrastructure. The possibility to use cloud HPC resources offers the unique advantage of exploiting high performance computational capabilities together with domain specific expertise without having to sustain the managing costs of standard cluster offerings.

The scalability of the system has been analysed, namely the improvement of calculation times resulting from parallelising the simulation on an increasing number of cores. Such analysis was performed by considering a mesh with 600k elements.

The speed-up for the nth simulation has been calculated by dividing the time used for the nth simulation by the time used for the simulation with the lowest number of cores. Therefore, if the calculation time is halved, the speed-up equals 2.

In the chart below we can see how the speed-up curve detaches from linearity: this is due to the achievement of the maximum obtainable results in parallelisation with this number of elements. In fact by subdividing the calculation domain in an ever increasing number of sub-domains, the number of elements per core would be too low to benefit from the use of greater calculation capacities. From the analysis of this curve we can then infer the optimal number of cores for a certain spatial discretisation of the problem with the aim of optimising calculation resources and times.

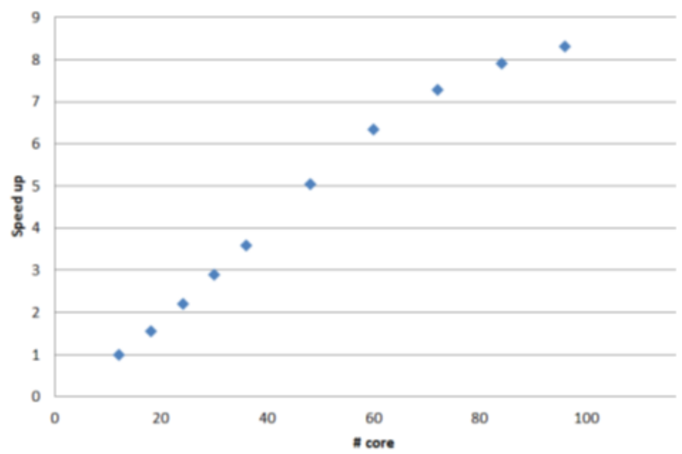

Speed-up 600k mesh

4.2.7 Feedback to Work Package 2

The experiment has identified the following requirements:

- need for interface offering to access appropriate infrastructure (both HW and SW)

- clear accounting and measurement of usage of resources

- introduction to HPC workflows is welcome (how to deploy jobs, MPI coding…)

4.2.8 Feedback to Work Package 3

The technical work should resolve the requirements from the paragraph above.

4.2.9 Input to Work Package 8

The pay-per-use option is the most appropriate usage model for this experiment. The shift form CAPEX to OPEX is clear as described in section 4. The business benefits for industrial end users (customers of the experiment SME) are also evident as described in sections 3 and 5.

4.2.10 Input to Work Package 9

Experiment 402 has used an OpenSource and free software solution. This is due to the fact that comparable licensed solutions would represent a major cost in the presented workflow if they were to be implemented.

If we want to even consider using a commercial software solution the ISV should also apply a pay-per-use licensing model with rates that are flexible enough. This licensing model should also be supported by clear performance benefits of using such software.

4.2.11 Any further input to Work Package 10

None at this point.

4.3 Experiment 403: Design of High-Pressure Vessels

4.3.1 Overview of the experiment

The objective of experiment No.3 (WP 403 Micro-Macro mechanical properties of high pressure vessels) is to improve end user capacity to satisfy the principal ISO 11439 standard and the ECE R 110 normative for Gas cylinders — High pressure cylinders for on-board storage of natural gas used as a fuel for automotive vehicles. In order to achieve this objective, the calculations for basic micro-macro mechanical properties of a composite material were implemented in GNU Octave. For the design of the composite used for production of high pressure vessels, parallel optimization is used. This is an innovative idea in the field of composite design and entirely new field within the scientific research MIKROSAM does, already improving the performance of the filament winding solutions it offers. The benefits of this experiment are fostered by the replicability of the parallel optimization outcomes in the aviation, space, medical, and other industries that increasingly substitute metal with composite high pressure vessels.
In this experiment, HPC Provider and Host Center roles are undertaken by ARCTUR and HPC Expert consultancy is provided by XLAB.

4.3.2 The industrial setting

Composites as multiphase materials offer the possibility to influence their properties or add new functionalities by a proper choice and combination of different phases. The biggest advantage of modern composite materials is that they are times lighter, and yet as strong or stronger compared to widely used metals. By choosing an appropriate combination of matrix and reinforcement material, a new material can be made that exactly meets the requirements of a particular application. Composites also provide design flexibility because many of them can be molded into complex shapes.

Filament winding has become a popular construction technique in a wide variety of industries for creating composite structures with high stiffness-to-weight ratios. Filament-wound composite pressure vessels utilizing high strength and high modulus to density ratio materials offer significant weight savings over conventional all-metal pressure vessels for the containment of high pressure gases and fluids. Among all materials, composite materials have the potential to replace widely used steel and aluminum, and many times with better performance. Composite components can save in component weight 60% to 80% by replacing steel, and 20% to 50% by replacing aluminum parts. Filament winding machines Mikrosam provides are custom-made and vary dependent on the final application (product) they are meant for. Regarding fossil fuel environmental concerns and optimization of their storage and transportation, composite high-pressure vessels gain popularity.

Composite design is an exhausting and very time consuming process. Experiment 403 “Mikro-Makro mechanical properties of high pressure vessels” targets and aims to optimize the design of the composite laminate structure. As a manufacturer of high-tech computer controlled machines for production of composite materials and parts, Mikrosam faces the need to test each of the final products its equipment is designed to produce, to verify the quality that is guaranteed to the customers and end users. Launching the simulation model subject to this experiment on an HPC infrastructure, reduces the usual time to design a composite with most effective elastic constants, for about 30%. With the technical advances subject to this experiment under the FF project, Mikrosam acquired competitive advantage offering potential customer accelerated speed in designing the composite of the final product.

4.3.3 The improvements in the industrial process through the use of an HPC Cloud

Composite pressure vessels should take full advantage of the extremely high tensile strength and high elastic modulus of the fibers from which they are made. One approach to model composite material behavior is micro-mechanics, where interaction of constituent materials is examined in details as part of the definition of the behavior of heterogeneous composite material. Working with base of materials leads us to more theoretical tests, decreasing the design and testing costs (less material used) by 10%. Theoretical tests have shown that using parallel computation on HPC system can reduce time for composite design for about 30%. The use of HPC resources allows more theoretical calculations to be made and reduces the amount of actual 'physical' testing, which has to be done when it comes to the design and use of composite materials. Improvement of composite laminate design timing would add 10% to 20% (depending on the previous experience in designing composite pressure vessels) to the positive evaluation potential customers give to our complete solution offer.

4.3.4 The cost benefits of using an HPC Cloud

For each filament winding machine for production of type IV CNG cylinders we need to anticipate different materials' combination and winding angles' behavior. With the algorithm that is subject to this experiment, we get the opportunity to accelerate the time of making complex computation of a given database of materials and different winding angles effect on the final product, to theoretically test the design upfront. This means that all costs incurred from the actual physical design testing are directly related to the theoretical design cost savings. Considering that the price for production of one unit is €250, and one filament winding machine requires on average 10 units to be tested to validate machine functionality and final product design, the total material price is €2500 and labor cost is €1459, for a machine manufacturer. These numbers and above mentioned 10% saving on material production and 20% saving on labor, lead us to €542 saving in order to physically assess the design of the final product and so to test the machine. On average the calculation of one single vessel costs about 100€. Thus meaning that the overall saving of 442€. Design cost savings for a machine buyer in actual production conditions multiply in accordance with the need for new product design testing, product design improvement and regular product design properties' and machine functionality examination.

4.3.5 Any Other Business benefits

Being actively involved in a research project to improve our technical and technological competence around which we build our equipment and software solutions, adds to the value and recognition as a company that continuously invests resources in enhancing the knowledge and experience for the benefit of our customers and their product end users.

Fortissimo enabled us to use HPC resources for the first time, and this successful application initiated our consideration to further apply them to other complex time-consuming calculations. This experiment experience encouraged our company to use self-service model to solve every day technical problems and advance our overall machine design and manufacturing capacity.

Using optimization in composite design in order to simulate micro-macro mechanical properties of filament wound high pressure vessel is an innovative idea which will reflect in improved services and together with that in higher competitiveness of our company.

4.3.6 HPC as opposed to standard cluster offerings

Given that we, as an SME have no experience in the use of standard cluster offerings we are not competent to make any comparison. However, if we need to choose between HPC and standard cluster, the trust and confidence that we can get from the HPC provider would be hard to achieve when standard cluster is used. Also, the direct communication that we had with our HPC provider during this experiment proved to be very useful in order to set up everything as needed and properly use the resources. The scalability of the code has been mostly addressed by the partner XLAB that managed to scale the code from running on a single CPU to running a much bigger and detailed case using the MPI parallelization method running over multiple computing nodes. While the overall scalability of the code cannot be expressed just by comparing the amount of CPU resources used we can state that the overall scale of the model and complexity have improved considerably. The amount of simulated elements has by about a factor of 4 without major increases in times to result. Therefore we can state that the achieved scalability of the case is at least 3-fold. Due to the specifics of the code this speed up probably couldn't be achieved on a standard public cluster since the cases need a tightly connected infrastructure, that cannot be easily obtained on a public cluster.

During the experiment we have developed a composite design simulation algorithm that can be replicated each time we need to update our database of materials, add another constraint or add another objective function. Super computing power and an HPC provider that will adapt the working environment to our needs (enabling Octave for the HPC usage and 'mpi' package for the parallel computations) are inevitable in order to fully exploit the advantages of this algorithm. Direct communication with the HPC provider is necessary, but unavailable when standard cluster offerings are used.

4.3.7 Feedback to Work Package 2

The end user provided the technical characteristics of the cluster that were needed for the experiment and the SW configuration that was implemented.

4.3.8 Feedback to Work Package 3

Suggestions for the envisaged functionalities of the system were provided based on the technical specifications delivered to WP2.

4.3.9 Input to Work Package 8

Flexible environment, user friendly interface and trial period for the use of services offered by the marketplace are very important in order to continue using HPC infrastructure. Easy and fast integration of a new job that needs to be done using super computing power would be very helpful for future use of HPC. Continuous support by forum and training would help us to submit new computational large processes that relay on super computing power.

4.3.10 Input to Work Package 9

The open-source software GNU Octave was used for the simulations of micro-macro mechanical properties and all optimization processes. No ISVs were involved in this experiment. Therefore, we can't present any relevant business models for ISVs.

4.3.11 Any further input to Work Package 10

In order to use all available resources from cloud based HPC infrastructure we need to have knowledge of parallel processing, so we think that training for parallelization of the processes would be very helpful for successful use of the performances. Also, open source software can be very useful, but this software is always very purely documented, so the user needs to put much time and effort in order to use it. The training and having available experts for the use of particular open source software in order to clear the issues for proper use and expected results will help to disseminate the usage of cloud based infrastructure.

One of the basic things when choosing the market in order to buy some service is success story of the market and positive experience of the buyers, so direct communication with the users (if they are available) and reading their success stories would encourage other users to use particular service from the same market and of course leave their story.

Almost always the end users look for privacy of the work (or product). Putting some job on cloud makes it public, so the level of guarantee for the privacy of the job is always needed.

4.4 Experiment 404: Simulation of Tightening Flanges

4.4.1 Overview of the experiment

The objective of this experiment is to test the feasibility of using cloud-based HPC simulation environment for the design condition of specialised mechanical components (flanges) which are used in complex structures and subject to complex loading.

Flanges are joint elements between metallic parts which are used by many manufacturing companies in various sectors such as oil and gas, chemical, pharmaceutical, etc. The Fortissimo experiment 404 works in a particular case which considers flanged joints that are used for high pressure and/or temperature gases or liquid piping applications (Figure 1). Due to sealing requirements for these applications, between the flanges a gasket is placed inside a groove (Figure 2) which is located on both flanges. During the process of closing and tightening the joint (Figure 3), the gasket and/or the flanges may suffer deformations and high levels of stress which, if higher than the materials elastic limit, may damage the flange so it cannot be reused once it has been reopened (i.e., for maintenance) or may result in sealing problems, leakages or a decrease in the joint's operational life. Even worse, if the closing process is done inadequately, the production plant could suffer fatal accidents which halt operation or may result in serious damage to the workers, the surrounding community, or the environment.

The experiment develops a model for simulating the tightening process and a front-end application to control the parametric simulations, which will permit to improve the tightening process design and comprehension for each case using the Fortissimo infrastructure.

Figure 1: Example of some flanged joints (those with bolts and nuts).

Figure 2: Image of two flanges and a gasket (in the centre)

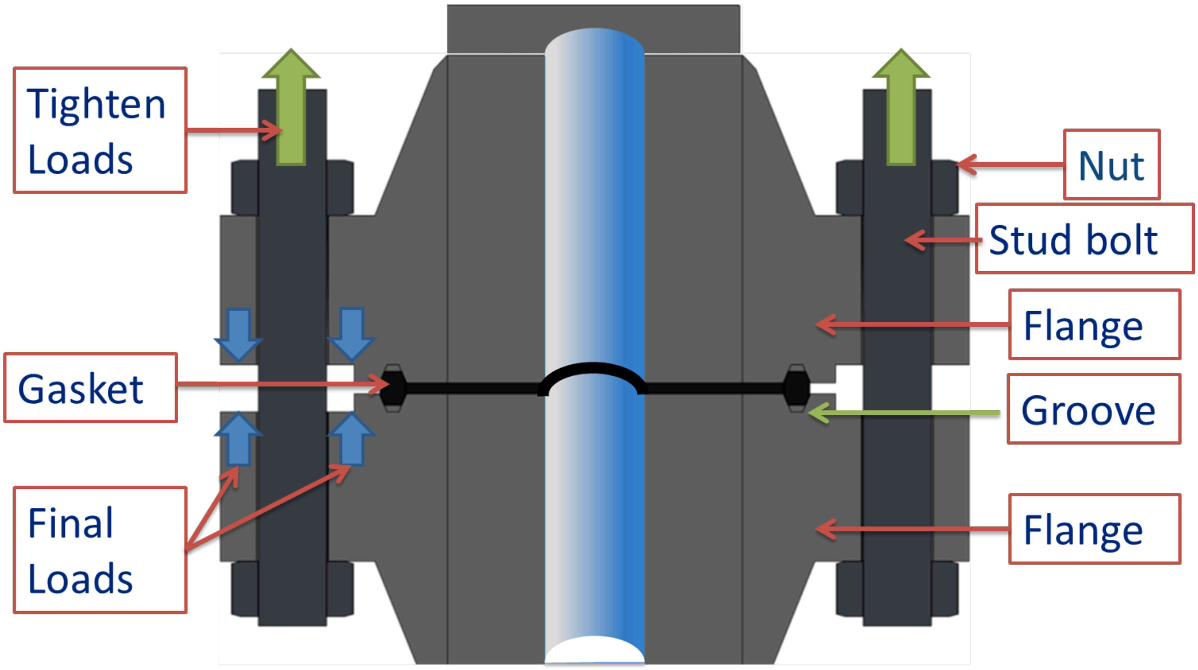

Figure 3: Outline of a flanged joint, its elements, and sealing process.

4.4.2 The industrial setting

Texas Controls is a Spanish SME that offers large industrial facilities technological products and services oriented to the instrumentation, tightening and sealing solutions. In the field of tightening and sealing processes, today Texas Controls is the sole provider of engineering solutions in this field. As such, one of our priorities is to provide the highest quality solutions.

Texas Controls offers these solutions mainly to the industrial and energy sectors, and to the R&D environment. These solutions are especially important to customers of gas and oil industry. In these sectors, probably more than in any other, it is imperative to avoid leaks in pipes and pressure vessels or reactors that could occur under extreme pressure and temperature conditions. The costs of preventing leakage are much lower than those generated by its consequences. However, many variables must consider in order to effectively prevent hazards: the sealing process involves plastic interactions between elements so an expected deformation may occur in certain areas. The response and the service life of each and all of the components will be determined by the extreme and varying pressure and temperature conditions. This is certainly a complex and fascinating engineering problem to be solved. Therefore, it is crucial to study and predict the behaviour of flanged joints, to understand how the elastic interactions between elements are produced.

This experiment develops a model for simulating the tightening process and a front-end application to control the parametric simulations. In this kind of simulation, the tightening process usually is not taken into account. During the tightening process damage to parts such as bolts, flanges, etc…could be inflicted. Any damaged suffered by any or all of this parts could be fatal. It is in this very process that final sealing conditions are established. In a fundamental way, optimal tightening depends on the initial loads and on the defined tightening strategy. Thus, it is of the utmost importance to minimise the load applied to bolts and flanges as well as to optimize the tightening sequence when critical applications such as the described above are involved. The simulation of each tightening case needs several executions depending on the number of different parameters and the number of levels for each parameter.

Industrial clients for our solutions often require immediate response times for maintenance and/or critical incident intervention since the infrastructures affected by tightening and sealing technologies regularly play a major role in providing energy to large populations as well as other industries. Thus, to run all the parametric cases simultaneously in order to get the urgently required results as soon as possible is therefore the best possible solution. This will allow Texas Controls to improve tightening processes design and gain comprehension and insight for each case using the Fortissimo infrastructure.

4.4.3 The improvements in the industrial process through the use of an HPC Cloud

Industrial clients for our solutions often require immediate response times for maintenance and/or critical incident intervention since the infrastructures affected by tightening and sealing technologies regularly play a major role in providing energy to large populations as well as other industries. Thus, to run all the parametric cases simultaneously in order to get the urgently required results as soon as possible is therefore the best possible solution.

From a financial standpoint, it is indeed unthinkable for a small engineering SME such as Texas Controls to make the regular investments needed to maintain an operational and updated HPC cluster to provide the necessary simulating infrastructure to respond to its unusually high yet sporadic demands.

The Fortissimo infrastructure will allow Texas Controls to improve tightening processes design and gain comprehension and insight for each case. The Fortissimo infrastructure provides the best possible technological solution available to satisfy our specific needs today. Not only for the highly scalable computing resources it provides, but also for the knowledgeable human capital accrued by the application experts available to us whenever needed.

With the simulation model, developed in the experiment, Texas Controls shall be able to conduct analysis before we start actual tightening process. Thus, time and money will be saved at the same time that damage to components will be avoided, and greater safety for facilities, workers, the community and the environment will be achieved.

4.4.4 The cost benefits of using an HPC Cloud

A key benefit of using the Fortissimo HPC Cloud will be the reduction of operational time in the maintenance of huge and critical oil and gas heat exchangers or reactors. Another fundamental benefit will be avoiding damages in flanges during tightening processes. As an example of the benefits of the solution, a real case can be presented (Figure 4) for a WELD NECK FLANGE 60" 900 lbs. The technical data of this case is shown in Table 1.

Based on previous field work experience (of a time past when we did not rely on simulations), a non-optimised tightening process on a 24 stud bolt flange consumed 108 man-hours while a current optimised tightening process consumes 72 man-hour. The optimisation allows the technicians to reduce the number of tightening passes from 6 to 4, in all this represents a 33% time saving per flange. Not to belabour the point, this obviously represents considerable savings in labour costs. Such savings provide a competitive edge for our company in international tenders.

With the simulation model, developed in the experiment, Texas Controls shall be able to conduct analysis before we start actual tightening process. Thus, time and money will be saved at the same time that damage to components will be avoided, and greater safety for facilities, workers, the community and the environment will be achieved.

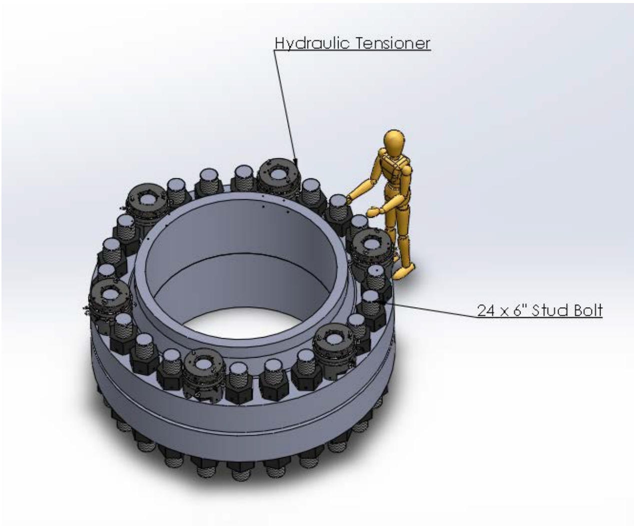

Figure 4: 24 bolts simple case WELD NECK FLANGE 60" 900 lbs (horizontal position)

Design parameter	Value of the parameter
STUD BOLTS	24 5 ½" A193 B7M
DESIGN PRESSURE	157,8 bar
DESIGN TEMPERATURE.	• 85ºC
HYDRAULIC TEST PRESSURE	238,6 bar
GASKET	RTJ SS
TENSIONER WEIGHT	155 kg.
MAX LOAD	7050 KN

Table 1: Design parameters of the case

4.4.5 Any Other Business benefits

The most important short-term benefits are qualitative rather than quantitative but in the mid-term they will have a quantifiable economic impact.

Relevant improvements are the reduction of the uncertainties of the process of tightening and the increasing safety of the final result of the process, as we could guarantee that in any step of the process we are over-stressing neither the flange nor the gasket. These factors results in reducing the maintenance cost of the facilities in long-term. . The cost of a leak is much higher than the cost of avoiding it.

Here is a list of extra benefits we expect from our participation in this experiment.

To acquire Scientific & Technical knowledge about a key business activity.

To further improve and optimize the design of tightening tools and procedures.

To create new revenue cash flows and to generate new services for new and existing customers.

To turn a "based-on-experience" activity into a "knowledge-based" activity .

To improve the technical background of the company engineers.

To gain access to computing technologies and reduce the "learning curve" for engineers.

To sound out new business lines: Offer FEM on HPC to third parties.

To live up to an image of a company fully compromised with R&D.

To grow and stress differentiating factors to set our company apart from competitors.

To collaborate with other organisations, not involved in our line of business, but important to improve our knowledge.

4.4.6 HPC as opposed to standard cluster offerings

Tightening work is by no means a regularly scheduled activity. Peaks and valleys in workloads are common. Tightening procedures in the same installation can involve several different tightening processes for flanges of different types. In many cases it will be necessary to perform many simulations in a short period of time, however there could be fairly long periods in which no simulation may be required. HPC cloud offers fast response and the number of resources required in each case.
Industrial clients for our solutions often require immediate response times for maintenance and/or critical incident intervention since the infrastructures affected by tightening and sealing technologies regularly play a major role in providing energy to large populations as well as other industries. Thus, to run all the parametric cases simultaneously in order to get the urgently required results as soon as possible is therefore the best possible solution.

From a financial standpoint, it is indeed unthinkable for a small engineering SME such as Texas Controls to make the regular investments needed to maintain an operational and updated HPC cluster to provide the necessary simulating infrastructure to respond to its unusually high yet sporadic demands.

The FEM model for this case was initially designed to be executed using the distributed computing capability offered by the MUMPS solver. After obtaining the final model and performing the scalability tests, the results indicated that there is no significant benefit of using parallelization for this model. This was because the calculation of contact areas does not benefit from parallelization and in this model the contacts affect most of the computational domain. However several simultaneously simulations (parametric jobs) are needed to obtain the best tighten strategy and these simulations are intensive in Ram usage, especially if the option IN_CORE is activated. This option avoids the data dumping to disk in order to improve the computing time.

4.4.7 Feedback to Work Package 2

Complete as indicated above.

4.4.8 Feedback to Work Package 3

Complete as indicated above.

4.4.9 Input to Work Package 8

The model can be used later to develop new services in the company. For example, as support for the training courses in Texas Controls.

4.4.10 Input to Work Package 9

None. The software is open source.

4.4.11 Any further input to Work Package 10

Complete as indicated above.

4.5 Experiment 405: Design of High-Voltage Cables

4.5.1 Overview of the experiment

The objective of experiment 405 "Cloud-based multiphysics simulation" is to showcase the use of cloud based HPC services to investigate multiphysics simulations using commercial and open-source software packages. The participants in WP 405 are Prysmian, world leader in the energy and telecom cables and systems industry, and Cineca, as service provider.

The end-user company is experiencing the need to scale-up the current simulations of energy cables and systems to a larger refinement and size, exceeding the computing power available to them. In addition they would like to enlarge the limits of what they are simulating, requiring new competences and tools. Demonstration of a cost-effective cloud-based HPC simulation would represent a breakthrough for this end-user.

The end-user is using COMSOL Multiphysics as a tool for their electromagnetics simulations and it will upgrade its licenses for the scope of the experiment. COMSOL Multiphysics is a modelling and simulation tool that allows engineers to accurately model systems in which one or even more than one physics are involved at the same time. The experiment will provide the end-user with a workflow integrating COMSOL software with pre/post processing visualization and open-source tools, such as OpenFOAM, in a cloud environment, together with an insight into how this could be exploited as a future business model.

The experiment will address porting and testing commercial and open-source software on the Fortissimo cloud-based HPC infrastructure, and integrating them in a workflow customized on the customer specific needs and habits. This setup addresses not only increased competiveness by providing affordable computing resources through a cloud-like business approach, but also the reduction of the time-to-market exploiting the expertise of HPC centres to build customizable high-productivity and high performance set-ups.

4.5.2 The industrial setting

High voltage cables for power transmission are source of electromagnetic losses during operation, when load electric current flows along the conductor. Losses can be divided into well-defined categories: joule losses inside the conductor, due to electrical resistivity of conductor's material (copper or aluminium), dielectric losses inside the electric insulation layer, losses in additional metallic layers (screen, mechanical armour) generated by hysteresis or currents induced by the magnetic field from the load current.

Heat generated inside cable's elements is then dissipated into ambient, and temperature distribution establishes inside cables as a results of equilibrium between generation and dissipation.

During steady-state normal operation conductor temperature is not allowed to exceed a prescribed limit, to not speed ageing of electric insulation and reduce cable operating life. Because most of the losses are directly affected by load current, this limitation raises the problem of determining the rating of a cable system, i.e. the intensity of the electric current that the system can sustain on a continuous basis without overheating.

From above discussion rating calculation then involves:

- Correct computation of all loss contributions inside the cable;
- Precise modelling of heat transfer phenomena with surrounding environment.

So far modelling and calculation activities have relayed on formulas and procedures mostly based on analytical or semi empirical models that have the advantage of being easily implementable into spread sheets or self-made routines, but in many cases are affected by a high level of approximation which may result in excessive overdesign of the systems.

The increasing needs of optimization of cable systems performances together with challenging installation and operation conditions demands now more evolved design methodologies. FEM-based advanced simulations together with the availability of user-friendly and powerful computation infrastructures offer provide designers with powerful tools to meet stringent requirements of future high voltage cable transmission systems.

4.5.3 The improvements in the industrial process through the use of an HPC Cloud

The use of Fortissimo HPC Cloud helps Prysmian to:

1. Improve calculation of loss contributions inside an energy cable, especially those induced by magnetic field from load current;
2. Model heat transfer form cables' surface to surrounding environment.

Point 1 can be addressed through implementation of an FEM electromagnetic model to compute magnetic field distribution from load current and its interaction with metallic structures surrounding the conductor.

Point 2 is a complex task especially when cables are installed inside tunnels with natural ventilation only (most of the applications). Proper cables rating calculation includes the detailed description of heat transfer mechanisms from the cable surface to the air and from the air to the tunnel wall. These mechanisms involve: heat conduction across the boundary layers adjacent to the cable surface and tunnel wall, air convection and buoyant movements, radiation between cables and walls and cables themselves.

CFD computer analysis can provide a great deal of information about these phenomena by solving Navier-Stokes equations and nonlinear heat transfer equation associated to radiative phenomena.

In conclusion advanced FEM modelling combined with HPC Cloud computing offers the opportunity to address in detail thermal and fluid dynamics phenomena that play a critical role in the design and operation of critical high voltage cable transmission systems.

4.5.4 The cost benefits of using an HPC Cloud

Prysmian moved to a design workflow for their products completely based on CAE simulations about 7 years ago already. At the time it was estimated that this would have been provided a ROI within two years, considering effort and money spent on producing and testing physical prototypes. It has to be noted however that the cost of CAE simulations is a small fraction (some %) of the total engineering costs of a typical Prysmian project (e.g. setting up a new inter-Atlantic communication cable).

Tests performed during this experiments have now shown that the use of HPC-based simulation using a Cloud and external expertise will result in considerable savings with respect of acquiring internal resources, in a three years framework, as described later in more detail.

It has to be stressed that the outcome of this experiment represent an enabling step for Prysmian to deal with new kind of simulation and services. Therefore, this can give Prysmian using such simulation a significant commercial advantage, but it makes difficult to quantify ROIs.

HPC-based simulation has been used to address optimization design of high voltage cables with a prototype service combining FEM simulations of new design and CFD simulations of cables in their installation environmental conditions, and adequate computing resources, which were not previously available to the company.

Considering the CFD component, this was a completely new issue for the company, no competencies were previously available. It was estimated that implementing a setup as provided by Cineca in the framework of the experiment, either hiring a consultant, or training internal personnel, would have implied an investment of at least €50K. Since competence in Open Source CAE applications is still rare in the market, it would have been also likely that the setup would have been built using commercial ISV applications, meaning a further cost of about €10K/year of software licenses.

Considering the FEM component, with current Prysmian IT infrastructure for design, consisting in a few high-level workstations, it is impossible to run 3D simulations, since a single run would last weeks. In some extreme cases the size of the FEM model requires an amount of RAM memory that exceeds those available in Prysmian's hardware. Therefore, to allow in-house 3D simulations in the Prysmian design process, it would have been necessary to invest in a HPC cluster. We estimated a suitable system to cost €150K in hardware investment and about €40K/year in operational costs, considering power, cooling and maintenance. In a 3 years operational lifecycle, this amounts to a total €300K. At the current average fees of HPC clouds services this equates to at least 3 million CPU hours if considering an alternative cloud solution. The benchmark test made in this experiment allowed Prysmian to have a more precise idea of computational needs for their design process, allowing estimating that using HPC cloud services instead of investing in an in-house system savings of around 30% are possible.

4.5.5 Any Other Business benefits

The provision of a Cloud-based HPC solution would offer significant innovation into the marketplace. In particular the innovation would be derived from the possibility for the end-user to tackle simulations at an unprecedented refinement level. This offer benefits to Prysmian allowing them to retain the competitive advantage that allowed them to be the market leader, with a time to market that does not affect the increase in complexity of the design process.

In summary, customer value and benefits from such HPC-cloud CFD simulation could be summarised as follows:

- Increased speed, as a result of obtaining simulation results within hours instead of weeks.
- Increased flexibility, occurring mainly through services utilization on demand and payment out of specific project budgets.
- Reduced cost, encompassing pay on demand capabilities, reduced depreciation costs, and more effective labour deployment. This has been estimated in Ch. 4 as about €30K per year.
- Improved quality, as a result of the production of an enhanced solution based on higher resolution models.
- Available know-how of CFD expert for cable environment installations. This has been estimated in Ch. 4 as a saving of about €50K as a starting cost, plus €10K per year.
- Faster design processes, driven by optimization made possible by increased CFD simulation throughput
- Increased competitive advantage on the basis of the above mentioned benefits. This means the potential edge for winning projects with €10-100M budget.
- Lower entrance barrier for company approaching CFD

4.5.6 HPC as opposed to standard cluster offerings

The need for cloud-based HPC simulation is of outmost importance. In this experiment, as by moving simulation of CFD or FEM models on a cloud based HPC environment, processing times can be impressively shortened. Moreover, with the possibility of a flexible use of the same license between a local infrastructure and remote cloud services, the anticipated royalties and utilisation costs of such services can be drastically reduced, enabling the end user company to have a measurable ROI moving from stand-alone and effort/resources consuming applications to more flexible SaaS solutions.

It is important to stress that, especially considering COMSOL simulations, the workflow described in this experiment is particularly suited for a Fortissimo-like infrastructure, since it requires computing nodes with particularly large RAM available (>256 GB) and remote visualization capabilities, not available from traditional providers like the Amazon EC2.

4.5.7 Feedback to Work Package 2

Experiment 4.05 has provided to Work Package 2, "Core Service Requirements Capture and Infrastructure Development", as described in Deliverable P405.1, detailed hardware, and software and license requirements, alongside with security and data management constraints and user management and accounting, as well user support needs for porting COMSOL software over CINECA’s HPC cloud infrastructures.

4.5.8 Feedback to Work Package 3

Experiment 4.05 has provided up to now only indirect feedback to Work Package 3, through the requirement collection document provided to Work Package 2 (Deliverable P405.1).

It is however willing to be an early tester of the marketplace middleware and ready to port the developed framework to the Fortissimo middleware as soon it will be available.

4.5.9 Input to Work Package 8

Experiment 4.05 input to Work Package 8, "Exploitation and Sustainability", will be collected in Deliverable P405.3 (due PM17) comprising a preliminary business model analysis with regard to the commercial exploitation of the experiment's service offering. The analysis will include a detailed presentation of the services to be offered, and identify all the stakeholders involved.

4.5.10 Input to Work Package 9

Experiment 4.05 exploits one of the most widely used multiphysics application in industrial contexts, COMSOL Multiphysics. COMSOL declared their interest in following the results of the experiment and to provide input for WP9 data collection activities and the annual ISV Forum.

4.5.11 Any further input to Work Package 10

Not at the moment.

4.6 Experiment 406: HPC Tools for Urban Planning

4.6.1 Overview of the experiment

Fortissimo Workpackage 406 (WP406) is titled “Experiment 6: Cloud-based HPC tools for urban planning”.

The participants in WP 406 are IES, developers of one of the world’s leading 3D building simulation software suites, The Virtual Environment (hereafter referred to as VE), and UEDIN EPCC, a leading European centre of excellence in advanced research, technology transfer and the provision of high-performance computing services to academia and industry.

The main objective will be to enable VE desktop installations and/or web-based interfaces to access the calculation resources (“engines”) hosted on an HPC cloud infrastructure. This will be achieved by:

- modifying and adding to the existing software and services, and,
- creating a structure which supports model and data transfer both from and to VE desktop installations and/or web-based interfaces.

4.6.2 The industrial setting

IES are currently developing a dynamic simulation modelling masterplanning tool for cities which will enable relevant stakeholders to assess, for example, the energy efficiency of a city, quality of living and more.

The tool can be used at any stage of a city’s life, and can be used in cities only beginning the journey towards a ‘Smart’ City or those that are well on their way towards sustainable advancement and integration with ‘Smart’ Technologies. The developed interactive decision support tool can be used by City Architects/Planners, Urban Planning Consultancy Firms, the Public and Business Community, ICT/RET system suppliers and potential investors in a City.

This tool will rely heavily on the existence of HPC cloud based simulation as a result of the potential depth of information that would be associated with multiple buildings and their interaction within the urban context.

4.6.3 The improvements in the industrial process through the use of an HPC Cloud

Masterplanning for cities, as you can imagine will require access to a huge mass of varied data; this data will in turn be simulated through IES's Virtual Environment Software on its' current desktop solution; simulating even a small scale city through this channel could take days, even weeks. An HPC cloud based solution would drastically reduce simulation time, making the process easier, simpler and more concise for anyone involved.

4.6.4 The cost benefits of using an HPC Cloud

To validate the experiment prior to launching the first version to customers we tested the upload of building model simulations of various sizes ranging from small to XXL. Figure 1 provides the timings for these simulations and some commentary on what these results mean.

Faster simulation times, compared to Windows desktop…

Suncast Samples Listed by 'Model Size'	Rooms	Surfaces	Windows	FAST Desktop (4 core)	INDY node (64 cores)	Credits
Small 1	18	162	231	1m 20s	5s	1
Small 2	408	3053	646	18m 19s	2m 50s	1
Medium 1	959	6966	916	4m 23s	43s	4
Medium 2 (Heathrow T5)	1114	9919	1284	33m 55s	7m 04s	4
Large (v. large hotel)	2304	18308	2727	8h 49m 56s	1h 4m 37s	10
XL 1 (v. large complex)	2306	19030	3077	Est. 2 days	13h 17m 30s	20
XXL 2 (large hotel complex)	4223	35471	5435	Est. 3days	18h 22m 29s	50
XXL 3 (largest ever model)	4231	35626	5538	Est. 13 days	52h 53m 08s	50

Notes: Very Large/complex sims take a long time
Long, complex surfaces & lots of openings make a huge difference.
We show the simulation queue status

Figure 5: Simulation performance

Figure 2 illustrates the largest of these simulations. This sample model would have taken 2 weeks to run the SunCast simulation through the desktop but through the Fortissimo experiment we were able to reduce simulation time to only 2 days.

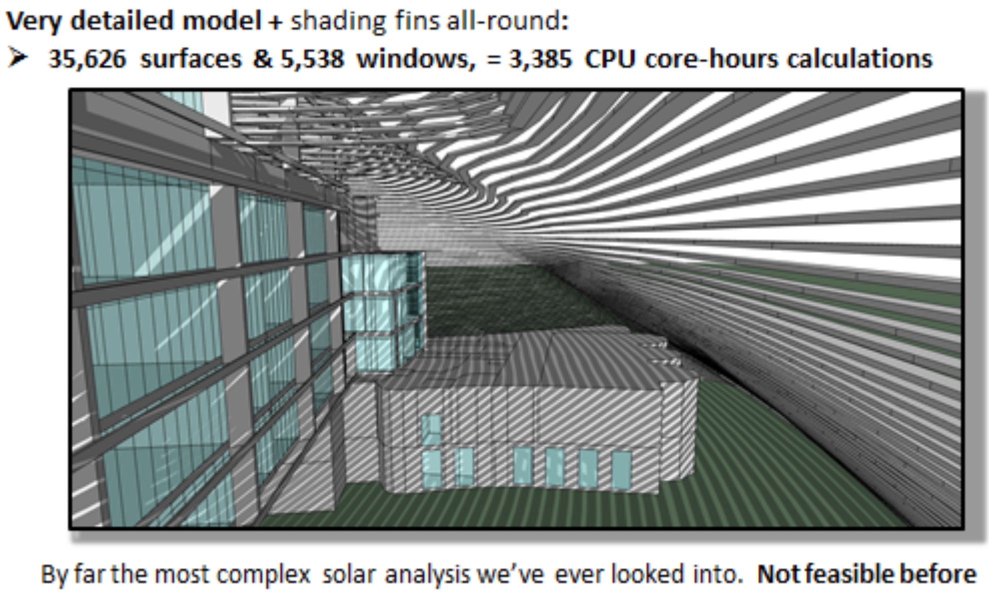

Figure 6: A previously infeasible simulation is now possible.

Based on these tests we were able to determine timing and therefore produce costs shown in Figure 3 below, which may over time change as the experiment adapts.

Credit cost (pay-as-you-go):

Cost	Credits
£40	10
£200	50
£400	110 (discount)

Credit consumption per model

Model size				Charge (credits)
Label	Surfaces		Windows	
Small	<5,000	&	<1,000	1
Medium	<10,000	&	<2,000	4
Large	<20,000	&	<3,000	10
XL	<30,000	&	<3,500	20
XXL	unlimited		unlimited	50

Figure 3: Cost Model

Since these initial tests happened, there were a number of real life projects where SunCast simulations were ran on the cloud. The system was used to run simulations on two schools in Qatar. These schools have circa 150 teaching spaces, coupled with sports facilities and ancillary rooms. Both of the designs had have lots of shading forms around the classroom façades as well as overhangs in around the main entrance areas. Looking at the schools, we were seeing a reduction from around 9 hours on a 4 core fast desktop to 1hour via the cloud. On a single core machine this could be closer to 3-4 days. So to put that in perspective the cloud calculations are bringing the analysis down into a 1 day window where simulations can actually be run daily whereas before the lead in period for the simulation would have been days to weeks. For a standard VE user that is a massive jump to be able to access that level of computing power on demand. There are the high profile companies that we all know who will have their own version of our terminal server configurations so for them the jump will not be as significant. However what it does do is allow the cloud server to offer a dedicated service and not consume their licence and their local servers' computing power.

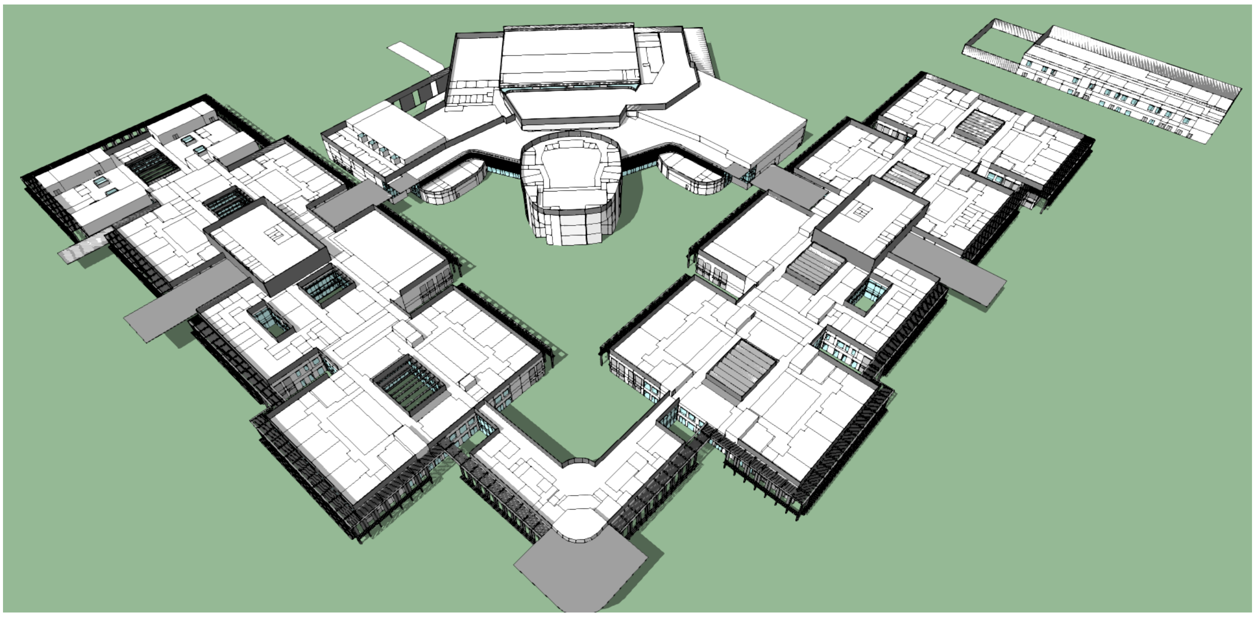

Figure 7 School design analysed in Qatar

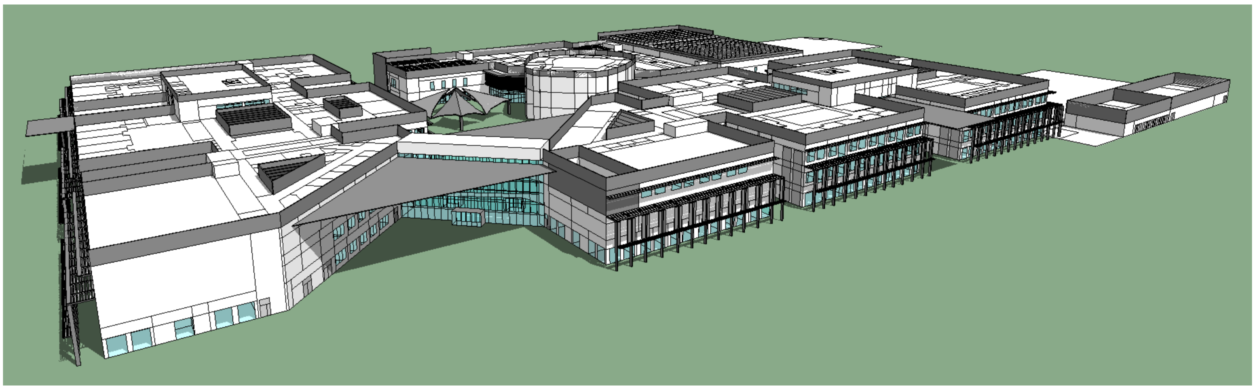

Figure 8 Second school design that was analysed

4.6.5 Any Other Business benefits

The following other business benefits have thus far been identified.

- Faster roll-out of new features to end-users.
- Less disruption to end-user when software is updated. For example, in large corporations it can take up to 6 months for desktop software updates to be implemented by corporate IT staff. With a web update this delay is minimised.
- Reducing update costs for end-users. That is, IT staff (internal and/or external) do not need to perform software updates thus reducing staffing costs. This is the same regardless of whether or not IT services are outsourced.
- End-users always have access to the latest software version.
- Support teams are better able to assist end-users since support staff are immediately able to access the user models on the server.
- A similar credits payment system can be used to control end-user payment for support.

4.6.6 HPC as opposed to standard cluster offerings

IES is concerned with time-related simulations. The nature of these integrated dynamic simulations requires that nodes are closely coupled for optimal performance. The IESVE Suncast engine has also been optimised to use MPI on HPC systems.

The savings and benefits achieved thus far in the Fortissimo experiment are not possible without access to the larger number of cores per node available on the EPCC HPC platform compared to Amazon EC2 standard offerings.

4.6.7 Feedback to Work Package 2

WP 406 have previously provided the following feedback to WP 2.

- Completion of the Aug/Sep 2013 WP 2 requirements questionnaire
- Replying to the WP 2 Complementary Requirements November 2013 request. This feedback is available on the Fortissimo wiki.
- Updates to the complementary requirements wiki response following the regular quarterly meeting of experiment partners. (eg OpenID)

The most recent feedback to WP2 concerned providing a link from the Fortissimo marketplace to the IES cloud offering via a web tile or similar mechanism. IES view the marketplace as an additional channel through which to direct traffic to their cloud offering. IES believe the Fortissimo marketplace is likely to be so varied that it is unlikely to be the usual route for IES users to access the IES cloud offering or indeed become aware of it.

4.6.8 Feedback to Work Package 3

Through the Fortissimo experiment, IES have been able to make the Suncast cloud offering in a production setting running on EPCC's Indy HPC platform. IES intend to expand this offering to include other IES software tools as well as external tools for solar intensity analyses.

4.6.9 Input to Work Package 8

IES' experience in Fortissimo experiment WP 406 has provided a template for migrating all IES' core applications to the cloud. IES envisage this enabling the creation of a sustainable platform where core simulation software is made available to end-users on a new computing ecosystem based on HPC. Currently, IES solutions are desk-top-based, with this new cloud offering, IES can deliver a more manageable, affordable and fluid solution to end-users whilst creating a brand new business model for IES itself. This has meant IES creating a new charging structure more suited to software leasing. IES view their cloud offering as being critical to their future strategy.

4.6.10 Input to Work Package 9

IES have completed and returned the Business Models questionnaire, giving their opinion based on experiences with current HPC capabilities. Other than this, this WP has not been in contact with IES and so, as yet, there has been no further interaction. This is viewed as a pity since IES have spent significant time considering business models suitable for their cloud offering and hence are interested in hearing about other approaches.

4.6.11 Any further input to Work Package 10

IES have received an email from Ulli Becker-Lemgau, the Fortissimo WP 10 leader that lists the courses from the Fortissimo HPC resource providers available to Fortissimo experiment partners. IES staff have however, already been on HPC courses at EPCC, most notably the MPI course. IES already provide extensive eTraining to their clients.

IES have produced two flyers relating to their involvement in Fortissimo. The first is concerned solely with Fortissimo and WP 406. The second is concerned more generally with R&D at IES and hence Fortissimo WP 406 is described in the context of other activities at IES. These flyers have thus far been distributed at two events in 2013 (Building Simulation Conference and BATIMAT) and a further event in 2014 (Ecobuild). The intention is for them to be distributed at a further two events in 2014 (Base London and Greenbuild).

IES have conducted an extensive global email campaign to make all their clients aware of the cloud offering being developed in WP 406 and that is now available as a production service. Further information on the numbers involved from IES is available on request. IES however are reluctant for these numbers to be made public.

IES believe that correctly identifying the Fortissimo target markets will ultimately determine its success or otherwise. Dissemination to those target markets is therefore crucial. IES will inform the Fortissimo WP 10 leader whom they believe the target markets are. IES however note that there are issues around data protection of client information that may prevent exchange of contact details.

4.7 Experiment 407: Airflow Over a Vehicle

4.7.1 Overview of the experiment

This experiment concerns the provision of cloud-based supercomputing infrastructure to enable more in-depth CFD simulations used within the context of automotive industry. OpenFOAM is utilised as the cloud-based CFD package hosted on Fortissimo platform to tackle some representative external aerodynamic problems widely used in automotive design.

4.7.2 The industrial setting

Automotive industry heavily relies on Computational Fluid Dynamics (CFD) for its design needs. Traditionally this is mainly achieved using commercial CFD software. Given the growing complexity of these simulations, nowadays concurrent parallel execution is often the only realistic way to produce results within an acceptable time frame. The commercial software licensing models, which often charge on a per core/CPU basis, can become very costly, making them beyond the reach of SMEs. For this reason, SMEs increasing move onto the alternative open-source solutions.

OpenFOAM is a well-established open-source general-purpose CFD package. It is known to have capabilities comparable to its commercial counterparts. It is widely adopted in both academic and industrial settings. And it is well supported by both dedicated commercial offerings and freely available community help. These factors make OpenFOAM a credible alternative to those commercial software packages.

4.7.3 The improvements in the industrial process through the use of an HPC Cloud

Even for existing high profile users of HPC, such as our end user Volkswagen, the use of an HPC cloud such as Fortissimo can pay off. Here the focus of the benefit is not purely on cost but of quick availability of resources. Currently Volkswagen uses a local HPC system with 28,000 cores which provides them with a peak performance of 400 TFLOPS. This system is used to calculate 15,000 jobs a day. Our HPC provider's system ARCHER has a peak performance of 1.56 PFLOPS (=1,560 TFLOPS) consisting of 72,192 cores. Thus the ARCHER cluster has, at least on paper, almost 5 times the processing power of the Volkswagen in-house solution. Thus such a system can be used to quickly increase the processing power to satisfy a short-term sudden burst of computing demand. Apart from the normal maintenance time (a few hours every fortnight), HPC clusters are almost always online. In case the end user has urgent need to use these systems, it can often be possible to make special arrangement with the HPC hosting providers to guarantee the available resource – it is common practice for HPC provides to implement job reservation system or high-priority job queues.

Looking at the SME side, the benefit is much more focussed on cost and access to HPC infrastructure. The fact that license costs often correspond to the number of cores that are used by the commercial CFD software and the fact that SMEs do not want to invest in an expensive computer cluster environment motivates the use of a cloud infrastructure in conjunction with the use of open source software. This allows SMEs to use a supercomputer environment without the huge initial investment into hardware and the constant need to upgrade. Pay-per-use in the cloud is often more cost-effective, as will be shown in section 4 and it is often carries lower risk than investing in an in-house cluster. In addition, the huge amount of data generated by million-cell simulations can be stored in a secure environment and can be shared with other SMEs in a secure way without the need of copying the data locally. The security and efficient access to the data is ensured by the cloud service provider.

4.7.4 The cost benefits of using an HPC Cloud

Typical license cost of commercial CFD packages can be thousands to tens of thousands of Euros. That is on top of hardware cost which varies on system size. Such initial investment is a great obstacle to SMEs. SMEs can instead use an HPC cloud based offer to satisfy their computational needs. In particular such offer can be cost-effective if it includes a pay-per-use charging model.

HPC platforms often offer state-of-the-art technology that is competitive -if not superior- to local clusters used in automotive industry. In addition, SMEs can rely on the software engineering expertise available at supercomputing sites to tackle software issues which are often time-consuming and costly to resolve within their own organisations, e.g. having a dedicated IT staff to maintain a local cluster can easily cost tens of thousands of Euros.

It was shown from experiment WP407 that typical simulations can scale well to 3000+ cores on the EPCC HPC facility. This is mainly due to the superior hardware settings (e.g. much better interconnect) at the HPC site. This translates into faster job turnaround, giving the end use a competitive edge.

To assess the cost benefits that can be obtained by using an cloud environment such as Fortissimo, the cost of running a hypothetical but representative job of 60 hours on 512 cores were evaluated on our host provider ARCHER at EPCC. At a cost of 10p per core hour (approximate 0.127€, based on ARCHER's standard rate for industrial users), this leads to a total cost of roughly 3900€.

As a comparison, the cost of running similar jobs on an in-house system was evaluated by scaling down the cost of using a RWTH cluster (20k cores). The 2nd column of Table 1 shows the approximate cost of running the RWTH cluster. For a smaller system with 512 cores, cost in most categories can be scaled down linearly, as shown in the 3rd column of the table, except the staff cost. It is assumed that 1.5 full-time employees are needed to maintain and administer the small system. So a annual cost of having an in-house cluster is about 220,000€. Note that this calculation only includes running costs over a year like updating hardware but does not include the upfront cost of first installing such a system which will be much higher (but can be amortized over multiple years).

Assume that an in-house system is running 24/7, then the 60 hour job in the above example costs 60/(365*24)*220000=1510€, which is cheaper than using a cloud-based HPC system. However, in practice big HPC systems are almost running at full capacity. But in-house clusters are rarely used 24/7. Assuming the in-house system is busy at 30% of time (8 hours a day), the cost of the example job would become 1510/30%=5033€ which is much higher than using HPC facilities.

So the decision if an in-house solution or a cloud based service is more cost efficient depends highly on the usage of in-house cluster, as the cost of running in-house over the year is roughly constant whether the cluster is used or not. For our example this means that the cloud based offering becomes cost efficient if utilisation of in-house system is less than 44%, or 10.7 hours a day. It is believed that for the majority of SMEs the trade-off will favour the cloud based offering, as it is quite hard to keep such a machine busy with a small user base.

Finally, cloud-based software solutions often have the advantage of user friendliness, which is particularly relevant to the current experiment – OpenFOAM is known to have a steep learning curve. On a cloud-based system, end users may rely on well-established and tested procedures (and web-based user interface when this is implemented) to set up and perform their simulations quickly, resulting in improved productivity and significant savings in training costs that can easily surpass thousands of Euros (e.g. an official OpenFOAM 2-day training by OpenCFD costs 1,200€). Setting up CFD simulations is a speciality that needs to be acquired by the end users, no matter where to run the simulations. However, Fortissimo marketplace can attract third party supplier with such specialised skills. End users can then via the marketplace pay for these value-added services.

Table Error! No sequence specified.**: Approx. Cost to Operate of the RWTH Cluster and for the example cluster with 512 cores**

	Cost @ 20k Cores (RWTH) in €	Cost @ 512 Cores (scaled from RWTH) in €
Investment Hardware	2,000,000	51,200
Maintenance	800,000	20,480
Power	1,564,000	40,038
Building (over 25yrs)	300,000	7,680
Batch System	100,000	2,560
HPC Software	50,000	1,280
Staff	720,000	97,500 (1.5 FE)
Sum	5,534,000€ / year	220,000€ / year

Cost for a 60h 512-core job:	On the in-house cluster	On supercomputer ARCHER
At 24/7 full utilization	1,510€ (in-house solution cheaper)	< 3,900€ (ARCHER)
Threshold ~ 44% utilisation	3,900€	= 3,900€ (ARCHER)
At 30% utilization	5,033€ (cloud solution cheaper)	> 3,900€ (ARCHER)

4.7.5 Any Other Business benefits

Cloud-based HPC offering can encourage collaborations among SMEs by giving them a platform to work together. Fortissimo platform can also offer standard facility to support common business needs, such as cloud-based data storage, secure mechanisms for data sharing among users, and additional tools for data analysis and visualisation.

4.7.6 HPC as opposed to standard cluster offerings

There are a number of reasons for the SMEs to choose cloud-based HPC providers over in-house solutions or standard offering (such as Amazon) to satisfy their computing needs

HPC platforms have larger capacity to carry out big simulations which may not be feasible on smaller clusters. Even small clusters are in theory able to handle such big simulations, running on HPC platforms can improve the turnaround time, giving the end user a competitive advantage over its peers.

For small simulations, improved productivity on HPC platforms can facilitate more complex design processes (such as parametric studies).

As already discussed in Section 4, simulations on an HPC platform can achieve very good scalability, therefore making more efficient use of computing resources, in particular if the cloud-based service is charged on a pay-per-use basis. On the other hand, in-house systems can be idle many times, so hardware investment is often not fully used.

HPC platforms are maintained by highly skilled professionals and often have dedicated support teams. There are also likely to be domain experts available to help certain applications (e.g. ARCHER's CSE team have several experts specialised in CFD). These are particularly helpful for applications with complex software packages like OpenFOAM. On the other hand, on local cluster maintaining centrally installed and up-to-date version of the software (which is not straight-forward) often needs to be done by ordinary users, which is not a good use of staff times. Similarly, standard cloud computing offers (such as Amazon EC2 or similar) are lack of domain experts.

4.7.7 Feedback to Work Package 2

Feedback to Work Package 2, Core Service Requirements Capture and Infrastructure Development:

Setting up an OpenFOAM simulation involves lots of user inputs. This is conventionally done via supplying text files in predefined format. On a cloud-based environment, we envisage that such inputs are collected via some user interface (e.g. web) and information required by OpenFOAM is then generated on the fly. For this to happen, it is desirable to have some form of programming interface to facilitate the implementation of such functionalities in the portal.

Of course, standard web-based functionalities, such as file uploading (for supplying files required by simulations, such as CAD model), text editing (for users to reconfigure their simulations), etc., are required.

4.7.8 Feedback to Work Package 3

Feedback to Work Package 3, Core Service Deployment and Facility Operation:

Ongoing efforts like enabling and supporting remote visualization should be carried on, as those techniques can significantly streamline the CFD simulation and evaluation process.

4.7.9 Input to Work Package 8

Input to Work Package 8, Exploitation and Sustainability:

Nothing additional to the Market Analysis Questionnaires we already completed.

4.7.10 Input to Work Package 9

Input to Work Package 9, Business Models for ISVs:

We have no input at this point.

4.7.11 Any further input to Work Package 10

Further input to Work Package 10, Dissemination and Training:

There have been plenty of offers to attend different Fortissimo related training courses. Due to time constraints and the fact that most of our project members already attended similar events at HLRS we did not yet attend one.

A remote visualization event might be interesting, either as a webinar or on-site.

4.8 Experiment 408: Optimisation of Aeroplane Wiring

4.8.1 Overview of the experiment

The current trend for SMEs in need of computational resources is to use public Cloud infrastructures. The main reasons being expected lower costs than those of maintaining an own infrastructure and an increased flexibility in response to changing computing demand. A closer look shows that the applications often are web-servers, shops, back-ends for mobile apps etc. SMEs running industrial simulation codes on Cloud resources can be observed less frequently. Three reasons can be identified: (i) industrial simulation codes are complex and the learning curve is rather flat than steep when the codes are not used frequently because of SMEs' different production cycles. (ii) The simulation codes are usually designed for parallel execution in HPC environments and achieving a good performance in a public Cloud environment requires knowledge that is often not part of the core competences of the SME. (iii) Industrial simulation codes are protected by software licenses that in general are tailored for executing the simulation using local resources.

Within the European funded Fortissimo project, the MORE.CLOUD experiment (Main Routing Architecture Optimisation Research Experiment in the CLOUD) implements a solution for:

- Formalizing and executing business processes as simulation tasks (Engineering as a Service) of the MORE experiment;

- Using license-protected simulation applications in off-site cloud-based HPC environments that leverage the elasticLM flexible licensing mechanism provided by the ISV;
- Solving the problem of integrating with the flexible licensing software and providing Engineering Services to manufacturing industries.

MORE or the Main Routing Architecture Optimisation Research Experiment is a documented industrial use-case provided by the SME KE-works and the process integration and design optimization company Noesis Solutions N.V. The MORE simulation workflow optimization experiments aims at reducing weight and cost for wiring systems for complex products such as aircraft (Figure 9).

Figure 9: Main Routing Architecture Optimisation Research Experiment is aimed at reducing weight and cost for wiring systems for complex products such as aircraft

The MORE experiment involves the execution of different applications that are orchestrated by the Process Integration and Design Optimization (PIDO) tool Optimus from Noesis. These applications have been modified to run within the HPC cloud infrastructure and enabled to use the flexible license manager component available in the Fortissimo platform. The simulation workflow of MORE contains three main applications developed by KE-works: MORE, AWARD, and CSI. The MORE application proposes and defines a main routing architecture, AWARD (Automatic Wiring and Routing Design) uses signals, which are routed across the proposed main routing architecture and CSI (Connector Selection Instrument) chooses the respective industrial connectors which are used to connect multiple wiring assemblies. The simulation workflow provides optimisation technologies to systematically search the design space for ‘better main routing architecture designs’. Hence MORE.CLOUD is the architecture that forms the outcome of this experiment, running the MORE simulation workflow in the HPC cloud infrastructure and enabled by the License as a Service licensing provider.

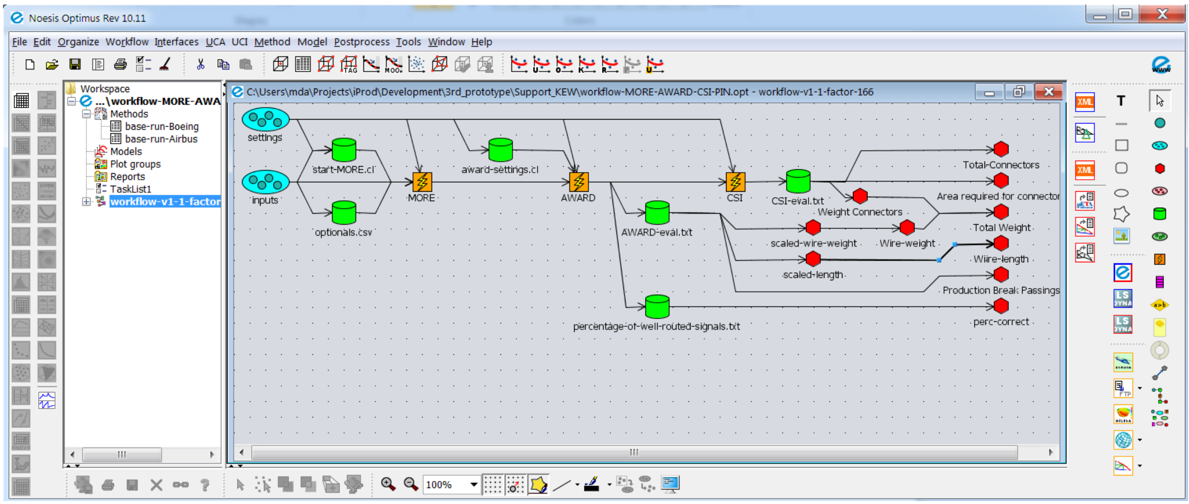

Figure 10: Simulation Workflow of the MORE experiment

The experiment has the following objectives:

1. Formalization and execution of business processes as simulation tasks (Engineering as a Service) of the MORE experiment;
2. Use of license-protected simulation applications in off-site cloud-based HPC environments using elasticLM;
3. Solve the KEW challenge of integrating the elasticLM software and providing Engineering Services to manufacturing industries.
4. Take advantage of using a workflow-based Cloud HPC environment dynamically and on demand with a pay-per-use business model without taking the burden to buy and maintain the entire computational infrastructure.

The rationale behind the experiment consists in providing two possibilities for the SME to use an application in a cloud-based HPC infrastructure:

- It will allow the use of already purchased licenses for running an application locally at the SME or in a cloud-based HPC environment without additional cost.
- It will allow a pay-as-you-go (operating expenditure) approach, only paying for additional licenses as needed when running the application using external the cloud-based HPC resources.

The innovative product design solution and license management solutions will be the two key ingredients of this experiment.

Figure 11 shows the functional architecture of the experiment setup including a provisioning part (Engineering as a Service frontend) and a computing part (HPC cloud-based simulation).

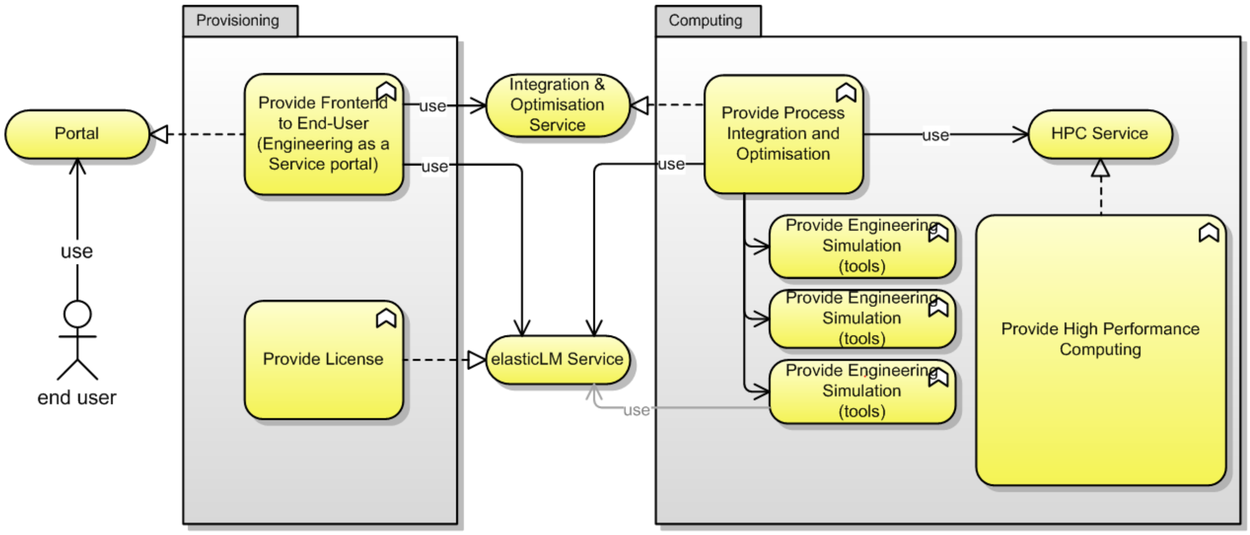

Figure 11: Functional architecture of the experiment setup

The relevant market sectors involved in this experiment are:

- **Manufacturing industry:** KE-works acting as the end user of this experiment aims to reduce weight and cost for wiring systems for complex products such as an aircraft.
- **Application Expert:** KE-works developed dedicated design applications for wiring systems for complex products such as an aircraft
- **ISV:** Noesis acts as the Software Vendor proving an optimizer tool applicable for the experiment objective.
- **ISV:** KE-works acts as a Software Vendor providing a web-based workflow management platform (KE-chain) for the formalization and execution of business to simulation tasks (Engineering as a Service) of the MORE experiment
- **ISV:** Fraunhofer SCAI acts as a Software Vendor providing a license manager tool which will manage the two different applications involved in the experiment, KEchain and Optimus.
- **HPC provider:** Gompute acts as the High Performance Computing provider making available the computing resources required for the experiment through their platform.

4.8.2 The industrial setting

There are many auxiliary companies furnishing the aeronautic industry. Some of these are small and medium-sized enterprises (SMEs). In order to cope with the technical design challenges, computing simulation has yet raised as an efficient working solution. Computing simulation applications such wiring optimizers, are very computational demanding tools.

The price to acquire a High Performance Computer is high. Moreover, the efficient usage of the HPC is not easily guaranteed, since the demand is dependent on specific projects, making it very unstable.

Current practice is to stick to a discrete size HPC workstation not to exceed the average computing load. Often, this short dimensioned workstation entails that the solution delivery is postponed. The collaboration between the SMEs and HPC centres will directly influence the response time and solution accuracy.

With regards to this supply chain, the present experiment will motivate companies with similar challenges to shift towards computational efficient frameworks that are easy to use and lower the barriers for engineering use. This is particularly true for SMEs, where the impact of the FORTISSIMO experiment is larger. In fact, those SMEs (here represented by KE-works) will be capable of handling complex multidisciplinary simulations without necessarily having to invest for internal hardware and software resources and still maintain a clear control over computational costs. In this way they will become capable of handling such large and complex simulation workflows (here provided by Noesis) with a relatively smaller investment. This will allow SMEs to both:

- Use existing computational services to execute complex multidisciplinary simulations, thus extending their capabilities and expertise towards more complex engineering scenarios;
- As a consequence of this, be able to offer extended engineering services, thus not being limited anymore to the amount of available in-house computational resources..

The outcome of this process will be a much lower barrier for SMEs as service providers to access multidisciplinary design optimization methods, allowing them to work with complex simulation workflows and offering more complete and competitive engineering services.

4.8.3 The improvements in the industrial process through the use of an HPC Cloud

The traditional trial-and-error approach of European SMEs to the design the wiring layout in aeroplanes lies in the past. Software optimizers have taken over that responsibility, reducing costs and improving the results. This practice is well accepted by all parties, although the HPC hardware acquisition is not an easy task. Prices are high and number of cores dimension is not clear upfront.

By accessing to the Fortissimo market place, SMEs can compete at equal position with the large enterprises, connecting to thousands of cores if reducing costs and delivery times are required.

The use of the Fortissimo one-stop shop which brings together supercomputing resources, domain specific expertise and licensed simulation packages can overcome these barriers. This enables European SMEs to compete globally through better design and quicker times to market.

The product design framework implemented in this Fortissimo experiment provides engineers with a flexible and scalable tool for the execution of all product design steps, including engineering simulations. The experiment shows flexibility in the usage of computational resources resulting in affordable cost for high computational capabilities by a cloud of HPC resources instead of maintaining an in-house HPC infrastructure. The experiment showed feasibility of running custom KBE applications in an optimization framework on a Cloud HPC. The optimal result was reached 20 times faster than on a single node when using 30 nodes in parallel.

Due to the integration with a cloud-based environment, end users can benefit from a “one-stop-shop” where they can define their product design steps from a business logic point of view and link them to the necessary simulation steps. In this way, end users do not need to have different and fragmented tools to define their design and simulation processes, but have a “one-stop-shop” where all these components are provided in one single, flexible and scalable solution.

Furthermore the simulation steps take advantage of the Cloud-based HPC environment for execution, optimizing time, resources and costs thanks to the coupling with the flexible licensing and license management component of this experiment.

At same time this solution allows the ISVs to access to a new market, giving the possibility to expand their customers’ portfolio and product offers opening up new business models related to the HPC cloud computing that became more affordable.

This will yield to the creation of new industrial opportunities based on the creation of an innovative HPC cloud-based approach to engineering that could use Clouds and flexible licensing schemes as the basic means to deliver new software services in the aeronautic industry and in other fields like automotive or transportation systems in general by providing software services to large, medium and small enterprises.

The elasticLM licensing mechanism integrated into the "one-stop-shop" solution offers the ISV the possibility to make tailored software licenses for its application available to its customer. The customer is no longer forced to sign a one-year contract for usage of the application as usual. Rather, the SME with occasional demand for the simulation application can acquire licenses for individual simulations. The ISV in turn is able to attract new customers that it would miss otherwise because of SMEs will avoid the high flat-rate costs for not permanently used yearly software licenses. In the next paragraph the resulting benefits of this experiment are described.

4.8.4 The cost benefits of using an HPC Cloud

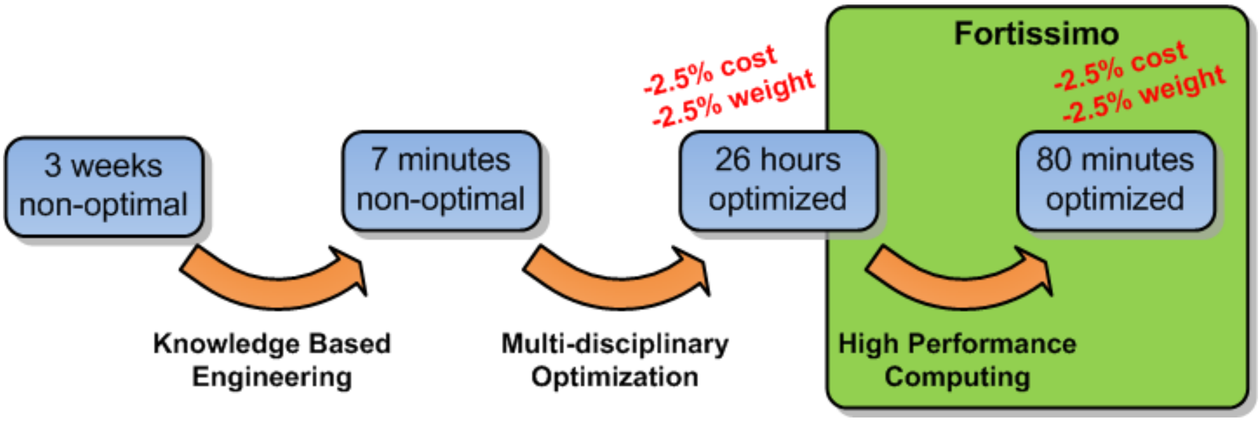

Figure 12: From non-optimal process to optimized cloud HPC process

Comparing the current non-optimal process with the automated & optimized process in a cloud HPC environment shows 2.5% reduction in cost & weight of the wiring system. This benefit is a recurring benefit for hundreds of ship sets that are produced across multiple decades. An aircraft manufacturing program runs for decades) and also a lead-time reduction of 90% (shown in Figure 12). This results in:

- Shorter design loops are possible with fast feedback to customer allowing for later incorporation of design changes or less Cost of Delay and thus more flexibility towards customer resulting in higher customer satisfaction (comparison between duration and number of threads shown in Figure 13).
- Possibility of inclusion of additional routing design variables (i.e. more complex analyses), resulting in even higher cost and weight savings in same amount of time.

Parallel thread	Duration [min]	CPU hours [hrs]
1	1560	26
30	80	40
60	56	56

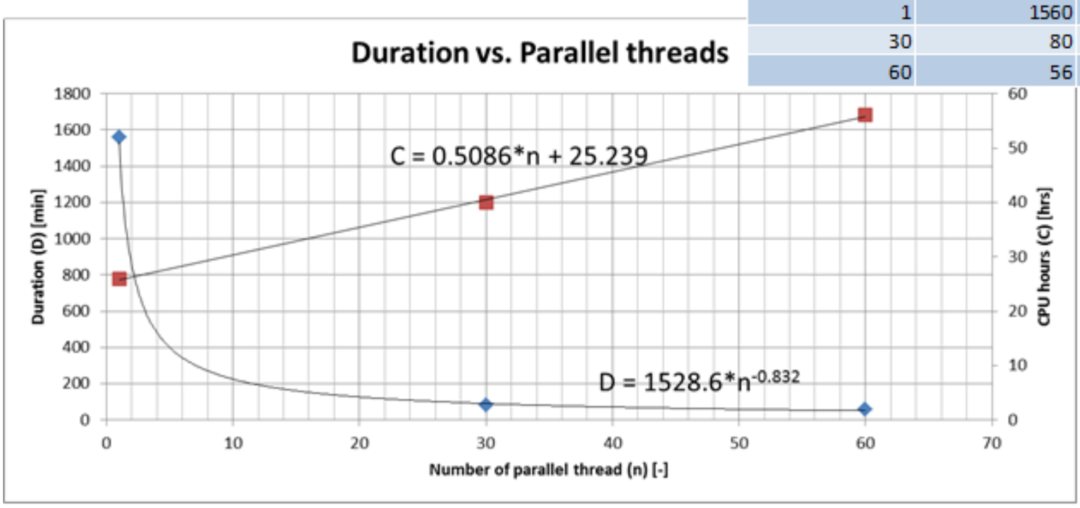

Figure 13: Cost of Delay vs computation time

Comparing in-house and cloud **HPC**; Cost of computation on cloud HPC is about EUR 660 for a single experiment run, while Cost of in-house HPC is about 61 kEUR / year. RFPs for design programmes are requested only once per year on average, so having an in-house HPC facility is much more expensive than using cloud HPC. Furthermore the in-house HPC facility is much less scalable if more computation power is required. This proves the feasibility of using cloud HPC for engineering simulations by SME's in order to compete with large enterprises.

The experiment has proven the feasibility and benefits of new models for software licensing when using an HPC cloud. Both the ISV as license provider and the end-user as user of the license-protected software of the ISV take advantage of the on-demand licenses. With the technology for flexible license management in Cloud computing environments new business models can be implemented that provide a significant advantage for the SME end-user and create new business opportunities for the ISV.

1. The SME end-user does not have to pay for an expensive annual license, which is only used during periods where optimised simulations for the customer of the SME need to be prepared. The other time of the year it just contributes to the basic cost of the SME without any return. Instead, the SME only pay for the time the license is required to perform the optimisation of the simulation.
2. The ISV in contrast would win a new customer with the pay-as-you-go model. The end-user would stay with the manual optimisation otherwise to be able to provide a competitive price to its customer.

It could be shown that an important prerequisite for the ISV can be provided by the elasticLM solution for flexible software licensing: the protection against fraud is on the same level as provided by other traditional license management systems while providing the functionality required in Cloud environments, which are both outside the administrative domains of the ISV and its customer.

For Fraunhofer SCAI its contributions to the Fortissimo experiment can be exploited for presentations and demonstrations at suitable events to convince more ISVs to move from the traditional, rather inflexible licensing and corresponding business models to a flexible solution that allows extending the market from in-house solutions to Cloud environments.

I discussed internally with our CCO if we can somehow quantify these benefit more.
His answer was that at this point it is impossible to quantify the business benefit more that we have already done.

The reason is that by far the main benefits (largely reduced lead-time and some improved quality) have perceived value by the customer of the Electrical Wiring Company with respect to the possible quick reaction to changes and requests, the higher quality of the proposed design and reduced risk of unfeasible design. These benefits are not possible to quantify at this point since no information is available and each project/customer is different.

With the results of this project we can start investigating what the perceived value (and thus the willingness to pay) of these results is by demonstrating the experiment results to our customer (the wiring company).
From that point on we might get more insight in the value of reduced lead-time and improved quality.

At this point putting numbers on these point is purely guessing. I can however add the explanation above to the success story.

4.8.5 Any Other Business benefits

The use of simulation flows for proposals result in faster response time to RFPs, high quality estimations or offering lower cost and lower weight wire harnesses due to optimized routing architecture, thereby having a lower risk to wire harness manufacturer and OEM, larger profit margin, and/or higher chance of winning contracts. Furthermore the cost and time savings result in increased competitiveness with BRIC countries.

4.8.6 HPC as opposed to standard cluster offerings

The HPC platform combines high performance hardware together with a support team who is entitled to maximize the efficiency of the resources.
In this case in particular, the experiment 408 has come across a series of integrations in the HPC platform that would not be possible without the help of the support team. Support has been a key factor during this experiment.
Furthermore, fast interconnection between nodes has remarkably reduced the simulation time. This offering is not available on standard commercial clusters.

4.8.7 Feedback to Work Package 2

From the software perspective, technical knowledge and effort is required to migrate the various softwares into a Linux environment and to integrate the Optimizer submission commands into the HPC queuing system.
From the HPC centre perspective, it has been required to provision a dedicated cluster composed by a private login node, private computational nodes and a virtual license server. All accessible by a public IP address. Root permissions have been required from the end users to do their own configurations.

4.8.8 Feedback to Work Package 3

Private license server and Root permissions for the end users are required to guarantee the operability of the system.

4.8.9 Input to Work Package 8

Given the successful results achieved by the experiment activities and previously described, there is a substantial exploitation potential of the Fortissimo solution into the current engineering industrial practices. This potential is mainly unleashed in the field of design engineering, especially for SMEs and ISVs as solution providers to larger industry. As such, experiment partners have prepared a first exploitation plan to progress towards the commercialization of the Fortissimo HPC-Cloud platform and the experiment results. In particular, the following main key results will be exploited after experiment end:

- Workflow as a Service solution, providing both business and simulation workflows in the Fortissimo marketplace.
- Integrated solutions for handling and solving complex and compute demanding multidisciplinary engineering simulations in a shorter time-frame, with better results and reduced infrastructure investments.
- Integration of elastic license management systems into Optimus for adoption of flexible business models
- Integration of elastic license management systems into KBE applications for adoption of flexible business models
- KBE simulation and workflow setup on a cloud HPC platform via Optimus
- HPC cloud computing integration into Optimus
- Linux based KBE applications previously available only on Windows machines

These results have addressed the main barriers that so far prevented SMEs from providing cloud and service solutions due to the high costs of building and maintaining them (e.g. HPC hardware, expensive long term licenses, availability of technical competences etc). For this reason a new set of business opportunities has now become accessible to the partners of experiment 408 and plans for exploitation and commercialization of the above results have been drafted.

In particular:

- Noesis will introduce the simulation workflow as a service feature in its product in 2015, with the introduction of a new web based product for design engineering
- KE-works will extend its customer base by providing KE-chain and the KBE applications as a service for its customers
- Fraunhofer SCAI will provide facilities for better integration of the license manager in new products and for future research on business models for HPC cloud services
- Gompute will benefit from the results by providing a better and more flexible system that could allow outsourcing of business and simulation processes instead of only simulation tasks, thus extending their business to new customers.

More generally, thanks to the extreme flexibility of the approach implemented in this experiment for what concerns the management and coupling of the business and engineering simulation processes, the support of a high computational efficient framework and the advantages of using a reliable and elastic license managements system, the core partners of experiment 408 on the base of their past experiences do believe that this methodology can be easily and proficiently applied in a number of other contexts and\or applications where there is a high demand of computational power and efficiency. The wiring harness problem addressed here, indeed, can be considered as an outstanding successful proof of concept that requires relatively small investment before a real deployment in an industrial environment. In the same time it can help into creating the conditions for increasing business for ISVs from all the value chain by motivating more and more companies with similar challenges to become users of the Fortissimo marketplace.

A special attention has been paid by the partners also to the sustainability of the platform proposed in Fortissimo. In particular, relying on preliminary results achieved in this experiment reported in section 4, and considering the great flexibility and applicability of the proposed solution in several other industrial settings there is enough confidence that there will be a positive balance between benefits\profits and costs\investments in bringing Fortissimo to the market. According to partners of experiment 408 those numbers represents an important selling point to bring to the attentions of the top managers of the ISVs potentially interested in participating and sustaining the Fortissimo marketplace as well as to attracting more and more new end-users.

4.8.10 Input to Work Package 9

The Fraunhofer Institute SCAI considers the Fortissimo Marketplace as an excellent opportunity to exhibit its license management solution elasticLM and to gain access to potential SME customers in Cloud environments. SCAI is plans to join forces with Independent Software Vendors, their customers and Cloud Service providers. Within the Fortissimo experiments elasticLM
SCAI plans to make elasticLM available as product in the marketplace. The offering will include both a stand-alone version with consulting for integration and use provided as needed by the potential customers. Moreover, SCAI will also bundle elasticLM with other products, e.g. software of ISVs, Cloud infrastructures, to provide end-users with one-stop-shop solutions.
For SCAI the Fortissimo Marketplace is a valuable extension of exploitation of project results beyond the end of the project.

4.8.11 Any further input to Work Package 10

Currently using HPC platforms is not really straightforward, a lot of deep-down technical knowledge is required when you want to get an experiment up and running. More detailed training and documentation of the use of cloud HPC (specifically the Fortissimo one-stop shop) is required before SMEs can start using HPC and optimization software.

Furthermore during the course of the experiment we did not get much insight in the goals and vision of the Fortissimo Marketplace, the experiment was executed between just one specific HPC provider and several other experiment partners, making it very difficult to see the experiment in the context of the complete Fortissimo project and provide correct feedback. More dissemination of results of other work packages in the form of presentations, web video tutorials etc. would benefit in giving better feedback.

4.9 Experiment 409: Simulating Automobile Components

4.9.1 Overview of the experiment

The objective of the project is to showcase the use of cloud based HPC services to investigate new mode of use for numerical simulation to find correlation between Scilab/Xcos models and real tests systems. SDI, the end-user company is a twenty years old SME working in test engineering for automotive and aeronautical industries.

Demonstration of new mode of use for simulation model sharing, correlation between simulation and real data would open new markets to develop.

The software provider in this project is Scilab Enterprises, the developer and publisher of Free Open Source software Scilab for numerical computation. The project opens new perspectives for a future version of the Scilab/Xcos that can be used in a cloud environment for industrial processes.

The project will address porting the Scilab software and testing it on the Bull extreme factory cloud-based HPC infrastructure. Scilab Enterprises and the Bull Fortissimo Competence

Centre will work together to produce the cloud-based version of the software that meets the needs of the end-user. An evaluation of the viability of this service on commercial HPC clouds will be carried out, and possible business models will be proposed.

4.9.2 The industrial setting

The experiment is about project and confidentiality management regarding HPC Cloud Simulation. This approach will allow to avoid consuming big data transfers and replications which is still taking too much time in a cloud based architecture.
We set up 2 different scenarios. The first one is based on a brake manufacturer and a vehicle modeler working together, remotely on a cloud-based infrastructure. The second one is about a main contractor wishing to compare the performance between 2 brake manufacturers. Each use case illustrates well different levels of confidentiality between main and sub-contractors.

4.9.3 The improvements in the industrial process through the use of an HPC Cloud

The possibility to launch simulations via the Fortissimo HPC cloud will help to increase the numbers of scenarios to evaluate, but also to reduce CPU time. For an end-user, the effect could be either an increase of productivity or a commercial advantage compared to company using a classic client-server model, without counting the significant cost reduction by avoiding the need to maintain a computing cluster in-house and the associated human expertise.

4.9.4 The cost benefits of using an HPC Cloud

Since user-cases have not yet started on HPC-Cloud, it is not yet possible to evaluate the cost benefits of such architecture/service. This evaluation will be performed on the last 6 months of experiment.

What both SDI and Scilab initially planned as 2 companies that are not HPC users were about two main objectives:

- Be ready to access HPC infrastructure.
- Use an HPC-Cloud infrastructure as a neutral place to improve the relationship between industrial main and sub-contractors.

As per simulation time gain, a study of one transfer function takes between 1hour and 4hours at SDI. SDI is compiling the cost benefit in terms of time and resources to reduce this simulation cost on a case study.

The results of the experiment are very promising. The overhead induced by a script managing the computation node is negligible in comparison with the time of simulation. The speed-up obtained is almost equal to the number of computation nodes (sum of the individual simulation times divided by the time taken by the longest simulation).

The test was realised using 4 nodes, but the parallelization used easily extends to as many nodes as available for the calculation. This paves the way to new optimization techniques where the parameters are chosen automatically in a great number of simulations instead of the current approach at SDI where a qualified engineer selects carefully the parameters they deem important.

With their current usage, SDI deems that the maximum cost of the offer for the marketplace is 100€ per month (with a hope of reducing case study time by about 70%). If the technology is satisfying they plan on a basis of 3 contracts with the use of Cloud co-simulation for 2015, 6 contracts in 2016 to 20 contracts by the end of 2017.

4.9.5 Any Other Business benefits

One of the goals of the experiment was to demonstrate that confidentiality in the use of co-simulation in a neutral place was achievable. Sharing a simulation with business partners in a unified manner without exposing the know-how was demonstrated with the use of the Functional Mock-up Interface (FMI) standards and the related Function Mock-up Units (FMU).
This approach interested FAW, third Chinese car manufacturer and a client of SDI, willing to make sure their intellectual property was protected. On the other end, SDI could propose working on a Cloud based test bench for the case study without having to share models but instead FMUs.
As a direct result of the experiment SDI is negotiating a contract with FAW with revenues ranging between 200k€ and 500k€. European car manufacturers and BRICS countries are also seduced by the idea.

SDI plans an increase in demand for the cloud based test bench in the future years to come and plans to extend their business to BRICS countries by 2017. (Currently 50% of their business is conducted in Europe, 30% in Asia).
Target market for cloud base correlation between simulation and test sytem (million euros for SDI)

2015	2016	2017	2018
0.1	0.2	0.25	0.5

Target test bench full market inducted by cloud base correlation (million euro for SDI)

2015	2016	2017	2018
1	1.5	2	3

4.9.6 HPC as opposed to standard cluster offerings

Scilab parallel scalability is not optimal today, and does not take full advantage of HPC systems, but this situation will not last and the upcoming Scilab 6 version will be perfectly suited to HPC systems.

4.9.7 Feedback to Work Package 2

Scilab and SDI have no further feedback to Work Package 2 having already supplied initial requirements in response to the requirements gathering questionnaire.

4.9.8 Feedback to Work Package 3

Despite the security aspect being a core problem on HPC platform, a fix IP constraint, or the inability to have tokens for fast limited access can reduce the attractiveness of a HPC Cloud platform.

During the project, SDI had to get a fix IP address, in France this require a formal demand to the internet access provider. It took SDI several weeks to get a proper set-up to work on BullXF. This setup is generally reduced in time for general Cloud clusters.

4.9.9 Input to Work Package 8

Scilab and SDI completed the Market Analysis questionnaire for Work Package 8 and don't have any additional input.

4.9.10 Input to Work Package 9

For Scilab Enterprises, the gain is not done by product selling as Scilab and Xcos are free, open-source software.

The business model of Scilab Enterprises is mainly based on services provided to customers. The services lead to focused training on Scilab/Xcos, specific training on some plugins available Scilab, targeted development to improve the performance of existing scripts. The project created important business bonds with SDI.

With Fortissimo, Scilab has gained the experience to run parallel scripts and simulation of Xcos for industrial partners. As HPC is one of the main items in Scilab Enterprises roadmap for the Scilab and Xcos software, having an industrial example of a run on Bull XF improves our credibility in the HPC field.

A business analysis leads to the following figures:

Installation Cost and verification of the parallel architecture: 3000€ (while the HPC Cloud marketplace with a Scilab offer is unavailable)

Training for ScilabMPI: 2000€ per person trained (based on 1 day training on SciMPI/SciGPGPU plugin)

Script Audit and Upgrade to SciMPI: 7500€ to 15000€ (based on the minimal estimation done with the audit already done in year 2014 for Scilab Enterprises)

This leads to a price of 12500€ to 25000€ per such package sold (estimated target benefits range from 6000€ to 18000€ worst to best case scenario) with an average of 12000€. Currently Scilab Enterprises has sold one such package in year 2014. We estimate the number will reach 3 units (benefits of 36000€) in 2015, 5 units (60000€) in 2016, 8 units in 2017 (96000€).

4.9.11 Any further input to Work Package 10

During dissemination and training, the aspect of confidentiality could be added to the benefits of having a co-simulation neutral place.

4.10 Experiment 410: Prediction of Air Quality

4.10.1 Overview of the experiment

The objective of the experiment is to showcase the use of cloud based HPC services to investigate the air-quality at city scale. The European regulations on air-quality protection plan require more and more to test and evaluate adaptation/reduction scenarios. From traditionally the evaluation of 2-3 scenarios, consulting companies and regional air quality agencies are facing now to launch tens of scenarios. This requires a large increase in the capacity for calculation by the end-user beyond that which can be managed at their scale.

Demonstration of a cost-effective cloud-based HPC simulation service would represent a breakthrough and solution for the end-user. In this experiment, the end-user is NUMTECH, a French SME whose activities are to provide support to cities and regional institutions on managing the question of local air-quality.

The software provider in this project is CERC. They will provide their software solution for evaluating air-quality at city scale: ADMS Urban. Today the solution is distributed as classic client software. Even if NUMTECH had its own computing facilities, the increasing need in computation capacity is beyond what they can provide if we consider modelling different cities at the same time. The experiment will provide NUMTECH with a version of the

CERC's software that can be used in a cloud environment, together with an insight into how this could be exploited (introducing more flexibility use for the end-user and interaction with them) as a future business model.

The experiment will address porting the commercial software and testing it on the Fortissimo cloud-based HPC infrastructure. CERC and the Fortissimo Competence Centres will work together to produce the cloud-based version of the software that meets the needs of the end-user. An evaluation of the viability of this service on commercial HPC clouds will be carried out, and possible business models will be proposed.

4.10.2 The industrial setting

The European regulations on air-quality protection plan require more and more to test and evaluate adaptation/reduction scenarios. From traditionally the evaluation of 2 to 3 scenarios, consulting companies and regional air quality agencies are now facing to launch tens of scenarios. This requires a large increase in the computing capacity for calculation, beyond that which can be managed by a typical end-user. Moreover, it is also required to quantify uncertainties associated to air-quality simulations. An increase capacity of CPU resource is a way to achieve such evaluation. Demonstration of a cost-effective cloud-based HPC simulation service would then represent a breakthrough and solution for the end-user.

4.10.3 The improvements in the industrial process through the use of an HPC Cloud

The possibility to launch urban air-quality simulation via the Fortissimo HPC cloud will help to increase the numbers of scenario to evaluate, but also to reduce CPU time. For an end-user, the effect could be either an increase of productivity or a commercial advantage compared to company using a classic client-server model, without counting the significant cost reduction by avoiding the need to maintain a computing cluster in-house and the associated human expertise.

Practically, the expected impacts are:

- For NUMTECH (end-user), the ability to access cost-effective HPC service, enabling them to serve existing customers better and gain new market share, especially at short-term by the capacity of providing a service to allow the testing and evaluation of tens of scenario required now by the regulation on air-quality, and to perform uncertainties analysis. At mid-term, the objective is also to use HPC service for operational system to monitor in near real-time air-quality over cities combining numerical simulations and measurement.
- For CERC (service provider), the ability to offer its software in a cloud-based HPC environment, expected to open further business opportunities.
- For Fortissimo Competence Centre (Bull), a business opportunity in the link with CERC/NUMTECH in a joint framework to propose new technical services by means of cloud HPC.

Technically, the CERC/Bull exchange on APIs and workload submitting processes is helping to fully specify the API that is already planned and will be offered to all ISVs and End-Users when ready.

- For Fortissimo project, a feedback on requirements for the Fortissimo Infrastructure.

4.10.4 The cost benefits of using an HPC Cloud

For the conditions of the experiment (that is to say considering the cpu available on the cloud service and on NUMTECH's Internal servers, considering the today cpu cost on Extreme

Factory, and considering the today cost of NUMTECH's IT people etc.), the analysis performed show:

- If there is no constraint on the time to produce results for a typical urban air-quality study (i.e. simulations can take 3 months or more), the cost to exploit HPC cloud platform is higher than to acquire and maintain recent internal server. At the reverse, if the time to perform a simulation is short (1 month or less), then the exploitation of HPC cloud platform (due to the activation and pay on demand) is economically interesting (from a factor 2 to 10). This can allow an increased margin while offering the client a reduction of time to do a study; or to propose additional services such as uncertainties evaluation or testing more scenarios. For example, for a time window of one month and for a cost similar to those using recent NUMTECH's servers, use of XF platform allow to have 2 or 3 more cores.
- When annual percentage of using recent internal server is below 40% (whatever is the number of cores to consider from few to hundreds), the system to pay on demand of cloud service is also economically interesting compared to acquire and maintain internal servers. If the annual percentage of using internal servers is low, then the benefits to use a Cloud service is important with, in this condition, a mean gain factor between 2 and 3. We can note that an important part of the cost for internal server is the cost of IT people, with a constraint in a SME like NUMTECH that there will always have at least one IT person in the company.

4.10.5 Any Other Business benefits

CERC can now offer the ADMS-Urban software as a cloud service, on a pay-for-use basis rather than requiring a customer to purchase an annual licence and install & run the software locally on PCs. This allows for an attractive pricing option for customers needing an infrequent use of the model.

The ADMS-Urban model has been improved in a number of ways and a new release version has been created. The updated version is able to handle the modelling of much larger urban areas in a significantly more user-friendly manner. The whole urban region can now be stored in a single input file and model runs automatically cut down the input to the area of interest. This allows a large urban area to be easily split across multiple processors. Additional utilities were created for the automatic recombination of output from these runs. These new features make the model more attractive to users as it replaces functionality currently performed manually.

CERC's Run Manager software has been extended for use with BULL's Extreme Factory. This made use of the Extreme Factory SSH command interface and scripts to manage job submission, monitoring and retrieval.

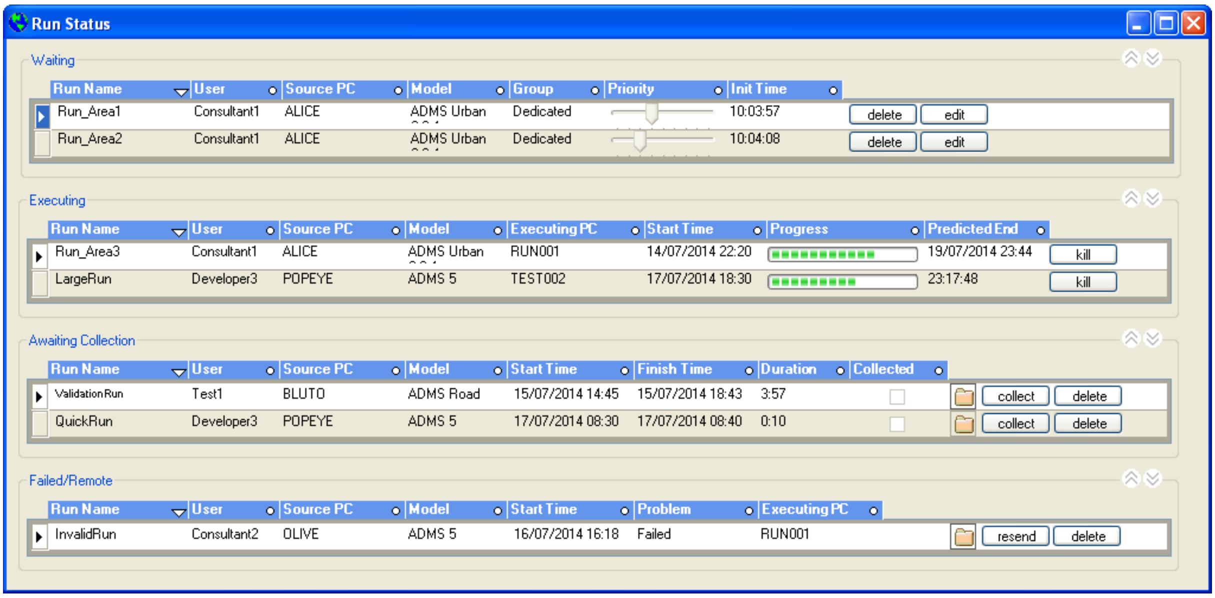

In today's operational needs for HPC Clouds, CERC's request to use an API to interact with the Run Manager fitted perfectly with an idea that was on Bull's roadmap for some time. A web API that can be used on extreme factory to launch remote simulation without the need to live interaction is a must for many computing environments : It minimizes numbers of interactive typing if many jobs need to be submitted and in general helps automate customer workflows.

A first version of the API was defined in accordance with CERC's specifications, and it will be enhanced and updated to cope with new needs and functionalities that we will meet on the market.

4.10.6 HPC as opposed to standard cluster offerings

Extreme factory is a pure HPC cloud offering, with optimized scalable architecture, dedicated management software and fast InfiniBand interconnect. It is also built from scratch with security and ease-of-use in mind.

The ADMS-Urban software could be run on less sophisticated architectures, but with less efficiency, less security, and a perceived cost which does not take into account the time and efforts to setup an operational HPC cluster on an Amazon-type infrastructure.

Also, Amazon-type clouds are not designed for CPU-intensive usage, as their low-cost offerings are typically Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline.

4.10.7 Feedback to Work Package 2

Core service requirements capture & infrastructure development - Defining and developing the market place

The experiment provided feedback to Work Package 2 through the various requirements gathering questionnaires sent out by WP2 throughout the experiment. One of the main requirements CERC have identified is the need for an API to allow automatic submission, monitoring and retrieval of jobs via the market place.

4.10.8 Feedback to Work Package 3

Core service deployment & facility operation - Deployment & operation of the market place

System down-time must be minimal (service-level agreements should be in place with ISVs and end-users).

Both the input and the output from the software can be very large, and require transferring from/to the user's local computer. Therefore the market place must have good connection speed, no data transfer limits and very high availability.
Customers who wish to use the software will demand that their jobs are kept confidential. The work is often legally privileged, so a high level of security must be ensured.

4.10.9 Input to Work Package 8

Exploitation & sustainability
CERC and NUMTECH completed the Market Analysis questionnaire for Work Package 8. For the long-term success of creating a one-stop-shop, the software offering must be allowed as wide as possible and not have any artificial restrictions on software provision. Therefore the market place should include software that doesn't necessarily benefit from HPC, but does require large processing over a cloud.

Adapting ADMS Urban to extreme factory provides opportunities to explore new business models and attract new customers. The work was however significant and it will not be possible to do it again for many new and different HPC cloud providers. A sustainable market place will need to mutualize formats, processes, workflows, OSes, APIs, interfaces …. in order to avoid duplicating efforts and generate intrinsic value.

4.10.10 Input to Work Package 9

Business model for ISVs
The choice of operating system available will affect ISVs business models. If the OS available does not match those the ISV currently develops in, there can be both an initial overhead and additional ongoing development & support costs which must be taken into consideration. This could make using the market place less attractive to an ISV with limited funds.

The ability to offer the ADMS-Urban model through a pay-per-use service needs careful consideration. CERC's current business model generates income from software licensing and training in the use of the software. The licensing is generally on an annual basis and includes benefits such as helpdesk support and product updates during the licence period. Users may run the software as often as they like on a set number of their own PCs without additional cost. For users, this means they have a known fixed annual cost for air quality modelling software fees.

Moving to a pay-per-use licensing model has the benefit of user's only being charged for the work they carry out. This can make software more attractive to users with a limited amount of modelling requirements. However, for users with a high amount of modelling requirements, this licensing model could significantly increase their costs, and they would therefore prefer to retain the existing licensing model.

If both licensing models are provided by CERC, then great care is required when setting the pricing points to ensure there are sufficient new users to offset any loss of income due to existing users switching to the pay-per-use licensing model.

The decision how to charge users on a pay-per-use licensing model is also a difficult one. Charging based purely on run-times may seem attractive, but this then has monetisation implications for an ISV when implementing speed-improvements to the model. Alternatively, charging based on the contents of the input file is another good option. However, this has privacy & confidentiality implications, particularly as CERC and some end-users are in direct competition for consultancy work.

4.10.11 Any further input to Work Package 10

Dissemination & training

The Experiment was advertised on each participant's web site and at various customer meetings.

Further dissemination is envisioned when the Fortissimo market place will materialize and the attached business models will be clarified

4.11 Experiment 411: Simulation of Risers in the Oil Industry

4.11.1 Overview of the experiment

The objective of this experiment is to adapt DeepLines, one of Principia's main software, to cloud-based HPC infrastructure and qualify the feasibility of new business models. To that extend, the following steps were taken:

- Improvement of the parallelization of the code to reduce computational time,
- Simultaneous access to file (read only) in order to perform post-treatments,
- Simulation in parallel of several hundreds to thousands of environments with combinations of different durations.
- Definition of a user case which analyses the business benefits of the experiment, and proposes a viable commercial model for future production work.

This experiment will open new doors to Principia for increasing its market by offering to its customers through Cloud services new and flexible access to its software portfolio. It will enable Principia to propose more sophisticated numerical schemes (multi-physic coupling) and consider very large load case matrices (several thousand of cases to consider for a given field). Currently such offer is not possible within the existing HPC facilities owned by Principia.

Based on PRINCIPIA's and IFP's experience in risers and hydrodynamics, DeepLines® is developed to provide oil and gas engineering companies with an integrated tool to design and optimize their risers, flowlines and moorings systems for shallow to deep-water offshore projects. DeepLines is used daily by engineering companies in the oil and gas sector as well as the marine renewable energy sector.

4.11.2 The industrial setting

The interest of having a version of DeepLines running on HPC cloud is to enable both Principia and its DeepLines customers to use these optimizations to obtain results more quickly, but also to do larger case studies to represent more accurately the offshore environment seen by the system being designed over two to three decades.

4.11.3 The improvements in the industrial process through the use of an HPC Cloud

The main objective of the Fortissimo experience is to perform simulations that are not accessible right now. Due to computer limitations, the current practice in the offshore industry is to simplify the meteocean data to limit the number of simulations necessary for the design but still being conservative. Due to the risk involved, all simplifications are overly conservative and can lead to design which are not optimized. This is a general trend in the offshore industry where systems are to stay in place for 20 to 40 years. Over the past decade, the number of cases to run has dramatically increased both in sheer number (now above 1000) but also in term of simulation length (from 500s to 10800s).

4.11.4 The cost benefits of using an HPC Cloud

At the present time, the cost advantage is limited since there is still half of the calculations (post-treatment) that are not available on HPC. Otherwise, once the post-treatments are available, the profit from the use of HPC will be a time saver for the whole calculation chain, thereby reducing delays and costs of studies, even if there will be an additional cost to licence and rental Cluster.

Therefore to propose a first usable HPC model, we made some changes in Deeplines. We developed a small preliminary post-treatment tool to extract some data directly on the cluster. This avoids recovering all the data and do the post-processing on a server or Windows PC. For example for a simple case studies (630 runs), we went from 60 days of calculations on the customer's hardware to three days on the cluster to have the same outputs.

The main benefit of using the HPC is to reduce the delays calculations, while increasing the number of cases in one study.

4.11.5 Any Other Business benefits

4.11.6 HPC as opposed to standard cluster offerings

The use of HPC cluster types (such as Bull) is more advantageous compared to a standard cluster (such as Amazon). The main reason is the opportunity and ability to tailor a cluster of views CPU, RAM and especially disc type, based on the specific needs. Also a management system of the volume of IO depending on the quantity of calculations to be performed is required. It is, in our case, necessary to use the GPFS system.

4.11.7 Feedback to Work Package 2

Currently, none of our competitors is positioned on this market. We are the first to propose such option; we will have our first order in January. Following the first order, we hope (based on customer's approval) to communicate feedbacks on this experience.

4.11.8 Feedback to Work Package 3

4.11.9 Input to Work Package 8

Currently business model: | HPC business model:

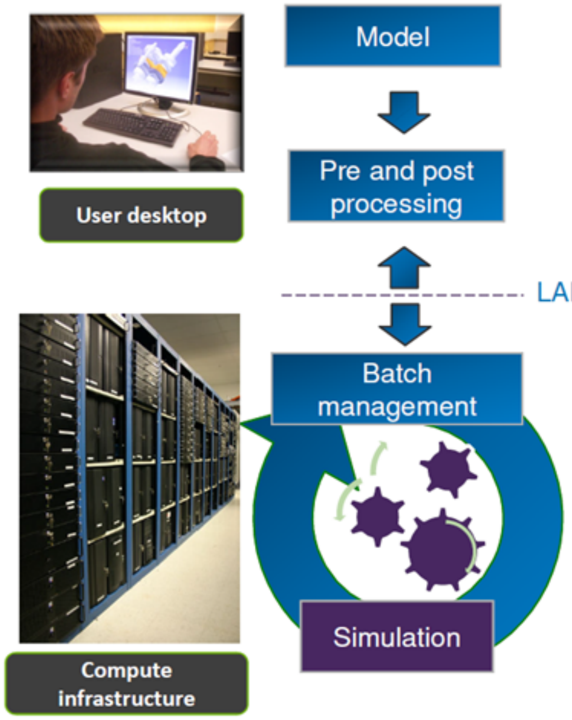

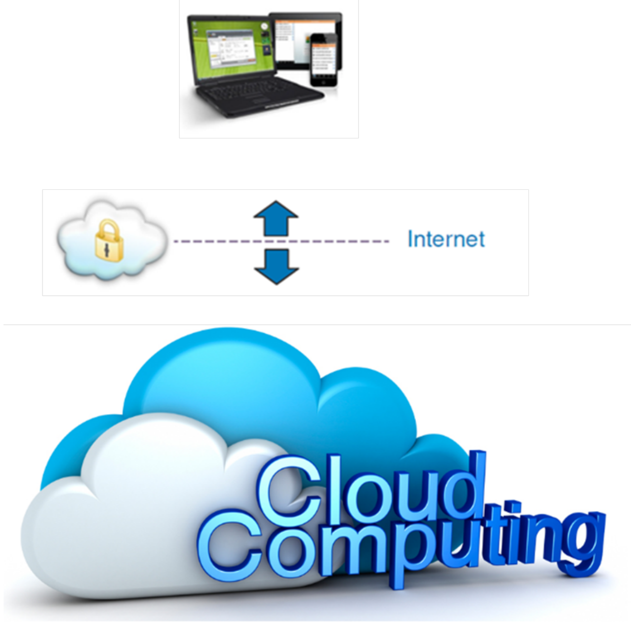

Focus on business model HPC

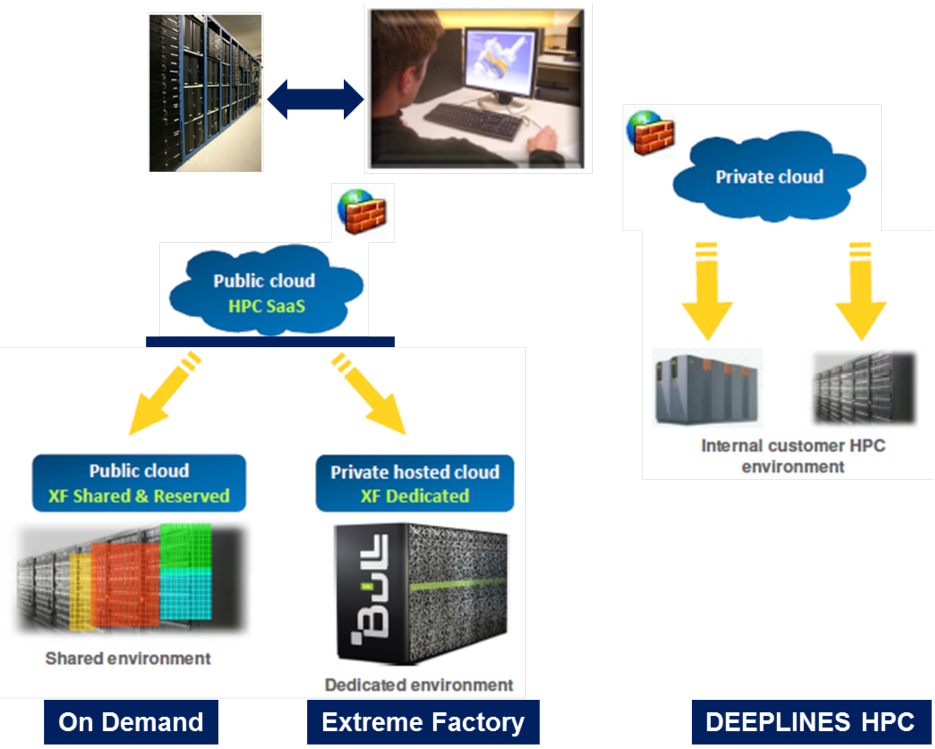

	HARDWARE	YOUR BENEFIT
DEEPLINES HPC : HPC Pack to extend your DEEPLINES multi-core runs (oceano-meteo multi-load cases)	Company Private	• Shorten analysis cycle • Manage yourself run-time
DEEPLINES Saas (Software As A Service)		
DEEPLINES On Demand	External Shared	• Full package Hardware/Software • Hundreds of cores ready-to-use with wide RAM and GPFS disk space • Unlimited DEEPLINES License Booking required to secure computation capacity
DEEPLINES Extreme Fact.	External Dedicated	Idem On Demand BUT permanent computational capacity allocated to Company

4.11.10 Input to Work Package 9

Not planned at this time.

4.11.11 Any further input to Work Package 10

In January we will perform a design study for a customer to decrease the conservatism of the design since there are questions on the validity of the system. Direct meteocean data with minimum processing are used. Thanks to the use of HPC, this approach will take about 3 weeks instead of 3-4 months and make it feasible with respect to the schedule of the project. Therefore, HPC is bringing new possibilities to the offshore industry to improve or qualify concepts.

4.12 Experiment 412: Reduced Vehicular Emissions of CO_2

4.12.1 Overview of the experiment

The experiment "Cloud-based simulation environment for CO2 emission prediction for automotive engines" deals with optimization of vehicle architecture in order to minimize CO2 emission. This is accomplished by running large scale parameter studies with vehicle model software simulation on HPC cloud infrastructure. Such an environment enables SMEs in automotive industry to perform high scale software simulation tasks without large primary investment in computational hardware and permanent software licenses.

Here is the outline of the experiment:

- Calculating vehicle CO2 emissions using simulation software tool AVL CRUISE
- Analyzing simulation results using post-processing software tool AVL CONCERTO
- Selecting optimal vehicle architecture by running large number of vehicle model variants in distributed cloud environment
- Extracting calculation/post-processing core from the desktop environment and preparing it for the remote computing
- Enabling large-scale calculation tasks by employing high power computational center HLRS at University Stuttgart
- Approaching end-user through a secure, user-friendly web-based front-end platform, featuring model/data sharing, result export etc.

4.12.2 The industrial setting

Many vehicle system optimization tasks in automotive industry require simulation of a large number of vehicle model variants. Vehicle system simulation software being analyzed here is a commercial software tool AVL CRUISE while the result analysis is performed with the software tool AVL CONCERTO.

Software description AVL CRUISE:

"AVL CRUISE is the industry's most powerful, robust and adaptable tool for vehicle system and driveline analysis.

CRUISE is AVL's vehicle and powertrain system level simulation tool. It supports everyday tasks in vehicle system and driveline analysis throughout all development phases, from concept planning, through to launch and beyond. Its application envelope covers conventional vehicle powertrains through to highly-advanced hybrid systems and pure electric vehicles. CRUISE offers a streamlined workflow for all kinds of parameter optimization, component matching - guiding the user through to practical and attainable solutions. Due to its structured interfaces and advanced data management concept, CRUISE has established itself as a data communication and system integration tool for different teams within world-leader OEM's and their suppliers."

Software description AVL CONCERTO™ - Data Post-Processing:

"CONCERTO is a generic data post-processing tool that focuses on quick and intuitive signal analysis, validation, correlation and reporting - for any kind of acquired data.

CONCERTO combines the flexible approach of being generic with the ability to implement application specific know-how accessible directly at the user interface. Its open interfaces and integrated comparison mechanisms offer a central data correlation tool for automation, simulation and measurement systems."

In order to accomplish large scale simulation tasks in automotive industry, SME service providers could either:

- run simulation of different vehicle variants sequentially on one machine with one software license (time restriction)
  or
- acquire larger number of CPUs and multiple software licenses (money restriction)

Many SMEs as suppliers in the automotive industry which are dealing with mid-scale, single-focus and short-term projects, can't afford acquiring and maintaining large software and hardware infrastructure. A new approach is needed to offer simulation resources to those SMEs in a competitive packaging.

4.12.3 The improvements in the industrial process through the use of an HPC Cloud

Opposed to that traditional approach, the solution applied in this experiment utilizes cloud computing power of an HPC which allows running many simulation tasks in parallel, according to a customizable schedule and obtaining the results in a shorter time.

This distributed simulation service together with a result analysis tool is offered to end-users within a user-friendly web-based front-end environment, enabling them to access and perform large-scale parameter studies as needed, when needed, while preserving control over progress and results.

4.12.4 The cost benefits of using an HPC Cloud

The most clear cost benefit using HPC cloud resources by SME service providers is the possibility to lease a powerful computing cluster for single projects instead of acquiring and maintain computational resources which would be underutilized the most of the time, and probably even not sufficient when really needed.

This leads to a better management of computational and development resources, opening a whole new set of applications such as extensive parameter variations.

The monetary benefit varies for different use cases. It could be the best described with a monetary benefit of buying an airplane ticket to fly from A to B, instead of buying and maintaining a private airplane which will be used just a couple of times a year.

Roughly, for example, if an end user needs 400 concurrent (annual) SW licenses and an appropriate computing infrastructure and engineering to run a large scale parameter study, it would cause an investment measured by millions of Euros. With all additional cloud overhead, a short term projects running millions of simulations on 400 cloud CPU cores for a period of a couple of weeks, several times a year would bring well over 50% cost savings.

4.12.5 Any Other Business benefits

Different business models (short-term lease, pay-per-use e.g. per CPU-time, per simulation time usage or per simulation run) can make cloud-based computing solutions attractive to SMEs and lead to new applications, better offers and sustainable growth.

On the long term, cloud based computational SW and HW resources, web-access and system administration allow for the global distributed development and service process.

4.12.6 HPC as opposed to standard cluster offerings

Opposed to standard computing clusters, HPC, in this case HLRS of University Stuttgart, is more suitable for this experiment because it offers more standardized utilization of parallel computing resources in terms of hardware and software architecture, queuing system, status check, error handling etc. and enables focusing on the engineering service task.

4.12.7 Feedback to Work Package 2

(Core Service Requirements Capture and Infrastructure Development)

Standardized and automatized interface for the external applications is the key requirement upon core service and infrastructure providers.

Opposed to the traditional "log-in-and-start-the-job" cloud approach, an automatized, web-based server approach is required.

Interaction with other applications as well as flexible and sustainable infrastructure usage business model are of significant importance too.

4.12.8 Feedback to Work Package 3

(Core Service Deployment and Facility Operation)

Stability of remote access, ability to run permanent server machine processes (e.g. MySQL), better status query response and resource availability estimation would significantly improve the deployment of application to HPCs.

4.12.9 Input to Work Package 8

(Exploitation and Sustainability)

Clear concept of FORTISSIMO Marketplace content, business models and maintenance strategy as well as clear definition of different contributor profiles and roles would improve acceptance of the joint service platform.

4.12.10 Input to Work Package 9

(Business Models for ISVs)

New business models for ISVs (new applications, charge-per-use, flexible time/resource based lease etc.) can be derived with the web-access support, standardized interface from HPCs to external applications, detailed usage record and cloud administration tool.

4.12.11 Any further input to Work Package 10

(Dissemination and Training)
For the contributors to the future FORTISSIMO Marketplace platform, comprehensive documentation and training programs should be offered, depending on their roles and activities.

4.13 Experiment 413: Multiphysics Simulations

4.13.1 Overview of the experiment

Project Goal
The idea of this experiment 413 ‘Fortissimo - Multiphysics Simulation in the Cloud’ is to provide multiphysics evaluation capabilities and related support services as a remote cloud service from HPC-centres. Such services will provide access to various simulation codes, an open code-coupling environment [MpCCI], a valuable and growing number of tutorials, and enough CPU resources to run even bigger jobs. Software licenses can be provided from the HPC centre (HPC centre has a contract with ISV which allows the use of software by third parties) or by the end-users (HPC centres re-routes its license request to external license server CPU). Experts for the software tools will also be available for support.

Target End-Users and Applications for WP413
The targeted end-users in 413 are manufacturers of electrical components like transformers, e-motors, electro-magnetic-controlled valves and injectors, etc. During the operation of such components local heat losses are a major issue and may lead to lower efficiency and even

damage. Such heat losses or local hot spots cannot be examined by real experiments due to complex structure of components and missing (or too expensive) measurement utilities sensors.
A joined simulation of the electromagnetic effects as well as the fluid effects can help to simulate these effects in one coupled multi-physics approach. In our experiment we will use the commercial codes

- JMAG from JSOL (Japan), distributed and supported by partner Powersys,
- MpCCI from Fraunhofer as coupling interface, and
- Fluent/ANSYS as 3rd party software.

The distributors for the software components JMAG and MpCCI are represented in the experiment and thus software and licenses are available without any problems. The vendor or distributors of the third software component (CFD code Fluent) are not part of the 413 consortium – and thus access to software and licenses needed to be organized separately in bi-lateral negotiations between Fraunhofer SCAI and Ansys. During the work for 413 several formal and license issues had to be solved before the technical setup could be used properly.

Manufacturer Types
In that target market we face two major different type of end-users:

- SMEs:
    - Only have limited simulation capabilities and experience (license, hardware, experts) Design is driven by experience and basic calculation methods (tables, etc) or simpler models (Matlab, Comsol, etc)
    - Not much time left to investigate in new analysis or design methods
    - They need a full service solution including long—term a customised and cloud based simulation environment
- Large Companies
    - Simulation experience and capabilities (i.e. Fluent, other EMAG codes) often available
    - Already have quite concrete ideas on what they would like to simulate
    - In most cases these companies already have partially prepared standalone models or benchmarks for a later co-simulation
    - Looking for an efficient way to test their ideas
    - They need an evaluation environment as decision support for a later purchase of new software

4.13.2 The industrial setting

A fast growing number of process engineers for manufacturing and production applications accept and know that they should enhance their currently used simulation environment with multiphysics features and tools. However, due to the daily workload on most engineers in industry the available time for running an evaluation of new methods and software solutions is very limited. Those engineers will only start software evaluations if they can be sure to get a valuable feedback from their tests after a very short time (i.e. few days).

A very basic but practically important hurdle for software tests often are the initial preparation steps - in the area of multiphysics simulation a specific complication is the co-simulation (coupling) of codes provided from different sw-vendors:

- software ordering,
- software download and installation,
- licensing,
- software configuration,
- integration in the user environments,

- availability of basic tutorials,
- obtaining help and support from multiple ISVs,

Although these initial preparation steps are just ‘stupid IT issues’ -- for a CAE engineer they often may lead to a termination of the evaluation if the in-house IT arena is not well prepared for such software tests.

4.13.3 The improvements in the industrial process through the use of an HPC Cloud

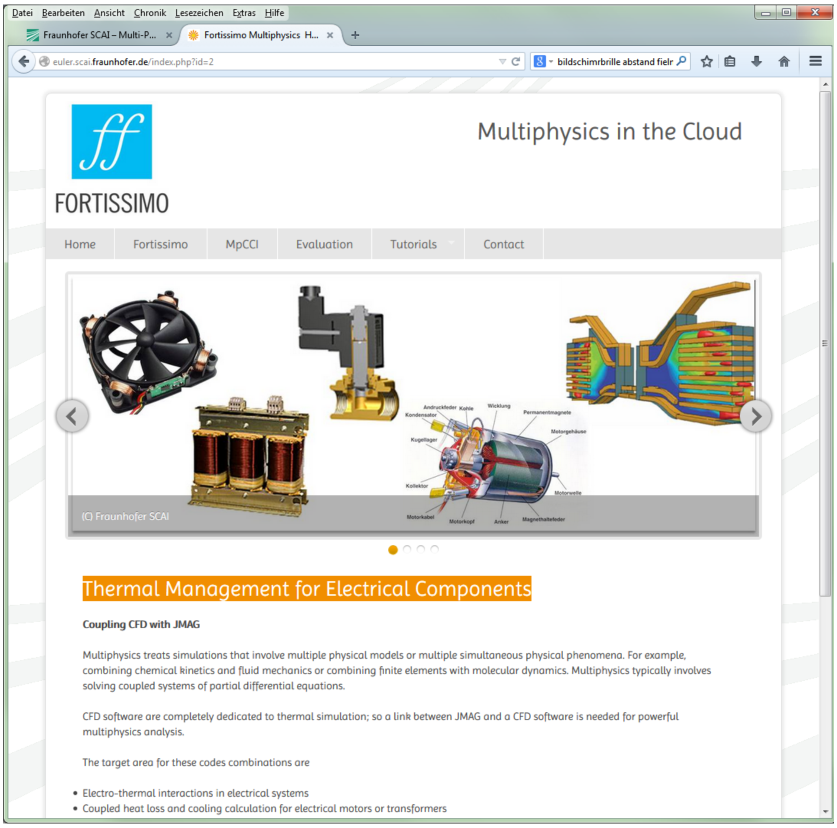

The idea of experiment 413 is to provide multiphysics evaluation capabilities and related support services as a remote cloud service from HPC-centres. Such services will provide access to various simulation codes, an open code-coupling environment [MpCCI], a valuable and growing number of tutorials, and enough CPU resources to run even bigger jobs.

Interested end-users from the manufacturing and production area may apply for an evaluation account at the HPC-centre. As soon as the account is available the engineer can start by going through the available training tutorials. Later on he can upload his own basic model data (previously used in mono-physics simulations) to run more realistic cases. Finally the HPC centre also offers resources for large and large-scale production jobs.

The overall approach of this experiment thus is not a solution for a single and specific end-user demand but to provide a more open and general way to promote the new technology ‘coupled multiphysics simulation’ to the market. Following that the focus of our work is on the evaluation of new cloud or web-based services which give multi-physics new-comers a more easy access to coupled software on appropriate hardware. Experiment 413 is a multi-CAE application, i.e. coupling and adjusting three soft-ware/licenses (1 proprietary, 1 open source, and 1 commercial) available in HPC/Cloud, includ-ing the development of cloud interface, to model, simulate and visualize a complex problem.

There are only a few steps for an interested engineer to get all information about the demo environment and to start his own evaluation:

1. More background information on ‘Thermal Management for Electrical Components’ and the concepts of MpCCI co-simulation
2. Detailed information on how to request for an evaluation
3. A selection of predefined tutorial cases which can be used online
4. Upload mechanisms to provide own models to the modeling experts at SCAI and Powersys

5. Parameter selection job submission
6. Retrieve simulation results and view by MpCCI Visualiser (available via download)

4.13.4 The cost benefits of using an HPC Cloud

In the mid-term 'Multiphysics services in the cloud' can be offered by HPC centres and ISVs as a complementary solution besides the traditional end-user specific rental / perpetual licenses. For engineers who need such high-end multiphysics capabilities only few times a year a 'pay-per-use-model' would be more acceptable than a full term annual license.

The major goal of this experiment is to demonstrate how HPC-cloud services can help to find and motivate new users for new multiphysics simulation solutions.

In the table below we have made a very exemplary calculation of hardware and software costs for a typical CFD-EMAG coupled model (we use ANSYS tools here because their pricing model is public and somehow typical for most of the CAE tools):

- 1 core with ANSYS/EMAG
- 8 cores with ANSYS/Fluent
- Duration 24h

Costs per parameter combination: 24h runtime; 1 core ANSYS/EMAG, 8 cores ANSYS/Fluent				
	€ / h*core	# cores	time[h]	Fees [Euro]
Pay-per-Use license for EMAG	20,00 €	1	24	480 €
Pay-per-use license for CFD base (incl. 4 cores)	20,00 €	1	24	480 €
Pay-per-use license for CFD HPC (4 additional cores)	2,50 €	4	24	240 €
Pay-per-use hardware Intel® Xeon® E5-2690 v2	0,15 €	9	24	32 €
TOTAL				1.232 €

Software license fees as listed by eCADFEM:
http://www.ecadfem.com/3403.html
Hardware fees as listed by CPU24/7:
http://www.cpu-24-7.com/en/resourcearea/overview-hpc-on-demand

These calculations for the software licenses differ strongly between the various CAE venros and their codes:

- CDadapco offers so called PoD licenses for STAR-CCM+ where the user can buy prepaid hours of license time – but with unlimited parallel capacities. Licenses are managed by a centralised license server from CDadapco.
- Numeca offers a very basic scheme where the core-hours for its CFD-codes FINETurbo or FINEOpen are used. Licenses are managed by offering HPC centre (eg CPU24/7)
- OpenFOAM can be used without any license fees – but the user is responsible for code adaption, extensions and maintenance.

If we would add any human service support to such a case some standard work fees of ~700-1000 Euro per day need to be added. It is obvious that (at least in the area of commercial CAE tools) the fees for software licenses will dominate all other fees – costs for hardware seem to be neglectible (from a user's point of view).
If compared to standard rental licenses such 'pay-per-use' fees are cheaper as long as you don't use more than 1000-1500 hours. Beyond that usage a standard rental license seems to be

more suitable. (*These numbers are only first estimations and by far not complete or valid for all CAE SW vendors*)

4.13.5 Any Other Business benefits

As experiment 413 does not target a single and specific engineering application the concrete benefits (in numbers) depend on the concrete demand of the end-user who is using the environment. The major advantage for the multiphysics newcomer will be the general possibility to test a new simulation environment through a cloud service – instead of installing new and complex software packages locally just for an evaluation run.

4.13.6 HPC as opposed to standard cluster offerings

In experiment 413 we target at engineers who are new to the area of multiphysics simulation. Our cloud-HPC offer wll enable to run some smaller tutorial cases and (depending on the user's experience) some extended user models. In both cases the number of required compute cores will be in a lower range; most cases only with 8-16 cores, some extended up to 64 cores.

4.13.7 Feedback to Work Package 2

Like for many other cloud-based services also WP413 would benefit from standard user management solutions:

- User registration
- User check
    - with respect to any kind of export regulations
    - to verify if a new users already is a licensed user of one of the commercial simulation tools (would need some feedback loop to these [non FOERTISSIMO] ISVs)
    - for any other type of relevant user-specific information
- Account creation
    - Access rights to specific tool environments
    - Access to pay-per-use licenses or Integration of own user-purchased licenses
    - Further access rights to shared data (in cooperation projects, user company data)
    - User groups
    - etc
- Pricing Information
    - Standard configuration & price table for sw-licenses and hardware usage
- Requested configurations
    - Required pre-processing and simulation tools
    - Data volumes
    - Frequency and response times for 'interactive usage requests'
- Access configuration
    - listed IP ranges which can get access to HPC systems
    - or just login data from anywhere (most suitable for short term end users)
    - local browser specification and plug-ins
    - local ssh, VPN-tunnel or similar installations
    - any other technical access related details

The target applications in WP413 will use standard commercial packages for EMAG and CFD simulation. So besides the above mentioned registration and account management details the most important requirements are on the user-interface and workflow management

- Availability of any kind of web-based pre-processing tools (e.g. Ciespace, SimScale solution, etc)
- Availability of CAE-preprocessing tools which can be used on the HPC login nodes (the target user might not have a simulation tool nor the related pre/post-processors)
- User friendly way to define a simulation workflow: *CAD ==> Meshing ==> setting BC and parms ==> starting simulation ==> post-processing ==> visualisation ==> reports*

4.13.8 Feedback to Work Package 3

==> same as in chapter 7 "Feedback to Work Package 2"

4.13.9 Input to Work Package 8

- Powersys got MpCCI distributor – December 2013
- NAFEMS UK Conference – May 2014
- Webinar: Why do coupled CFD-JMAG - July 2014
- JMAG UGM Frankfurt - Sept 23-24, 2014

4.13.10 Input to Work Package 9

Fraunhofer SCAI has its own local partner licenses for ANSYS/Fluent to be used for interface development works and MpCCI-related customer support. During WP413-installation phase the partners tried several options to enable license access on the HLRS compute systems:

- **Using HLRS ANSYS/Fluent licenses**: Access to HLRS's Fluent licenses is limited to users from HLRS or University of Stuttgart. The EULA between HLRS and ANSYS does not allow any further access to external users (see also https://wickie.hlrs.de/platforms/index.php/Application_software_packages)
- **Re-routing of SCAI's partner licenses**: SCAI is a member of ANSYS's partner system and has access to Fluent licenses for its MpCCI development. One of these partner licenses is dedicated to a mobile Laptop from SCAI which can be used outside the Fraunhofer campus Birlinghoven. SCAI and HLRS tried various ways (tunneling, etc) to re-route this license (at least for first installation tests) to HLRS/HERMITE. For various technical reasons (incompatible firewall policies at HLRS and SCAI, etc) this approach did not result in a technically stable solution.
- **Extending SCAI's eCADFEM IP range**: Besides its partner licenses SCAI also uses a regular eCADFEM (http://www.ecadfem.com/) license-on-demand since a number of years. These licenses allow an easy pay-per-use model for ANSYS/FEA (also CFD and EMAG) in research projects at SCAI. Access to these licenses is controlled through IP-ranges or numbers in SCAI's internal network – various of SCAI's workstations have access to eCADFEM license server. SCAI discussed with CADFEM to extend this IP-range by some HLRS-servers – but CADFEM and/or ANSYS did not provide any solution here.
- **Purchase of a dedicated project license**: Finally SCAI had to purchase a new license from ANSYS which may be used at a remote site – i.e. HLRS. SCAI can use this license within the FORTISSIMO project and run the Fluent code on HLRS systems. The only limitation is that only SCAI people may use this license – i.e. the <START>-button for Fluent calculations may only be pushed by SCAI people.

A very general question for SCAI and potentially for many other CAE service providers is:

"**How can an independent engineering service provider use a commercial sw-license from a 3rd party sw-vendor**

for his own Cloud-HPC-Services??"

4.13.11 Any further input to Work Package 10

Powersys and Fraunhofer SCAI have organised a joint webinar where the new solution JMAG-MpCCI-CFD was promoted to ~20 participants.

Webinar	2014/07/03
JMAG	Webinar: Multiphysics analysis with JMAG (EM) - MpCCI - CFD software

4.14 Experiment 414: Calculating the Properties of Hazardous Materials

4.14.1 Overview of the experiment

The Chair of Thermodynamics and Energy Technology (ThEt) at the University of Paderborn (UPB) develops the molecular simulation tool *ms*2 that is aimed at thermophysical properties of condensed fluids. This simulation software is currently used by ThEt for commercial applications on its workstation computing systems. Based on molecular force field methods,

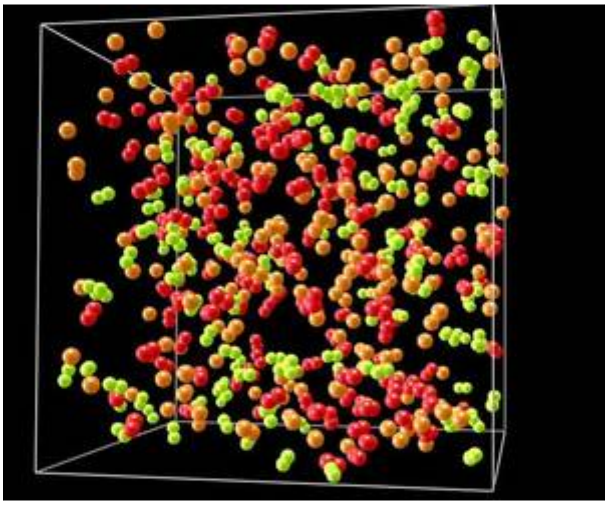

nowadays powerful predictive methods exist that yield thermophysical data which are the basis for the design and optimization of chemical engineering processes. Typically, the chemical industry measures the required data experimentally, however, if they are needed for hazardous substances (explosive, toxic or mutagenic), the associated costs may be prohibitive. In such cases, molecular dynamics (MD) and Monte Carlo (MC) simulations on the basis of optimized force fields are very attractive alternative routes. To carry out such predictions with molecular simulation, typically a larger number of simulation runs is necessary. The advantage of cloud-based simulations for this experiment is, that the individual simulations are mutually independent and that therefore, the whole experiment features very loose coupling. Therefore, this experiment can make use of cloud infrastructure, which is a lot less cost intensive than conventional HPC resources. Furthermore, the experiment will feature dynamic adaptivity to the obtained results – zooming in on relevant areas of the phase space. Here, the dynamic allocation possibilities in clouds will be very powerful.

The company LONZA is a leading supplier to the life-science industry and acts as an end-user in the experiment. The prediction of thermodynamic properties of hazardous fluids is of significant interest to LONZA. The High performance computing centre (HLRS) of USTUTT acts as HPC expert, HPC provider and host centre in the experiment.

A main goal of this experiment is to introduce a fairly new technology to a wider public. The cloud-based molecular simulation will lower the inhibition level of the prospective users because they are far more familiar with web interfaces than with supercomputer-terminals. Since the user is guided through the molecular simulation process, less expertise is required. Furthermore, this project will simplify the workflow of molecular simulations since less human interaction is needed. The chemical industry will have the greatest benefit from this project due to the fact that they are enabled to investigate substances that would only be examinable experimentally at considerable expense.

4.14.2 The industrial setting

The Chair of Thermodynamics and Energy Technology (ThEt) at the University of Paderborn develops the molecular simulation tool ms2 that is aimed at thermophysical properties of condensed fluids. This simulation software is currently used by ThEt for commercial applications on its workstation computing systems. Based on molecular force field methods, nowadays powerful predictive methods exist that yield thermophysical data which are the basis for the design and optimization of chemical engineering processes. Typically, the chemical industry measures the required data experimentally, however, if they are needed for hazardous substances (explosive, toxic or mutagenic), the associated costs may be prohibitive. In such cases, molecular dynamics (MD) and Monte Carlo (MC) simulations on the basis of optimized force fields are very attractive alternative routes. To carry out such predictions with molecular simulation, typically a larger number of simulation runs is necessary.

4.14.3 The improvements in the industrial process through the use of an HPC Cloud

At the moment, the efficient usage of *ms*2 requires expert knowledge. The User has to be familiar with the Linux terminal and the HPC-environment. Furthermore, there is no graphical user interface present. The Cloud-based molecular simulation will lower the inhibition level of the prospective users because they are far more familiar with web interfaces than with supercomputer-terminals. Since the user is guided through the molecular simulation process, less expertise is required. Furthermore, this experiment will simplify the work flow of molecular simulations since less human interaction is needed. All in all, this fairly new technology can be introduced to a wider public and will arouse the interest of more SME in the chemical industry resulting in new business opportunities for the technology supplier.

In order to inter- and extrapolate the simulation data, the Cloud-based molecular simulation is able to create correlations (**E**quations **O**f **S**tate, EOS) in a short amount of time automatically. These EOS are very helpful since they contain all the thermodynamic data in a consistent manner and allow adapting the data for the specific purposes of the SME. The manual creation of such EOS with experimental data is normally very tedious and time-consuming.

Another advantage of Cloud-based simulations for this experiment is, that the individual simulations are mutually independent and that therefore, the whole experiment features very loose coupling. Therefore, this experiment can make use of cloud infrastructure, which is a lot less cost intensive than conventional HPC resources. Furthermore, the experiment will feature dynamic adaptivity to the obtained results – zooming in on relevant areas of the phase space. Here, the dynamic allocation possibilities in clouds will be very powerful.

4.14.4 The cost benefits of using an HPC Cloud

Based on molecular force field methods, nowadays powerful predictive methods exist that yield thermophysical data which are the basis for the design and optimization of chemical engineering processes. Typically, the chemical industry measures the required data experimentally, however, if they are needed for hazardous substances (explosive, toxic or mutagenic), the associated costs may be prohibitive. In such cases, molecular dynamics (MD) and Monte Carlo (MC) simulations on the basis of optimized force fields are very attractive alternative routes.

All in all, the use of HPC- and cloud-based molecular simulation yields a considerable monetary advantage for chemical industry companies. Experimental pure component densities cost around 2700 € per substance for a very limited temperature and pressure range, when bought from an external supplier. Compared to that, 60 molecular simulations carried out in the entire fluid region up to arbitrary high pressures will cost around 1600 € yielding not just the density but every static thermodynamic property simultaneously.

When it comes to mixtures, the difference in prices is more extreme. 60 experimentally measured gas solubility data points of a binary mixture can cost up to 50000 €, while the cost for the molecular simulation usually increases by a factor of two compared to a pure component. It has to be noted that the prices for the experimental data only apply for moderate conditions and non-hazardous substances. Measurements at high temperatures or pressures can be much more expensive or even impossible to conduct.

For Lonza, which does the measurements for the design of their apparatuses by itself, massive savings can be expected. Considering a thermal separation task of a mixture of moderate difficulty, the experimental measurements can easily take up to half a year to be done. Molecular simulation yields the same data in about a month and only few experiments have to be conducted afterwards in order to verify the results. A rough calculation of the savings can be found in Table 1, it can be seen that a molecular simulation study is around 80,000 € cheaper, which is a substantial part of the total cost of a distillation column (~1.5 million €). Note that, depending on the competitive situation, around 3-5 projects of this size are carried out per year.

4.14.5 Any Other Business benefits

For a chemical engineering company like Lonza, the time-to-market will be decreased due to the faster gathering of fluid properties for distillation and absorption significantly.

4.14.6 HPC as opposed to standard cluster offerings

Due to the volume and complexity of the calculations with the *ms*2 simulation code, researchers typically rely on high performance computing (HPC) infrastructure. While a single calculation is relatively time intensive, the calculation of an EOS requires a large amount of independent simulations to be performed. Thus, besides the scalability aspects of the individual application, multiple instances will be executed typically in an embarrassingly parallel fashion. Therefore, the simulations require scalability in both horizontal and vertical direction, i.e. with respect to number of instances and their performance.

A highly accurate calculation of thermodynamic properties at a given state point currently requires around 20 hours on 4 computing nodes (i.e. 32 cores). In order to acquire enough data to generate the full information set as needed for the accurate prediction of the EOS, around 100 state points need to be evaluated. This means that even with exclusive access to small compute cluster, the calculation would take up to 2000 hours (close to three months).

4.14.7 Feedback to Work Package 2

- A basic access authentication with a username and password would be great in order to associate experiments with users.
- In order to distinguish between user-induced errors or system errors an event-oriented audit trailing is needed.
- As a user a workspace with private data and a history of experiments used would be beneficial.
- Results from previous simulations should be storable for some time

4.14.8 Feedback to Work Package 3

- Additional to the included help page, a method for contacting experts has to be present. This could be a forum for example.
- A possibility to upload tutorial videos could be helpful

4.14.9 Input to Work Package 8

(Exploitation and Sustainability)

The Fortissimo marketplace concept needs to clarify and refine possibilities for service offers and accounting as well as the legal framework surrounding it.

4.14.10 Input to Work Package 9

(Business Models for ISVs)

Load-balancing over several HPC providers very attractive in order to reduce the overall time of a simulation and improve time-to-market for its end-users noticeably. The challenge in regards to a load-balancing approach is to spread the calculations over providers and divide them the way the simulation completes as fast as possible, i.e. reduce number of nodes at a certain location if a job can then start faster: Meta-Scheduling in the Fortissimo HPC-cloud under consideration of current queues and costs associated (a fast queue can be more expensive).

4.14.11 Any further input to Work Package 10

(Dissemination and Training)

- Scientific Paper covering the HPC-cloud based approach for ms2 is the pipeline, intended to be published in the journal "chemical theory and computation"
- Training offers for ms2 simulations planned

Table 1 Calculation of the savings for experiment 414

Experimental study			
	cost/day	*#days*[1]	*total cost*
Measurements	1,000 €	100	100,000 €
(cumulated cost)[1]		(5 months)	
SUM			**100,000 €**

Molecular Simulation study					
	cost/CPUh	*#Sim*	*Duration (h)/Sim*	*#Cores/Sim*	*total cost*
HPC calculations	0.15 €	200	20	16	9,600.00 €
			20		
	cost/day	*#days*			*total cost*
Workload personnel	550 €	4			2,200 €
(for setting up Simulations)					
Subsequent measurements	1,000 €	5			5,000 €
(for verification)					
Software/Support and Consulting					Subject to negotiation
SUM					**16,800 € +X**

[1] Estimated cost from personal communication with LONZA

[2] Assuming that one month is 20 business days.

4.15 Experiment 415: Design of Moulds for Casting

4.15.1 Overview of the experiment

IMR produce foundry equipment for brass alloys and bronze. The plants are used in 80% of the cases for the production of taps and valves and in 20% for other high performance artefacts.
The moulds mounted on the installation are filled by injecting the liquid metal at temperatures of about 1000 ° C. The inside of the castings is obtained by preformed cores in sand.

4.15.2 The industrial setting

The success of the entire system depends on the quality of the piece obtained.
Therefore essential to design the moulds and the casting channels so as to obtain a laminar flow of the liquid metal and a cooling gradient constant throughout the final piece to avoid cracks and defects.
Currently, the design of moulds and filling channels is entrusted to the experience of mould makers in collaboration with experts from the foundry.
Generally, the exchange of experience obtained good results, but it is often necessary to modify the mould several times and repeat tests before committing to production.
The application of simulation programs with HPC avoid changes of the mould prototype obtained empirically by trial and not with the objective data of the simulation.
Use the simulation with HPC reduces by 20% the time of development of the mould and 20% the cost of testing before mass production.

The average cost for a set of moulds for a product designed and tested, it is currently about € 40,000.00. It is therefore save about € 8,000.00 per set and 3 weeks of testing and modifications.

Cost:		
Mould drawing	€ 8.000	
Mould Material (CuBe)	€ 6.000	
Core box drawing	€ 5.000	
Core box material	€ 4.000	
Testing (3 Days)	€ 3.000	(€ 1.000 saving test Time)
Drawing Modif. after testing	€ 2.000	(€ 1.000 saving drw. Time)
Mould/core modif.	€ 8.000	(€ 4.000 saving modif. Time)
Final test and sample	€ 5.000	(€ 2.000 saving final test Time)
Total	€ 41.000	€ 8.000

€ 41,000 is the current cost for with standard development of a prototype mould.
€ 8,000 is the savings generated by the use of casting simulation with HPC.
€ 2,000 is the cost for each set of mould done with HPC simulation. (HPC rental cost and personnel involved).

IMR has about 8 sets of moulds per year to develop (save € 6.000 x 8 = € 48.000)
Total saving per year € 48.000 by using casting flow simulation.

4.15.3 The improvements in the industrial process through the use of an HPC Cloud

SME's operating in our sector using the traditional method of test-error-correction-new test. Very few use simulation programs in the market because of the high cost and difficulty of use of the software.
In addition, the processing times with normal processors are very long, about 4-6 hours for each change of parameters.
The possibility of using a HPC with open source programs allows a considerable saving of time and investment.
The participation of the IMR program "Fortissimo" demonstrated that an SME with limited knowledge of the world HPC can get good results from the simulation by reducing design time and test the moulds.
The competitiveness and the impact on the market will bring great benefits in terms of reduced lead times, reduced development costs, expansion of the market thanks to the collaboration at the European level of SME's operating in the field of mould construction.

4.15.4 The cost benefits of using an HPC Cloud

The use of HPC Cloud, with open source software, allowing a saving on the purchase of hardware and software currently on the market (no HPC), evaluated by simulation programs like Magma or Procast approximately € 100,000.00.
The cost of using HPC, rental, only for the time necessary for the development of moulds particularly complicated, allows for hardware and software always update at affordable cost.

4.15.5 Any Other Business benefits

Benefits:
The rental of HPC and open source software simulation:

- Reduces the development time of the moulds
- Reduces the cost and the number of tests before putting into production
- Improving the competitiveness of the SME with the image of high-tech and service offered to the customer
- Improves collaboration between those who designed and built the mould and foundry expert, technologist process
- SMEs that offer the service of the simulation will have many other opportunities to offer turnkey systems not only to its own niche industry, but also to new markets and / or products.

The simulation offer service can improve numbers of set mould possible sale per year. (Estimated n°4 x € 40.000= € 160.000).
New company profile, with high tech. image can improve sales of turnkey plant. (Estimated n°2 x 400.000= € 800.000)

4.15.6 HPC as opposed to standard cluster offerings

IMR had no experience in this field since it has started using HPC technologies only thanks to the Fortissimo project.

As we explored during the experiment, the code has a good scalability if appropriate preconditioning is taken. Nonetheless, the simulation performed on a single CPU can take up to more than a day for completion.

In order to complete the calculation in a timely manner, the use of homogenous, many-node infrastructure is required, e.g. the same simulation on a 42-nodes infrastructure (84 CPUs) has the execution time reduced to less than 3 hours

For this reasons an HPC infrastructure is required instead of a standard cluster system.

4.15.7 Feedback to Work Package 2

The experiment has identified the following requirements:

- need for interface offering to access appropriate infrastructure (both HW and SW)

- introduction to HPC workflows is welcome (how to deploy jobs, MPI coding…)

4.15.8 Feedback to Work Package 3

The technical work should resolve the requirements from the paragraph above.

4.15.9 Input to Work Package 8

The pay-per-use option is the most appropriate usage model for this experiment.

4.15.10 Input to Work Package 9

Experiment 415 has used an Open Source and free software solution. This is due to the fact that comparable licensed solutions would represent a major cost in the presented workflow if they were to be implemented.

Constraints to using commercial software are described in section 4.

4.15.11 Any further input to Work Package 10

We are confident that further training opportunities will help to disseminate the information about the Fortissimo project. Training events will be especially beneficial to non-expert HPC users that lack the experience and confidence in using such systems.

4.16 Experiment 416: Structural Crash-Test Simulations

4.16.1 Overview of the experiment

The objective of the experiment is to showcase the use of cloud based HPC services to investigate the behaviour of automotive part components involved in structural crash test simulation using a commercial software package.

Currently, market requirements related to vehicle weight reduction and cost cutting are driving the industry to accelerate their innovation and to introduce new design and new material and manufacturing processes. The challenge for supply chain usually SME's, is to handle conflicting requirements and bring revolutionary changes to vehicles, while at the same time cutting the development costs and time drastically.

The software provider in this project is ESI-GROUP with its VPS (Virtual Performance Solution) Software Suite. The experiment will provide them with a version of their software that can be used in a cloud environment, together with an insight into how this could be exploited as a future business model.

The experiment will address porting the commercial software and testing it on the Fortissimo cloud-based HPC infrastructure. ESI-GROUP and the Fortissimo Competence Centres will work together to produce the cloud-based version of the software that meets the needs of the end-user. An evaluation of the viability of this service on commercial HPC clouds will be carried out, and possible business models will be proposed.

4.16.2 The industrial setting

The Green Vehicle initiative is a major driver for the automotive industry to answer environmental requirements and regulations. This challenge calls for disruptive innovation in vehicle design and material. Steel remains the main cost effective material for Body In White (stage in automotive design in which a car body's sheet metal components have been welded together) mass production for the foreseen future. This challenge is a good business opportunity to emerge new industry leaders like GESTAMP, who have developed a new grade of steel called Press Hardening with the associated manufacturing processes. It is one of the major innovations introduced to the market with success to answer the weight reduction of the vehicle with recognized WW advantages.

4.16.3 Improvements in the industrial process through the use of an HPC Cloud

However, the industrialization and the customization of this innovation also required a new design and development process in order to enable Engineering & Business trade-offs between rather conflicting requirements such as weight reduction, performance requirements and manufacturing constraints.

Business scale-up is one of the main competitiveness factors. This required mastering an efficient and flexible process for customization and localization. Divergence of the models are inherent phenomena with the current silo approach due to the complexity of the interactions; CAD-CAE, CAE-CAE.

To answer this challenge, we will apply Virtual Prototyping with a rather holistic view, several full vehicle simulation models all based on one unique central Body In White subsystem, named the Single Core Model approach. This disruptive approach will enable to support collaborative decision making process during the Product Development phase;

including engineering stakeholders from 3 different disciplines: Crash & Safety, Durability, NVH. Figure illustrates the drawback of the silo approach and the single Core Model advantages.

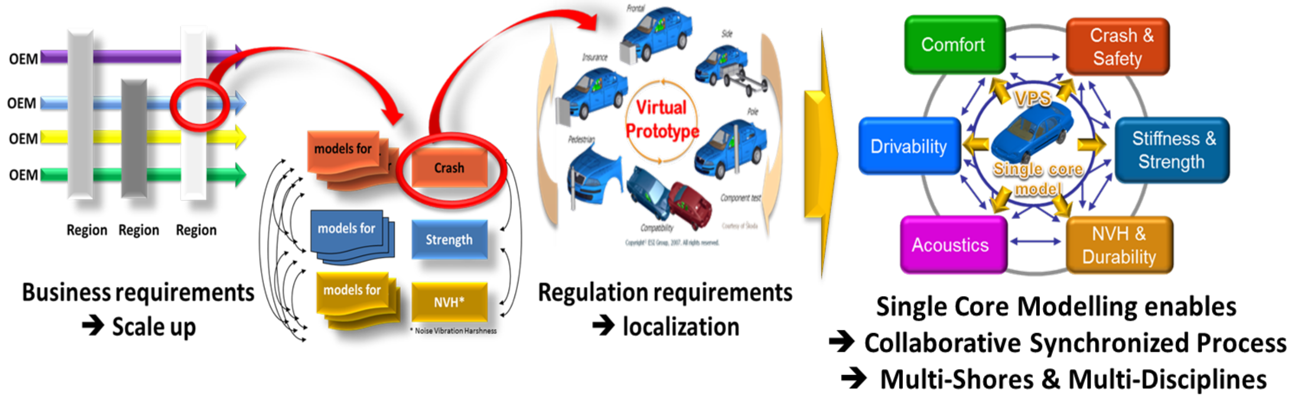

Single Core Modelling will reduce the CAD-CAE synchronization efforts, and eliminate the conversion & consolidation between different CEA-CAE Models for Design reviews and trade-offs. Today's approach is a very consuming task and source of errors. The Single Core Model approach will replace the human efforts by computing resources on low value tasks. It will enable better use of the engineering expertise to improve the product quality.

Therefore, this innovative development methodology calls for extensive computer resources, which was considered as a major obstacle to go through this experience. Cloud Computing seems to be the ad-hoc solution as of today for fast growing manufacturing European-based companies like GESTAMP. This is a major enabler; it offers the necessary flexibility to access to HPC resources when it is needed at an affordable cost.

This breakthrough will help to avoid last minute engineering design changes, and it should result in improved competitiveness in order to reduce the development period, to shortening time to market, to improve quality at competitive pricing

4.16.4 The cost benefits of using an HPC Cloud

The Benefits will be measured in terms of Time saving thanks to the leverage of HPC resources which enable the introduction of the Single Core Modeling as new collaborative engineering approach to support Decision making process.

Single Core Modeling enables the handling of only one Body In White numerical model for several applications (four applications in our example). Thus, once concepts of Body In White have to be updated, instead of implementing the changes in the 4 former models, only one single core model needs to be modified. Time spent on setting up and afterwards updating the model is then divided per 3.5 to 4 (see table below). Indeed, the simulation model setup for one application last approximately 50 hours, so 200 hours for all four applications. With the Single Core Model, only 60 hours are needed to build up the Core Model (10 hours more due to the higher complexity of the system). Moreover, as only one Core Model is used, the risk to have four divergent models, in term of concepts, is consequently reduced to zero. The compute cost difference between a classic crash model and one using the Single Core approach is close to zero. Thanks to this innovative approach, the expert team involved in such tasks will be able to spend more time in developing and optimizing lighter and better cost-effectives solutions for Body In White.

Time Saving / Single Core	Application dedicated models				Single Core Model
	Crash		Stiffness	NVH	Crash/Stiffness/NVH
	Front	Side			
Set Up	50h	50h	50h	50h	60h
Maintenance/Evolution	1-4h	1-4h	1-4h	1-4h	1-4h
Model Convergence percentage for several test cases	30%	30%	30%	30%	100%
Synchronization CAD-CAE	Need to do it 4 times				1 time
Synchronization CAE-CAE	Need to do it 4 times				1 time

4.16.5 Any Other Business benefits

The use of HPC cloud services will also be a huge benefit for multi-site companies. Indeed, Gestamp R&D Centers located in Paris (France), Luleå (Sweden) and Barcelona (Spain) have increasing needs in handling projects over several locations. Therefore, the HPC enables to access and share the data, to run simulations and to analyse the results in all the relevant locations through a conventional web browser. In summary, this HPC cloud application will enable a more efficient project handling by reducing data transfer time and by improving the communication between the different stakeholders of a project.

4.16.6 HPC as opposed to standard cluster offerings

Extreme factory is a pure HPC cloud offering, with optimized scalable architecture, dedicated management software and fast InfiniBand interconnect. It is also built from scratch with security and ease-of-use in mind.
The VPS software could be run on less sophisticated architectures, but with less efficiency, less security, and a perceived cost which does not take into account the time and efforts to setup an operational HPC cluster on an Amazon-type infrastructure.

4.16.7 Feedback to Work Package 2

ESI-GROUP and Gestamp have no further feedback to Work Package 2 having already supplied initial requirements in response to the requirements gathering questionnaire.

4.16.8 Feedback to Work Package 3

N/A

4.16.9 Input to Work Package 8

There is no formal plan at this time to include VPS offerings on the Fortissimo Market Place

4.16.10 Input to Work Package 9

Evolving from simulation per domain to a holistic view of virtual prototyping is an industry challenge. It required an in-depth review and upgrade of the industrial process. The Fortissimo project helped ESI Group to validate an industrial service approach and to adapt the methodology for Gestamp's specific business context.

In addition to the Product portfolio, this customized service will be available for the industry partners during the deployment phase. It will enforce their smart specialization, in particular the Mid Cap and the SME's.

A prototype of Cloud Based Licensing Management has been developed in Fortissimo. It enables the appropriate efficiency in term of flexibility to access a wider scope of applications; when needed with the necessary computing resources. The early experiments show that this represents a new business opportunity with a new Business Model with significant potential impact. It lowers the barrier of the initial investment to implement the Virtual Prototyping methodology for OEM, Mid-Caps, SME and also start ups.

4.16.11 Any further input to Work Package 10

The WP416 experiment was presented at several occasions :
Participation to I4MS conference dedicated for SME in Berlin – June 2014
Gestamp paper at the user conference for ESI Group – EGF, May 2014
ISC'14 : Mark Parsons presentation June 2014

More communication is planned around this experiment :
ESI Talk : success story in 2015
NAFEMS 2015 : Joint paper
Gestamp : internal communication xxx

4.17 Experiment 417: Optimised Design of a Supercar

4.17.1 Overview of the experiment

Experiment 4.17 "Cloud-based Computational Fluid Dynamics" belongs to the first tranche of experiments to drive the development of the Fortissimo infrastructure and focuses around simulation processing for Computational Fluid Dynamics (CFD) and Computer Aided Engineering (CAE) which is a vital domain for both big industrial domains such as aeronautics and defence, automotive manufacturing, and SMEs, ranging from complex manufacturing of goods, to consulting services in engineering.

The CFD market can be characterised with customers having two distinctly different drivers: heavy CFD users with high budgets (typically OEMs) who are moving into more complex physics with improved accuracy with associated business impact and the followers (typically SMEs who have modest CFD resources). The latter typically have engineers with limited skills in CFD and use commercial software with limited licenses for budget reasons, often not adequate to the science they attempt to simulate.

Both market groups are linked by a latent need for access to scalable, cost-effective CFD. Cloud-enabled technologies can fulfil that need. The case for cloud computing in CFD simulation is clear: it will not only increase the cost effectiveness but will accelerate the uptake and application of more complex CFD, that in many cases is required to capture the science influencing product development.

Experiment 4.17 will use the Fortissimo infrastructure to port existing simulation code solutions over HPC cloud infrastructures to enable end users to have instant access to a "one-stop-shop" which will offer simulation services on demand, while users will be able to obtain the results of their simulations in less time, and with a pre-defined cost per simulation (based on various criteria). More specifically, Experiment 4.17 will pursue the following objectives:

- Validate that the Cloud "one stop shop" solution is accessible, secure, technically sound and transparent in its economics for CFD End-Users
- Enhance usability for end-users to set up, submit and review their results in a logical and straightforward manner
- Provide access to user account status so it is clear what costs have been incurred and estimate likely further costs against specific tasks.
- Obtain feedback from end-users that this cloud solution is a viable supplement or replacement to their existing HPC resources

The experiment brings together CINECA, a leading HPC centre, ICON, a ISV/business owner/simulation code support provider, NTUA, a research organisation with expertise on interoperability, system integration design and decision support systems, and KOENIGSEGG, a niche, high performance sports car manufacturer, which requires greater access to affordable, scalable, HPC resources for its CFD requirements, as ICON's end user.

4.17.2 The industrial setting

Computational Fluid Dynamics (CFD) is widely used in industry; from aerodynamic performance improvement to duct system optimization or passenger cabin comfort analysis. CFD can bring significant cost savings to the product development phase. As a result, it is currently the fastest growing segment in Computer Aided Engineering (CAE) simulation.

Despite the necessity of continual experimental testing, in order to validate the mathematical models used and quantify the error made, CFD brings indisputable benefits. For example, within the automotive industry one of the major advantages of CFD analysis is to compute and optimize the drag and down force components for different car configurations (pitch/yaw angles, additional aero devices, overtaking, etc.). These improvements then lead to fuel consumption reduction, range increase and/or overall improvements of the aerodynamic performance. Moreover, in the Formula One industry intensive application of CFD is carried out in order to reduce cost of wind tunnel testing. Models used to replicate real life cars with high accuracy, consisting of every geometric detail, rotating wheels and integrated components (such as heat exchangers, fans, condensers, etc.), can be very large. HPC Cloud solution has recently become a possible alternative to large clusters. Fully time dependant simulations of such models are very complex and the use of an HPC solution, with a suitable support of domain specific expertise, makes a significant difference. Past experiences have shown that the full aerodynamic design of a hyper car can be entirely conducted using CFD with minimal road/wind tunnel testing. Such a task in a production environment is possible only if tight time deadlines are met.

4.17.3 The improvements in the industrial process through the use of an HPC Cloud

The use of the ICON software via the Fortissimo HPC Cloud helps the prestigious car manufacturer Koenigsegg to reduce or even substitute wind tunnel testing in some circumstances. Accessing powerful computing resources remotely also reduces the hardware expenses and maintenance costs. In 2013 100% of the aerodynamic development of the Koenigsegg One:1 has been conducted using CFD. In less than eight months, hundreds of simulations to test various configurations have been carried out. The results were an impressive 250% increase in down-force with only 15% increase in drag at 250km/h with a 50% higher down-force at 440km/h maximum speed.

4.17.4 The cost benefits of using an HPC Cloud

Tests have shown that the use of HPC-based simulation using a Cloud and external expertise will lead to a return on investment in less than three months of production of a new car configuration. Significant costs can be saved and re-injected into other critical parts of the car development and production.

The benefits obtainable by the use of the Fortissimo HPC-Cloud have been quantified as a 5% saving in operational costs, a 30% saving in design costs, 50% in wind tunnel/physical testing a 60% saving in prototyping costs, and a 30% shortening of the time to market. Furthermore, savings in development were about €90K Euro/year on the design process, corresponding to a 1.5% reduction in overall development costs.

4.17.5 Any Other Business benefits

The provision of a Cloud-based HPC solution would offer significant innovation into the marketplace. In particular the innovation would be derived from the ability to assemble significantly cheaper compute farms for CFD analysis (combined with open source solving) than ever before. In addition, ICON has already demonstrated the maturity of automatic optimisation technologies when combined with open-source CFD solver technology, (industrially supported by ICON) can offer significant price point advantage over competition and existing manual CFD simulation methods.

On the other hand, customer value and benefits from HPC-cloud CFD simulation could be summarised as follows:

- Increased speed, as a result of obtaining simulation results within hours instead of weeks.
- Increased flexibility, occurring mainly through services utilization on demand and payment out of specific project budgets.
- Reduced cost, encompassing pay on demand capabilities, reduced depreciation costs, and more effective labour deployment.
- Improved quality, as a result of the production of an enhanced solution based on higher resolution models.
- Faster design processes, driven by optimization made possible by increased CFD simulation throughput.
- Increased competitive advantage on the basis of the above mentioned benefits.
- Lower entrance barrier for SME's approaching CFD.

4.17.6 HPC as opposed to standard cluster offerings

The equations underpinning a CFD analysis have to be solved by the computer in an iterative manner and require thus very powerful and costly computer infrastructures and very long processing times. Simulation of CFD models has traditionally lent itself to parallel or cluster computing (multiple CPU working on the same overall problem by distributing equations and models into 'sub' calculations). As hardware costs have plummeted over the past ten years, the real cost of CFD simulations may have been expected to decrease. However, although it is cheaper than ever to procure multiple CPU (e.g. Linux clusters), CFD license costs did not decrease. Thus, as a company adds extra CPU to its computer resources, the software vendors have priced their solvers on a, per CPU or banks of CPU basis. This has resulted in a high proportion of companies having to severely restrict the type or size of the computer simulations that they undertake.

The need for cloud-based HPC simulation is of utmost importance. In this experiment, by moving simulation of CFD models on a cloud based HPC environment, processing times can be significantly reduced. Moreover, the anticipated royalties and utilisation costs of such services can be drastically reduced, enabling the simulation code provider to reach to new markets, while SMEs that were unable to "purchase" such services can be given the opportunity to do so, moving from stand-alone and effort/resources consuming applications to more flexible SaaS solutions.

4.17.7 Feedback to Work Package 2

Experiment 4.17 has provided to Work Package 2, "Core Service Requirements Capture and Infrastructure Development", as described in Deliverable P417.1, detailed hardware, and software and license requirements, alongside with security and data management constraints and user management and accounting, as well user support needs for porting ICON's CFD simulation software over CINECA’s HPC cloud infrastructures.

4.17.8 Feedback to Work Package 3

Experiment 4.17 has provided up to now only indirect feedback to Work Package 3, through the requirement collection document provided to Work Package 2 (Deliverable P417.1).

It is however willing to be an early tester of the marketplace middleware and ready to port the developed framework to the Fortissimo middleware as soon it will be available.

4.17.9 Input to Work Package 8

Experiment 4.17 input to Work Package 8, "Exploitation and Sustainability", is collected in Deliverable P417.2 and comprises a preliminary business model analysis with regard to the commercial exploitation of the experiment's service offering. The analysis includes a detailed presentation of the services to be offered, sketching out four types of service provision, i.e. trial offering, bronze, silver and gold subscription, identifies the stakeholders involved, analyses extensively the CFD market, and exposes alternative business models, focusing mostly on the aspects of the service provision scheme and the pricing model.

4.17.10 Input to Work Package 9

Experiment 4.17 input to Work Package 9, "Business Models for ISVs", is collected in Deliverable P417.2 and comprises a preliminary business model analysis for an ISV like ICON to implement a SaaS service on the Fortissimo marketplace, discussing possible licensing and pricing models.

4.17.11 Any further input to Work Package 10

Experiment 4.17 has been featured as Experiment-of-the-Month” on the Fortissimo website.

4.18 Experiment 418: Design of Centrifugal Pumps

4.18.1 Overview of the experiment

Experiment 4.18 “Cloud-based CFD turbomachinery simulation” belongs to the first tranche of experiments to drive the development of the Fortissimo infrastructure and focuses around simulation processing for Computational Fluid Dynamics (CFD) and Computer Aided Engineering (CAE) which is a vital domain for both big industrial domains such as aeronautics and defence, automotive manufacturing, and SMEs, ranging from complex manufacturing of goods, to consulting services in engineering.

The CFD market can be characterised with customers having two distinctly different drivers: heavy CFD users with high budgets (typically OEMs) who are moving into more complex physics with improved accuracy with associated business impact and the followers (typically SMEs who have modest CFD resources). The latter typically have engineers with limited skills in CFD and use commercial software with limited licenses for budget reasons, often not adequate to the science they attempt to simulate.

Both market groups are linked by a latent need for access to scalable, cost-effective CFD. Cloud-enabled technologies can fulfil that need. The case for cloud computing in CFD simulation is clear: it will not only increase the cost effectiveness but will accelerate the uptake and application of more complex CFD, that in many cases is required to capture the science influencing product development.

Experiment 4.18 will use the Fortissimo infrastructure to port an existing commercial software (ANSYS CFD) over HPC cloud infrastructures and to build on top of that a customized solution to enable end users with limited knowledge of CFD and HPC to explore an appropriate parameter space of the design of a turbomachinery (centrifugal pump) to achieve the desired optimization level. The solution may be implemented on the Fortissimo marketplace as a service, to complement Enginsoft service portfolio. More specifically, Experiment 4.18 will pursue the following objectives:

- Validate that the Cloud "one stop shop" solution is accessible, secure, technically sound and transparent in its economics for CFD End-Users;
- Enhance usability for end-users to set up, submit and review their results in a logical and straightforward manner;
- Provide access to user account status so it is clear what costs have been incurred and estimate likely further costs against specific tasks;
- Obtain feedback from end-users that this cloud solution is a viable supplement to their existing HPC service portfolio.

The experiment brings together CINECA, the Italian HPC centre, and Enginsoft, the Italian leading consulting company operating in the field of computer aided engineering (CAE), virtual prototyping, advanced simulation, process integration and design optimization (PIDO) and, more generally, scientific IT targeted at the optimization of design and production processes.

4.18.2 The industrial setting

Centrifugal pumps are widely used in many industrial applications, from oil & gas to water treatment, automotive and home appliance. Such devices may be required to operate over a wide flow range and prediction of operating characteristic curves is essential for a designer. Unfortunately, all theoretical methods and experimental tests are unable to determine the source of poor performances. In these terms, Computational fluid dynamics (CFD) becomes an important and common tool for pump designers. Many tasks can numerically be solved much faster and cheaper than by means of experiments and, most important, the complex internal flows in water pump impellers can be well predicted. As a result of these factors, CFD is now an established industrial design tool, helping to reduce design time scales and improve processes throughout the engineering world.

4.18.3 The improvements in the industrial process through the use of an HPC Cloud

Numerical simulation of centrifugal pumps is not easy due to some difficulties: complex geometry, unsteadiness, turbulence, secondary flows, separation, boundary layer, etc. These aspects require a high fidelity CFD model, or, in other words, very fine computational grids and transient analysis. This approach is quite prohibitive for a typical SME, so the aim of the Fortissimo HPC Cloud solution is to provide an extremely attractive solution in terms of business and technological perspective for any SME which can hardly deal with high fidelity calculations by its own efforts.

4.18.4 The cost benefits of using an HPC Cloud

Tests have shown that the use of HPC-based simulation using a Cloud and external expertise will result in a return on investment in less than six months, this meaning that a single pump project with optimization performance can be run in a timeframe of six months, rather than the usual 2-3 years with new hired junior engineer resident in the SME, and often involved in other activities, which implies a long learning curve. This can give Enginsoft using such simulation a significant commercial advantage.

HPC-based simulation has been used to address optimization design of centrifugal pumps with a prototype service combining a customized workflow to make available to Enginsoft customers for co-design activities, and adequate computing resources, that were not previously available to the company. Due to these improvements, considering the average number of project lost in the last couple of years because Enginsoft could not be compliant to the requests of the end-user in term of time-to-result, Enginsoft expects to increase its market share by at least 1% with a total profit of about €100K. Considering the additional expenses for Enginsoft to use a cloud service instead of the internal resources, this represents a positive return on investment in only six months.

4.18.5 Any Other Business benefits

The provision of a Cloud-based HPC solution would offer significant innovation into the marketplace. In particular the innovation would be derived from the possibility for the end-user to tackle simulations at an unprecedented refinement level. This offer benefits to both Enginsoft as a Consultant firm allowing the opening to a market segment hitherto precluded, and the SMEs end-users which have now an instrument allowing a better, more refined, design with a reduced time to market (up to 50%).

In summary, customer value and benefits from such HPC-cloud CFD simulation could be summarised as follows:

- Increased speed, as a result of obtaining simulation results within hours instead of weeks.
- Increased flexibility, occurring mainly through services utilization on demand and payment out of specific project budgets.
- Reduced cost, encompassing pay on demand capabilities, reduced depreciation costs, and more effective labour deployment.
- Improved quality, as a result of the production of an enhanced solution based on higher resolution models.
- Available know-how of CFD turbomachinery expert
- Faster design processes, driven by optimization made possible by increased CFD simulation throughput
- Increased competitive advantage on the basis of the above mentioned benefits.
- Lower entrance barrier for SME's approaching CFD

In fact, assuming a fast learning curve of a junior engineer in a SME of 2 years, the total cost on the SME would be of (all data are in k€):

	Year 1	Year 1+2
Personnel	50	100
Training	10	10

Hardware and maintenance	20	30
Software licenses	50	100
Total investment	140	240

With an estimate cost for following years of 120 k€ per year (assuming a suitable salary increase).

With an HPC cloud service training and hardware cost would be "hidden" as part of the service. In a pay per use scenario, where a single project can be set up with a token of an EnginSoft know-how engineer to set up the case and a HPC cloud service fee (per CPU_hour consumed) the cost would be:

	Per project
Token know - how engineer	15
75000 HPC CPU hours	15
25% cost of internal design Engineer	10
Total investment	40

Furthermore the learning curve is a risk for the SME, since the formed CFD expert could leave the position, hence an outsourcing politics would mitigate such risk. In fact it is quite a rare situation that a CFD engineer stays in the same position for more than 5 years.

The following table and graphs illustrates the concepts, considering a 7 years' timeframe. The column "Investment" in yellow represents the costs for a company using internal Infrastructures, hiring a junior engineer for design. The column "Value Project/Internal" represents an estimate on the return, considering a first year of training (no return), a second year with 4 projects for a total value of 80 k€, and a regime situation of 8 projects/year having a total value of 160 k€ from the third year on. The red area represents the "high-risk" region when according labour market data the engineer is likely to change position, with a subsequent drop in return for the SME. So, using resources in-house BEP would be acquired in 5 years, but in reality the dynamics of the labour market never allows reaching this point. The final column represents at the contrary the expected situation using outsourcing and Cloud resources, with an immediate regime of 8 projects per year, for a total 160 k€ value, and a likely cost reduction from the fifth year on for software improvements. In this case BEP is reached in 7 years, but risk is considerably limited with respect to the previous scenario.

			Value project	
	Months	Investment	Internal	Cloud
Lease + HW	0	90	0	0
Salary	12	140	0	160
Lease	12	190	0	160
Salary	24	240	80	320
Lease	24	290	80	320
Salary	36	350	240	480

Lease	36	400	240	480
Salary	48	460	400	580
Lease	48	510	400	580
Salary - risk	60	570	560	680
Lease	60	630	560	680
Salary - Risk	72	680	720	780
Lease	72	730	720	780
Salary - Risk	84	790	880	880

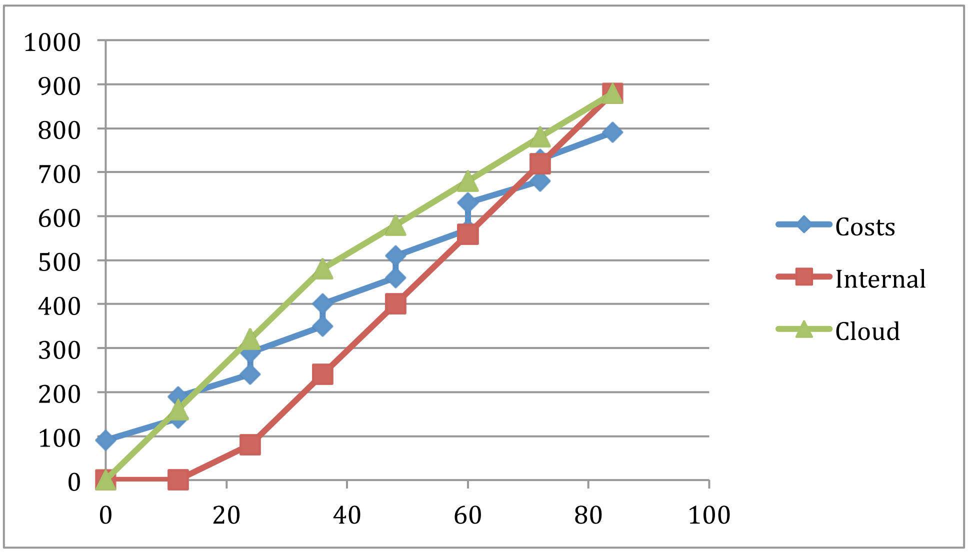

4.18.6 HPC as opposed to standard cluster offerings

Numerical simulation of centrifugal pumps is not easy due to some difficulties: complex geometry, unsteadiness, turbulence, secondary flows, separation, boundary layer, etc. These aspects require a high fidelity CFD model, or, in other words, very fine computational grids and transient analysis.

This translates in extremely long calculations, that can be addressed in timescales suitable to a modern industry lifecycle management only with the careful combination of a software with high parallel efficiency and a hardware that can exploit the application efficiency with a high number of powerful CPUs and high-bandwidth low-latency interconnects, i.e. a HPC system.

4.18.7 Feedback to Work Package 2

Experiment 4.18 has provided to Work Package 2, "Core Service Requirements Capture and Infrastructure Development", as described in Deliverable P418.1, detailed hardware, and software and license requirements, alongside with security and data management constraints and user management and accounting, as well user support needs for porting the ANSYS CFD simulation suite over CINECA's HPC cloud infrastructures.

4.18.8 Feedback to Work Package 3

Experiment 4.18 has provided up to now only indirect feedback to Work Package 3, through the requirement collection document provided to Work Package 2 (Deliverable P418.1).

It is however willing to be an early tester of the marketplace middleware and ready to port the developed framework to the Fortissimo middleware as soon it will be available.

4.18.9 Input to Work Package 8

Experiment 4.18 input to Work Package 8, "Exploitation and Sustainability", will be collected in Deliverable P418.8 (due PM17) comprising a preliminary business model analysis with regard to the commercial exploitation of the experiment's service offering. The analysis will include a detailed presentation of the services to be offered, identify the stakeholders involved, and analyse extensively the CFD market.

4.18.10 Input to Work Package 9

Experiment 4.18 exploits one of the most widely used ISV CFD application in industrial contexts, ANSYS CFD, and Enginsoft is a top ANSYS reseller on the Italian market. ANSYS declared their interest in following the results of the experiment and to provide input for WP89 data collection activities and the annual ISV Forum.

4.18.11 Any further input to Work Package 10

Experiment 4.18 has been featured as Experiment-of-the-Month" on the Fortissimo website.

4.19 Experiment 419: Simulation of Drifting Snow

4.19.1 Overview of the experiment

This is an overview of experiment 419, which consists of CFD simulations of eolian snow transport for civil engineering applications. Such simulations allow the assessment of snow loads on building rooftops due to structures and machinery, something that is not possible to assess properly with present Building Codes procedures. The overall objective of this experiment is to assess the performance of unstructured meshes compared to structured meshes for accuracy, cost and global solution time. Unstructured meshes are an attractive commercial mesh format since they allow meshing the most complex geometries in less than a day, something which could take a week or more with a structured mesh if possible at all. The CFD software to be used is snowFoam, a CFD program based on the CFD library OpenFOAM but not part of the official and community distributions.

4.19.2 The industrial setting

Roof collapses due to accumulated and drifting snow are responsible for losses in the order of hundreds of millions of Euros as well as bodily injuries and loss of lives every year in Northern Europe in particular, and the Northern hemisphere in general. The solver technology tested here can help identify rooftop areas that are prone to increased snow loads and provide the corresponding snow loads. Such analysis can help civil engineers design rooftop structures in perspective and avoid potentially catastrophic events. Additionally, drifting snow simulations can help design protective barriers and fences for roads and highways that help prevent or mitigate drifting snow and reduce its disruptive impact on transportation.

Drifting snow loads are not properly addressed by present Building Code procedures. Moreover, the state-of-the-art in drifting snow simulations is highly empirical and limited to very specific situations. The two-fluid technology used by Binkz in its snowFoam CFD program does not suffer from such limitations but is too computationally costly to use on desktop computers and even advanced scientific workstations. Typically one resorts to medium or large-sized clusters.

The snowFoam CFD program is initially based on the CFD libraries of the OpenFOAM toolkit which is licensed under the GPL. The parts of snowFoam that consist of reproduction and/or modification of original GPL source code are therefore GPL code as well under the terms of GPL.However, there are also standalone independently developed parts of snowFoam that are proprietary and cannot be distributed under the GPL.

snowFoam works with both structured and unstructured meshes, the former being preferred for better numerical performance and a lower computational load. However it is not always possible to build a structured mesh for complex CAD as encountered in typical rooftops with machinery and industrial structures. In situations like that it is much more advantageous to use an unstructured mesh which is much faster to build and always possible. The downside is the larger size of unstructured meshes, and their relatively poor numerical behaviour which requires a host of numerical corrections to be applied. As a result the computational cost for unstructured meshes is an order of magnitude larger than structured meshes, hence the need for even larger HPC systems.

4.19.3 The improvements in the industrial process through the use of an HPC Cloud

The use of snowFoam via the Fortissimo HPC Cloud allows engineers to design their rooftops for all applicable snow types and roof configurations. The downstream savings as compared to rooftop collapse and the potential for bodily injury, loss of lives, downtime of industrial operations as well as insurance and compensation costs is orders of magnitude larger than the analysis and computational cost that is typically 10,000-50,000 Euros. Moreover, drifting snow analysis can be completed within a few weeks to a couple of months. The aforementioned reasons make the use of snowFoam via the Fortissimo HPC Cloud highly advantageous for engineers, engineering firms being the chief contractors of such a CFD service.

4.19.4 The cost benefits of using an HPC Cloud

Building a structured mesh for a complex CAD geometry can take several man-days, or even man-weeks of highly specialized engineering staff working at a very high daily rate. The use of unstructured mesh technology on an HPC cloud reduces the man hours required to build a mesh to a couple of days but increases the computational cost drastically. This can still be cost-effective if the proper computational strategies are used since CPU-hours are more than a hundred times cheaper than engineering man-hours. The downstream savings as compared to rooftop collapse and the potential for bodily injury, loss of lives, downtime of industrial operations as well as insurance and compensation costs is orders of magnitude larger than the analysis and computational cost that is typically 10,000-50,000 Euros.

Typical CFD consulting projects, such as computing wind loads and evaluating structural integrity, are provided with costs in the order of 15kEuro, out of which 5-10kEuro is spent on computational resources, and typically take up to four weeks of overall throughput time. For the snow drift CFD consulting service, the cost and throughput time will have to be in the same order, in order to be sufficiently attractive for civil engineering companies.

Typically, a snow drift CFD consult requires 50k computing hours. With a price of 0.20 euro per CPU hour, the resource costs of 10k Euro falls within the limits for these consulting projects. The scalability of the HPC cloud allows the simulations to be run in two weeks, well within the expected throughput time.

4.19.5 Any Other Business benefits

Complete as indicated above

4.19.6 HPC as opposed to standard cluster offerings

Profiling results show that one of the main bottlenecks of the OpenFOAM runs is the MPI communication. The use of a high-speed and low-latency network as available in a HPC system is therefore important.

4.19.7 Feedback to Work Package 2

The in-browser desktop integration is nearly completed. Once finished, the user can login on the secured webpage with a desktop environment that provides access to the computing resources, storage and visualization facilities. The user does not need any additional software.

4.19.8 Feedback to Work Package 3

The in-browser desktop integration mentioned in the previous section is generic and can communicate with the PBS and the SLURM batch schedulers.

4.19.9 Input to Work Package 8

If successful the engineering service offered through this experiment will contribute a sustainable commercial activity to be exploited within the Fortissimo marketplace, with revenue for the end user Binkz, the cloud provider SURFsara and the Fortissimo marketplace.

4.19.10 Input to Work Package 9

OpenFOAM is open-source software and is therefore not related to an ISV, nor are there any licensing issues.

4.19.11 Any further input to Work Package 10

None.

4.20 Experiment 420: Computational Chemistry

4.20.1.1 Overview of the experiment

The objective of experiment 420, *Cloud based access to commercial computational chemistry codes,* is to partially rewrite existing density-functional theory (DFT) and especially density-functional based tight-binding (DFTB) as well as (reactive) Molecular Dynamics (MD) codes to optimally use CPU-GPU hybrid supercomputers architectures. Two-level parallelization is being implemented in existing codes as well as a newly written MD engine, which will make truly massive parallel calculations possible on supercomputers with many thousands of cores.

This will showcase the added value of cloud-based HPC services, which can then be made available to SMEs. Once a working implementation of the accelerated code is in place on CPU-GPU and multiprocessor supercomputing architectures, this can then become a one-stop-shop for SMEs that want to start modelling large-scale molecular systems. With SCM's fully integrated GUI, end users will be able to prepare jobs, submit them and analyse results in the cloud. In collaboration with SURFsara a cloud-based infrastructure will be developed to sustain the delivery to a larger group of SMEs.

The experiment addresses the materials modelling sector. Albemarle Corporation - an end user from the chemical engineering industry - is testing the whole setup from the user point of view on an industrially relevant application.

4.20.2 The industrial setting

There is a clear industrial demand for faster, predictive computational chemistry simulations. Modelling is a proven powerful tool in chemistry research, providing key information for the design of new chemicals and materials. The software for modelling of large-scale molecular systems has applications in sectors such as electronics (including semiconductors and other electric elements), organic chemistry and other areas of chemistry (fluids and polymers for foods, paints, home and personal care formulations, dyes, adhesives, petrochemistry, etc.) and materials design (such as alloys and ceramics for aerospace and automotive industry).

Our end-user is Albemarle Corporation, a perfect example of a research and innovation-driven company. It runs important and increasing R&D budgets (of $77.1 million, $78.9 million and $82.2 million during 2011, 2012 and 2013, respectively), with research and development efforts supporting each of their business segments. And those R&D budgets are not optional, since the company must continually innovate to comply with ever more stringent health, safety and environmental laws and regulations (such as the EU *Regulation for the Registration, Evaluation and Authorization of Chemicals*, REACH). Albemarle's refinery catalysts research is an important part of its green chemistry efforts, focused primarily on the development of more effective catalysts and related additives to produce clean fuels in commercially viable way. Catalysis research relies on the kind of insight that only computational simulations can bring.

4.20.3 The improvements in the industrial process through the use of an HPC Cloud

However, the *materials by design* approach requires powerful computational resources and significant expertise to use them. This places a burden on innovative companies seeking to increase their competitiveness and reduce times to market by means of computational simulations. In particular, SMEs have smaller budgets and fewer IT staff than larger companies, and the increasing complexity of modelling platforms often outstrips the growth of their IT resources. In the materials modelling sector there exists a plethora of alternative methods and software implementations, and factors such as licensing costs, setup and maintenance make it harder for industrial modellers to use the best tool for the specific problem. Cloud-based HPC services, with their reduced costs and overheads, constitute a unique chance to provide dynamic European SMEs with powerful, easy-to-use and cost-efficient modelling tools, helping them to innovate and compete on the global scale.

In order to showcase the added value of cloud-based HPC services, in this experiment the industrial end user is accessing the GPU hardware recently deployed at SURFsara to model an industrially relevant process, using SCM's software via a remote GUI also developed by SCM. More specifically, the end user is modelling the behaviour of molybdenum sulfide clusters, as used *e.g.* in catalyst-enhanced HydroDesulfurization (HDS) processes. Desulfurization of feedstocks is an important process in refineries, preventing SO_x emission from transport fuels that are derived from these feedstocks, and is an important component of modern clean fuel technologies. The cloud-based setup to remotely access SCM's software on SURFsara's hardware is already allowing simulations which would have been computationally prohibitive with in-house resources.

SCM has created a web-based GUI to its modelling suite, which allows end users to perform simulations without any need to install or set up anything locally (other than a working browser). This is the typical workflow:

1. End users log into the web interface,
2. prepare jobs from within that remote environment,

3 submit the jobs to the cloud, choosing from available computing queues,
4 run massively parallel calculations on supercomputers with many thousands of cores,
5 and analyse the results within the remote environment, as they become available.

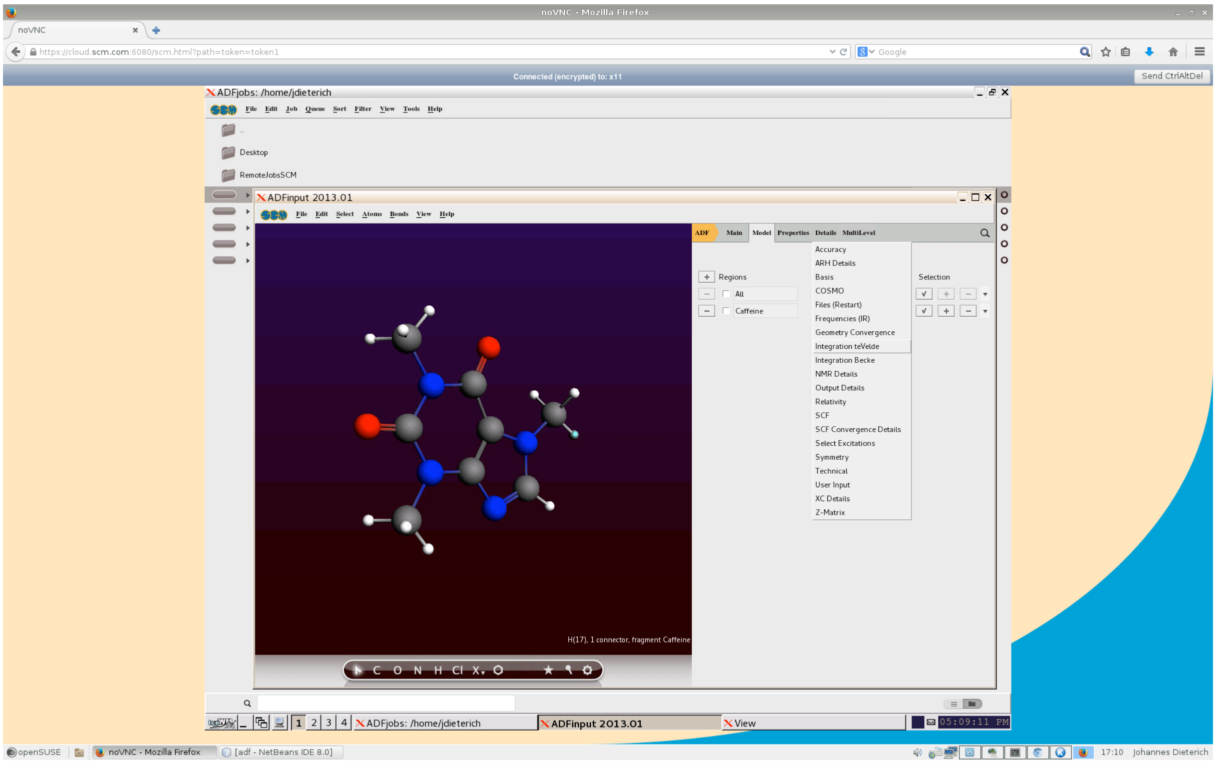

Figure 1: Screenshot of the ADF Modelling Suite web-based cloud interface.

This web front-end to cloud-based computing will allow industrial modellers to perform demanding simulations on cutting-edge platforms, without the need for deep expertise nor investing in expensive computing resources. The availability of easy-to-access modelling software on the cloud means that modellers will be able to test a wider range of models and software implementations, making for more flexible and adaptable industrial processes. Outsourcing the setup and maintenance of software and hardware to HPC providers frees industrial modellers from the burden of technical aspects, allowing them to focus on their research challenges.

4.20.4 The cost benefits of using an HPC Cloud

The benefits associated with using an HPC cloud are apparent, with clearly reduced hardware and software costs. The cloud-based HPC setup is allowing the industrial end user to easily scale up simulations and reach system sizes that would have demanded far too long computation times with their in-house resources. In a traditional, non-cloud setup, this would have required an important investment in new hardware and software licenses, as well as the manpower required to set up and maintain the software. For example, a low-end computing cluster can easily cost around 60,000 EUR, and a commercial license for a typical materials modelling package can well go over 80,000 EUR. On the other hand, commercial HPC providers can charge around 2,500 EUR/year for the usage of 8 cores, clearly a more cost-efficient option. In general, end-users of cloud services avoid hardware licenses or hardware maintenance and upgrade costs, and have reduced up-front software license costs. Cloud computing vendors can manage all of their customers on a single instance of the software,

which allows to amortize costs over many customers and yields substantial economies of scale and skill.
In our particular experiment, Albemarle's area of expertise is in chemical engineering, not hardware and software deployment and maintenance. The company can cut significant costs by outsourcing the latter to HPC specialists. The following table illustrates the yearly costs associated with a typical computing environment for materials modelling as required by our Albemarle's simulation team (~128 CPU cores), comparing two different setups:

- An in-house computing system that includes hardware and software licenses costs (which have been depreciated over 5 years), as well as the staff required to manage it.
- A dedicated HPC provider, in a business model where the end-user has 16 fully dedicated 8-core machines, and the usage of the software is charged as an additional 100% on top of the HPC costs. The CPUs are Xeon Intels, the interconnect is InfiniBand and the GPU accelerators are the AMD FirePro. The price corresponds to a yearly contract, and includes an important discount with respect to PAYG options.

	In-house	**HPC**
Manpower	50,000 EUR	20,000 EUR
Hardware	12,000 EUR	
Software licenses	16,000 EUR	20,000 EUR
Total	78,000 EUR	40,000 EUR

It is apparent that hardware costs are not necessarily the most important ones, easily eclipsed by the staffing expenses required to manage a computing system (or even the software licenses). It is easy to see how this experiment has convinced Albemarle to allocate a sizeable budget for HPC cloud computing next year. SURFsara is not available to industry per se, but an up-and-running marketplace could fulfill their needs.

4.20.5 Any Other Business benefits

However, the value of cloud computing goes beyond reducing hardware and software costs. Cloud-based HPC services allow companies to focus on their core skills to maximize growth and value. The agile and scalable character of cloud computing allows even small businesses to move in new and innovative directions to capture new markets, or to keep up with the market by growing as quickly as the market will allow. In our particular experiment, the end-user can test different modelling approaches (DFT-based as well as reactive molecular dynamics) without the full, up-front license costs associated to each option.

4.20.6 HPC as opposed to standard cluster offerings

- HPC clouds offer a number of key advantages with respect to generic cluster offerings:One limitation is of technical nature: scientific applications such as ours usually require super-fast interconnects, while standard cloud provider usually offer slow communications between nodes, and worse scalability.

- Dedicated HPC providers usually offer better pricing options for heavy "number crunching". A cloud service such as Amazon's EC2 may be cost-efficient when running services such as databases, but the kind of intense, peak computing power demanded by HPC applications quickly becomes comparatively expensive. For example, the commercial HPC providers offer 16 dedicated CPUs for less than 5,000 EUR/year. For the sake of comparison, the closest EC2 option (with interconnect speeds high enough for our typical scientific simulations) charges an upfront fee of 5,850 EUR plus an hourly price of 0.50 EUR for the same number of CPUs, amounting to more than 10,200 EUR/year.

- Another advantage of HPC platforms *vs* standard ones lies in their level of HPC expertise. Users of scientific computational software have special requirements and can outsource comparatively more of their expertise to HPC providers, e.g. installing the very software. Customers of standard cluster offerings usually need to take care themselves of the creation of software images and their installation.
- Software in standard clouds is normally run in virtual machines, while HPC alternatives allow running it directly “on the metal”.

4.20.7 Feedback to Work Package 2

The system at SURFsara is configured such that the ADF Modelling Suite can be run remotely on one of its GPU accelerated nodes. In this way also 3D rendering and visualization can be performed on the remote node. The setup at SURFsara therefore allows end users to perform computational chemistry computations using ADF without having to install any additional software on the client side, apart from a VNC viewer.

4.20.8 Feedback to Work Package 3

Molecular dynamics simulations require big and fast data transfer and storage. They may lead to result files of many GBs. The user may want to analyse them further on his/her local machine. The connection speed should be outstanding and not be limited on the marketplace side

4.20.9 Input to Work Package 8

It is crucial to let industry know that the marketplace exists and offers a top-notch secure environment for HPC cloud computing. The demand should come from industry itself for which they should first know that this solution exists.
Recommendation is to start with demanding users, such as pharmaceutical companies, and interview them on their needs and requirements, such as very secure data protection.

Offer the end users initial free access to collect practical feedback and improve the system.
Go out to the end users.

4.20.10 Input to Work Package 9

We may want to use it to offer our software in the cloud, provided it is a well-respected, widely used and trusted marketplace. We will always weigh the benefits of this marketplace to all the other options of offering our software in the cloud, either via commercial vendors, or by ourselves (the latter not looking very attractive right now). In that sense, it would be very helpful if all paperwork and accounting were handled on the marketplace side. The ISV should simply get a payment every 3 months depending on much the software was used and by which institutes. The ISV should get a usage overview per module per user.
The system should allow for a wide variety of business models and have policies for failed calculations (reimburse or not).

4.20.11 Any further input to Work Package 10

None