FORTISSIMO

D4.3
Initial Application Experiments Public Report

Workpackage:	4	Application Experiments 1	
Author(s):	Mark Sawyer (Editor)		The University of Edinburgh
	All experiments		
Authorized by	Mark Sawyer		The University of Edinburgh
Reviewer	Francis Wray		scapos
Reviewer	Ullrich Becker-Lemgau		Intel
Reviewer	Nejc Bat		Arctur
Dissemination Level	PUBLIC		

Date	Author	Comments	Version	Status
2015-01-23	M Sawyer	Initial draft	V0.1	Draft
2015-02-04	M Sawyer	For internal review	V0.2	Draft
2015-02-19	M Sawyer	Response to internal review	V0.3	Draft
2015-02-25	M Sawyer	Approved	V1.0	Draft

Executive Summary

Fortissimo is a collaborative project that aims to help enable European SMEs to be more competitive globally through the use of simulation services running on a High Performance Computing cloud infrastructure. The importance of advanced simulation to the competitiveness of both large and small companies is well established. However, the simulation of, for example, high-pressure gas cylinders or the moulding of plastics requires enormous computing power and specialized software tools and services. Generally, large companies have easier access to advanced simulation than SMEs, which are facing both technological hurdles and financial challenges. This means that SMEs are often not able to take advantage of advanced simulation, even though it would clearly make them more competitive.

Fortissimo gives users access to a cloud-based HPC infrastructure by means of a Marketplace, which is currently under development. The initial requirements for this Marketplace have been provided by an initial set of 20 experiments, each of which was aimed at solving a specific business problem that required the cost-effective HPC that the cloud-based approach can provide. This report summarises the results and findings of the experiments.

The results of the experiments strongly support the case for cloud-based HPC and the Fortissimo approach. Significant business benefits have been identified by the experiments (in many cases these have been directly quantified) and useful suggestions for how the Fortissimo Marketplace should develop have been made. Furthermore, many of the experiments have concrete plans to make business solutions available through the Marketplace. These include the provision of niche applications specific to the business needs of the experiment partners to more general use of HPC-enabled software via the cloud, and added-value services.

The experiments have been a valuable source of further requirements for the subsequent development of the Fortissimo Marketplace, and several features and areas of development have been suggested as being of value to Marketplace users. These further requirements address, among other things, security and privacy, ease of use, management of licences and supporting cooperative working. These requirements will be fed into the process of development for the Marketplace. New features and components that meet the requirements will be implemented following a process of feasibility analysis and prioritisation based on the benefits gained and the cost of implementation.

For more information see http://www.fortissimo-project.eu.

The Fortissimo project is an initiative supported by the European Commission under the 7th Framework Programme.

Table of Contents

1 Introduction

This document is the public final report on the first tranche of experiments in Fortissimo. It focuses on the achievements and results of the work, and details the actual and potential benefits and impact of each experiment. The relevance of each experiment to the Fortissimo cloud-based HPC Marketplace, which is currently under development, is also presented. In-depth technical details of the work carried out by each experiment are out of scope for this report.

The document is structured as follows:

Section 2 lists the experiments, their purpose and the roles taken by the partners.

Section 3 lists the benefits and impact of each experiment.

Section 4 gives details of the motivation (the business background), the relevance to Fortissimo and the impact and benefits of each experiment.

2 Objectives of the Experiments

Workpackage 4 of the project consists of a set of 20 experiments that were included in the original proposal. These experiments were intended to provide requirements for the Fortissimo Marketplace and to evaluate the approach taken by the project to implement it. Each experiment does this in the context of a real business problem defined by the experiment partners.

The experiments and the role of each partner are listed in Table 1.

There were a total of 30 partners taking part in the experiments (in addition to the core Fortissimo partners). A typical experiment consisted of an end-user, a technology provider (which may be, for example an Independent Software Vendor (ISV)) and an HPC expert.

Experiment Title	**Partners**
Experiment 401: Cloud-based simulation of aerodynamics of light aircraft	End User: Pipistrel HPC Expert: XLAB Research HPC Provider: Arctur Host Centre: Arctur
Experiment 402: Cloud-based simulation of continuous casting	End User: Ergolines HPC Expert: Arctur HPC Provider: Arctur Host Centre: Arctur
Experiment 403: Micro-Macro mechanical properties of high pressure vessels	End User: Mikrosam HPC Expert: XLAB Research HPC Provider: Arctur Host Centre: Arctur
Experiment 404: Improved mechanical design of metal flanges	End User: Texas Controls Application Expert: Asociación de Investigación Metalúrgica del Noroeste. (AIMEN) HPC Expert: Fundación Centro Tecnológico de Supercomputación de Galicia (CESGA) HPC Provider: CESGA Host Centre: CESGA
Experiment 405: Cloud-based multi-physics simulation	End User: Prysmian HPC Expert: Cineca ISV: COMSOL HPC Provider: Cineca Host Centre: Cineca

Experiment Title	**Partners**
Experiment 406: Cloud-based HPC tools for urban planning	End User: Integrated Environmental Solutions (IES) Application Expert: IES HPC Expert: The University of Edinburgh ISV: IES HPC Provider: The University of Edinburgh Host Centre: The University of Edinburgh
Experiment 407: Cloud-based CFD simulation in automotive design	End User: VolksWagen) Application Expert: RWTH Aachen HPC Expert: Numerical Algorithms Group (NAG) HPC Provider: The University of Edinburgh Host Centre: The University of Edinburgh
Experiment 408: Main routing architecture optimisation research experiment in the Cloud	End User: KE-works Application Expert: KE-works HPC Expert: Gompute ISV: Fraunhofer HPC Provider: Gompute Host Centre: Gompute
Experiment 409: Cloud-based correlation between simulation and tests for mechatronics	End User: SDI sas HPC Expert: Grand equipment national de calcul intensif (GENCI) ISV: Scilab Enterprises HPC Provider: Bull Host Centre: Bull
Experiment 410: Cloud-based simulation of air-quality over cities	End User: NUMTECH HPC Expert: GENCI ISV: Cambridge Enviromental Research Consultants CERC HPC Provider: Bull Host Centre: Bull
Experiment 411: Global analysis of risers, moorings and flowlines	End User: PRINCIPIA HPC Expert: Bull ISV: PRINCIPIA HPC Provider: Bull Host Centre: Bull
Experiment 412: Cloud-based simulation environment for CO2 emission prediction for automotive engines	End User: AVL LIST GmbH HPC Expert: Universität Stuttgart ISV: AVL LIST GmbH HPC Provider: Universität Stuttgart Host Centre: Universität Stuttgart

Experiment Title	**Partners**
Experiment 413: Multi-physics simulation in the Cloud	End User: External Application Experts: Fraunhofer, Powersys Solutions HPC Expert: Universität Stuttgart ISVs: Powersys Solutions, HPC Provider: Universität Stuttgart Host Centre: Universität Stuttgart
Experiment 414: Cloud-based molecular simulation for industrial chemical engineering applications	End User: Lonza Group HPC Expert: Universität Stuttgart ISV: University of Paderborn HPC Provider: Universität Stuttgart Host Centre: Universität Stuttgart
Experiment 415: Cloud-based simulation for low-pressure die-casting of copper alloys	End User: IMR Engineering&Technologies HPC Expert: Arctur HPC Provider: Arctur Host Centre: Arctur
Experiment 416: Cloud-based simulation of structural crash tests	End User: Gestamp Automoción Application Expert: N/A HPC Expert: GENCI ISV: ESI Group HPC Provider: Bull Host Centre: Bull
Experiment 417: Cloud-based Computational Fluid Dynamics Simulation	End User: Koenigsegg Automotive AB Application Expert: National Technical University of Athens HPC Expert: Cineca ISV: iconCFD HPC Provider: Cineca Host Centre: Cineca
Experiment 418: Cloud-based CFD turbomachinery simulation	End User: EnginSoft Application Expert: N/A HPC Expert: Cineca ISV: ANSYS HPC Provider: Cineca Host Centre: Cineca
Experiment 419: CFD simulations of eolian snow transport for civil engineering applications	End User: BinkZ Application Expert: BinkZ HPC Expert: VORtech bv HPC Provider: SURFsara Host Centre: SURFsara

Experiment Title	**Partners**
Experiment 420: Cloud based access to commercial computational chemistry codes	End User: Scientific Computing & Modelling (SCM) HPC Expert: SURFsara ISV: SCM HPC Provider: SURFsara Host Centre: SURFsara

Table 1: List of initial experiments in Fortissimo

3 Impact and Benefits

The partners reported the following impacts and benefits that were demonstrated through the experiments.

Experiment Title	**Impact and Benefits**
Experiment 401: Cloud-based simulation of aerodynamics of light aircraft	• 12-fold reduction in cost when using cloud-based HPC compared to in-house resources. • Speed up of design phase, giving competitive edge. • First use of HPC by the SME
Experiment 402: Cloud-based simulation of continuous casting	• Better understanding of physical process • Fast return on investment • Estimated 1% (€3M) annual saving for a medium-sized steel plant.
Experiment 403: Micro-Macro mechanical properties of high pressure vessels	• Decrease in labour costs of 20% • More optimal designs improving SMEs' reputation.
Experiment 404: Improved mechanical design of metal flanges	• 33% saving in time (labour cost) for the tightening process. • Process transformed from "based-on-experience" activity into a "knowledge-based" activity
Experiment 405: Cloud-based multi-physics simulation	• Reduction in time for analysis (from weeks to hours). • Estimated €50K start-up plus €10K per year saving for the SME by using cloud-based HPC

Experiment Title	Impact and Benefits
Experiment 406: Cloud-based HPC tools for urban planning	• Reduction in simulation time (2 weeks to 2 days) • Capability to run simulations that previously were too large to be feasible
Experiment 407: Cloud-based CFD simulation in automotive design	• Faster job turnaround giving competitive edge • Cost effectiveness proven for realistic levels of utilisation.
Experiment 408: Main routing architecture optimisation research experiment in the Cloud	• Optimized process using cloud-based HPC shows a 2.5% reduction in cost and weight of the wiring system. • Lead-time reduction of 90%.
Experiment 409: Cloud-based correlation between simulation and tests for mechatronics	• Parallel optimizations with a speedup by 3.9 in our test benches (for 4 parallel runs). • High value contracts being negotiated as a direct result of the experiment.
Experiment 410: Cloud-based simulation of air-quality over cities	• Twice as many cores can be used in the cloud-based scenario at the same cost as internal provision. • Profit margins increased by 5% for the end-user.
Experiment 411: Global analysis of risers, moorings and flowlines	• Prototype HPC version of DeepLines available • Potential reduction in processing time from 60 to 3 days.
Experiment 412: Cloud-based simulation environment for CO2 emission prediction for automotive engines	• Computing costs reduced by up to 90%. • Cost-effective use of licences.
Experiment 413: Multi-physics simulation in the Cloud	• Removed need and cost for HPC infrastructure for the user. • Pre-installed simulation environment available.
Experiment 414: Cloud-based molecular simulation for industrial chemical engineering applications	• Estimated 5% reduction in the cost of distillation column. • Possible to carry out previously infeasible simulations.

Experiment Title	Impact and Benefits
Experiment 415: Cloud-based simulation for low-pressure die-casting of copper alloys	• Reduces the development time of the moulds • Reduces the cost and the number of tests before going into production by 20%.
Experiment 416: Cloud-based simulation of structural crash tests	• Removes synchronisation step of CAD/CAE • Speed up factor of approximately 3.5
Experiment 417: Cloud-based Computational Fluid Dynamics Simulation	• Reduced cost and improved quality, as a result of using higher resolution models. • Faster design processes leading to increased competitive advantage.
Experiment 418: Cloud-based CFD turbomachinery simulation	• Simulation results in hours instead of weeks. • Expected increase in market share for end-user of 1%
Experiment 419: CFD simulations of eolian snow transport for civil engineering applications	• The commercial feasibility was proven for a snow drift simulation service. • Cost estimates can be guaranteed, thus minimizing the risk of budget overrun in the computational resource.
Experiment 420: Cloud based access to commercial computational chemistry codes	• HPC cloud solution approximately half the cost of the in-house solution. • Cloud-based HPC services allow companies to focus on their core skills to maximize growth and value.

Table 2: Impact and Benefits of the first 20 Fortissimo Experiments.

4 The Experiments

4.1 Experiment 401: Cloud-based simulation of aerodynamics of light aircraft

4.1.1 Motivation

There are two ways to examine the behaviour of airflow over 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 dynamics (CFD) simulation that is run on the 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. CFD simulations are usually used throughout the whole design phase, whereas the wind tunnel tests are used only at certain phases of the design for evaluation purposes, 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 flow 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 need to be employed. Such improved fidelity aerodynamic simulations can be of high time and space computational complexity. The use of HPC services is therefore very valuable for an aircraft designer like Pipistrel when CFD simulations are the only tool to give insight into aircraft in-flight behaviour.

4.1.2 Relevance to Fortissimo

Pipistrel is a proficient user of the OpenFOAM software, which is used as the main CFD tool at the company. Typically RANS simulations with fully turbulent flow models were used during the design phase. Since Pipistrel wanted to go a step further, the Fortissimo project came at the right time. It enabled the company to explore what methods of higher complexity can offer. The use of HPC allowed Pipistrel to run simulations that describe the real world more precisely and give better information on how the airplane behaves in-flight. These kinds of simulations are too demanding, in terms of computer power, to be run on a small in-house cluster.

In the first phase of the experiment Pipistrel successfully used a low-Reynolds k-kl-omega turbulence model simulating the airflow past a complete aircraft at cruise speed conditions. Such a turbulence model is able to estimate the location of the laminar-turbulent transition that significantly affects the performance. For this purpose Pipistrel learned how to make a mesh with more than 10 times larger number of cells and more than 100 times thinner first wall layer than used in the past. At the same time Pipistrel learned how to run, handle and post-process such big cases on a cloud-based HPC. In the second phase another type of simulation was tackled, namely Large Eddy simulation (LES), which uses turbulence models to take into account the fine flow structures, but calculates the large flow structures directly from the Navier-Stokes equations.

During the experiment Pipistrel also showed that CFD simulations of a typical complexity that were run on an in-house cluster exhibited an almost linear speed-up up to at least 120 cores used in the HPC system and that much more complex simulations could be obtained in approximately the same time.

On the other hand, experiment 401 represents a successful example of cloud-based HPC usage in the design phase of light sports aircraft by SME manufacturers. Since one of the objectives of the Fortissimo project is to show a versatile usage of cloud-based HPC, experiment 401 helped to cover also this industrial sector. During the experiment Pipistrel increased its knowledge in the field of CFD simulations and HPC computations, and upgraded its expertise in aerodynamic simulation software OpenFOAM. All this know-how will be offered in the Marketplace, where it will, also because of the increasing popularity of CFD software OpenFOAM, help other users of this application.

4.1.3 Impact and Business Benefits

Buying and operating its own cluster with 120 cores would cost Pipistrel approximately €250,000 per year if an assumption is made that the cluster gets depreciated over the time period of two years. On the other hand, rental of HPC resources for approximately 600k CPUh would cost Pipistrel approximately €20,000 Rental is therefore more than 12 times cheaper than buying, operating and maintaining its own cluster.

CFD simulations done at Pipistrel are typically of a research nature or new-aircraft-design nature, which means there is not 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 resources.

The Fortissimo experiment in particular enabled Pipistrel to use HPC for the first time and in this way gave the opportunity to learn what the capabilities of HPC are. Pipistrel ran more time and space computationally complex simulations that improved the fidelity of aerodynamic simulations and that were previously not possible with its in-house system. During the experiment Pipistrel learned what is the proper workflow using cloud-based HPC, it learned how to produce much finer mesh of better quality and it learned approximately how long it does 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 run on HPC and in this way to decide if a use of HPC is justified or not.

Pipistrel learned from the Fortissimo experiment results that HPC services will be very valuable also during a design phase of some future aircraft. On one hand HPC can be used to run much more demanding simulations that improve the fidelity of results. The time needed for such simulations on HPC doesn't increase substantially compared to simulations currently run on Pipistrel's in-house cluster (2-3 days). However they give more data of a higher quality 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 a 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 becomes available in this way to use in the design phase.

4.2 Experiment 402: Cloud-based simulation of continuous casting

4.2.1 Motivation

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 damages 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 (High-Performance Computing).

4.2.2 Relevance to Fortissimo

Participation in the Fortissimo Experiment 402 - "Cloud-based Simulations of Continuous Casting" provided Ergolines with the necessary HPC resources 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 two-phase configuration. Realization of a dedicated laboratory scale model allowed for successive validation of the numerical results. Excellent agreement between numerical predictions and experimental data 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.

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 currently does not possess the necessary computational infrastructure, the possibility to use cloud-based HPC resources proved fundamental. In fact, participation in the Fortissimo experiment allowed Ergolines to exploit supercomputing resources to reduce computational times without having to sustain the relevant costs of operating a dedicated infrastructure.

4.2.3 Impact and Business Benefits

Participating in the Fortissimo Project brought significant business benefits to the company: in fact, as a result of dedicated HPC-based numerical simulations, Ergolines created the premise for the future development of a novel automated slag detection system relying on vibration analysis, which in turn can provide considerable benefits to continuous casting plants. Such an innovative technology has a high commercial potential: the system would in fact bring key advantages to the casting industry, by improving both occupational safety and productivity.

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 *breakouts*, 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 enabled a thorough analysis of the different ladle emptying mechanisms, thus providing Ergolines' researchers with key insight into the physics of the system. As a result, it has been possible to establish a correlation between the shroud vibrational signal and the fluid dynamics of the system.

Based on this achievement, it will be possible to develop a novel technology for ladle slag detection relying 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 €300M. Envisioned annual savings for a medium size factory are of the order of €3M 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.

4.3 Experiment 403: Micro-Macro mechanical properties of high pressure vessels

The objective of the experiment 403 (Micro-Macro mechanical properties of high-pressure vessels) was to improve the capability of the end-users to satisfy the principal ISO 11439 standard and the ECE Regulation 110 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 laminate for production of high-pressure vessels, parallel optimization was used.

This was an innovative idea in the field of composite design and an entirely new field within the scientific research performed by Mikrosam that improved the filament winding solutions the company offers.

The benefits of the experiment - parallel optimization outcomes - are expected to be replicable in design and manufacturing processes of composite high-pressure vessels for the aviation, space, medical, and other industries that increasingly substitute metal alloys with composite material structures.

4.3.1 Motivation

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 many 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 moulded into complex shapes.

Filament winding has become a popular manufacturing technology 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. Composite materials have the potential to replace widely used steel and aluminium and offer better performance. Composite materials can save 60% to 80% in component weight by replacing steel and 20% to 50% by replacing aluminium. The filament winding machines Mikrosam provides are custom-made and vary depending on the final application (product) they are meant for. In response to fossil fuel environmental concerns and the increasing demand for optimization of their storage and transportation, composite high-pressure vessels are gaining popularity.

Composite design is a very time consuming process. Experiment 403 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 and 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 the required elastic constants. With the technical advances subject to this experiment under the Fortissimo project, Mikrosam has acquired a competitive advantage through offering potential customers accelerated speed in designing the composite of their final product.

4.3.2 Relevance to Fortissimo

Experiment 403 has shown that HPC simulation not only has a positive impact on the design processes, but also solves an important manufacturing problem. Aside from the improved composite-laminate design process, the experiment has shown that HPC simulation can speed-up the filament winding machine and accelerate the FEA analysis performed to optimize the overall machine construction.

The HPC simulation and the algorithms that have been used in this experiment address not only technical, but also commercial needs of Mikrosam as tools that would help the company provide more reliable solutions to its customers.

Advances in manufacturing design, by being an integral part of the solutions Mikrosam offers, add even more value to the equipment of end-users and final part manufacturers, as they can be further implemented in regular manufacturing processes.

While the calculations that are subject to this experiment and that benefit from HPC are not operations performed on a daily basis, HPC resources appear to be advantageous but unaffordable to an SME. Hence, the cloud-based self-service platform that is one of the key elements of this project infrastructure proves to be an excellent solution available to any interested HPC user. Being easy to access, user friendly and available on-demand, cloud-based HPC directly increases the resources and capacity available to any application that involves complex computations.

4.3.3 Impact and Business Benefits

For each filament winding machine for production of type IV CNG cylinders it is necessary to anticipate different combinations of materials and the behaviour of different winding angles. With the algorithm that is developed in this experiment, Mikrosam gets the opportunity to reduce the time of making complex computations for a given database of materials and taking account of the effects of different winding angles on the final product, to simulate and evaluate the design in advance.

Working with a database of materials leads to more theoretical tests, thus saving material costs incurred by physical design and testing. According to the tests done during this experiment Mikrosam expects that using parallel computation on an HPC system can significantly reduce the time for composite design. The use of HPC resources allows more 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. Since the design of the composite laminate requires multiple guesses in order to obtain the appropriate combination of angles, all of the calculations for each of the guesses are theoretical. This means that selected samples of obtained results from the theoretical calculations should be further physically tested. Using the application developed in this experiment starting with only one initial guess for the angles’ design, the application makes theoretical calculations involving multiple angles for different designs and returns the samples for physical testing.

This means that all costs incurred from the actual physical design testing are directly related to the theoretical design cost savings. The optimized simulation of the composite pressure vessel design, subject to this experiment, decreases the number of physical tests Mikrosam necessarily makes whenever it manufactures a filament-winding machine for high-pressure vessels. 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 its functionality and final product design, a machine manufacturer would need €2,500 for material cost and €1,459 for labour cost to actually produce all of the units to test the product and the machine. These numbers together with a 10% saving on material production and 20% saving on labour costs that the use of HPC gives, lead us to €542 saving when assessing 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 which is over 10% of the total design cost.

With the simulation of the composite design of the pressure vessels, a number of physical tests can be avoided, leading to saving on material and labour (designers, engineers and production workers), and hence reduction in the total cost for actual physical design and testing of the final product and of the machine. 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. This means that business benefits from the simulation of the composite design of individual high-pressure vessel, subject to this experiment, transfer to the machine manufacturer and the high-pressure vessel manufacturer i.e. the machine buyer.

Being able to offer optimized and time-efficient composite design of the final product will better convince potential customers to choose Mikrosam as a supplier of the overall filament winding solution.

Being actively involved in a research project to improve Mikrosam's technical and technological competence around which they build their 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 their customers and their product end-users.

Fortissimo enabled use of HPC resources by Mikrosam for the first time, and following this successful application the company is considering using HPC in further applications for complex time-consuming calculations. This experiment experience encouraged Mikrosam 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.

4.4 Experiment 404: Improved mechanical design of metal flanges

4.4.1 Motivation

Texas Controls is a Spanish SME that offers technological products and services oriented to the instrumentation, and tightening and sealing solutions. In the field of tightening and sealing processes, Texas Controls is today the sole provider of engineering solutions in this field. As such, one of their 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 within the industries. These solutions are especially important to customers in the gas and oil industry. In this sector, 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, to understand how the elastic interactions between elements are produced.

This experiment developed a model for simulating the tightening process and a front-end application to control the parametric simulations. Simulations typically performed in the industry do not take the tightening process into account. During the tightening process damage to parts such as bolts and flanges could occur. Any damage suffered by any or all of these parts could have devastating consequences. It is in this very process that the 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 those 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 to other industries. Thus, in order to get the urgently required results as soon as possible the best solution is to run all the parametric cases simultaneously. This will allow Texas Controls to improve the design of tightening processes and gain understanding and insight for each case using the Fortissimo infrastructure.

4.4.2 Relevance to Fortissimo

From a financial standpoint, it is unthinkable for a small engineering SME such as Texas Controls to make the 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.

Tightening work is by no means a regularly scheduled activity. Large variations 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. Cloud-based HPC offers fast responses and can supply the appropriate level of resources required in each case.

The Fortissimo infrastructure will allow Texas Controls to improve the design of tightening processes 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 computing resources it provides, but also for the knowledgeable human capital accrued by the application and HPC experts available to us whenever needed.

With the simulation model developed in the experiment, Texas Controls will be able to conduct some analysis before start tightening process for small/medium sized flanges. 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.

4.4.3 Impact and Business Benefits

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 damaging flanges during tightening processes.

The following is a list of extra benefits that the SME expects from its participation in this experiment:

- An improvement in the company's knowledge of FEM calculation. Turn a "based-on-experience" activity into a "knowledge-based" activity.
- To further improve and optimize the design of tightening tools and procedures.
- New Scientific & Technical knowledge about a key business activity.
- A new tool with an easy interface that will help to improve the design products and manufacturing procedures.
- To gain access to computing technologies, reduce the "learning curve" and improve the technical background of the company engineers.
- To create new revenue cash flows and to generate new services for new and existing customers.
- To sound out new business lines: Offer a tightening simulation service on HPC to third parties.
- To grow and stress differentiating factors to set the company apart from competitors.
- To collaborate with other organisations involved in the Fortissimo project, not in the same line of business, but important to improve knowledge.

4.5 Experiment 405: Cloud-based multi-physics simulation

4.5.1 Motivation

The objective of experiment 405 "Cloud-based multiphysics simulation" was to showcase the use of cloud-based HPC services to investigate multiphysics simulations using commercial and open-source software packages.

Prysmian 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 represents a breakthrough for Prysmian business.

Prysmian adopted COMSOL Multiphysics as the main tool for their electromagnetics simulations years ago and it uses a computing infrastructure consisting of a few high-end workstations. COMSOL Multiphysics is a modelling and simulation tool that allows engineers to accurately model systems in which multiple physical disciplines are involved at the same time. To satisfy the need for new and finer simulations, the experiment provided Prysmian technicians with a workflow which integrates the 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 technical part of the experiment consisted of 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 by exploiting the expertise of HPC centres to build customizable high-productivity and high-performance set-ups.

High-voltage cables for power transmission are sources of electromagnetic losses during operation, when the 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 the cable's elements is then dissipated into ambient heat, and a temperature distribution is established inside cables as a result of equilibrium between generation and dissipation.

During steady-state normal operation, the conductor temperature is not allowed to exceed a prescribed limit, so as not to speed-up ageing of the 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 the above discussion the rating calculation then involves:

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

So far, 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 spreadsheets 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, which means a waste of money.

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

4.5.2 Relevance to Fortissimo

The use of the Fortissimo HPC Cloud helps Prysmian to:

1. Improve the calculation of loss contributions inside an energy cable, especially those induced by the magnetic field from the load current;
2. Model heat transfer from the cable surface to the 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 (as in most of the applications). Proper cable rating calculations include 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 buoyancy movements, radiation between the cables and walls and between the cables themselves. CFD computer analysis can provide a great deal of information about these phenomena by solving the Navier-Stokes equations and the nonlinear heat transfer equation associated with radiative phenomena.

The need for cloud-based HPC simulation is of utmost importance. In this experiment, by moving the simulation of CFD or FEM models to a cloud-based HPC environment, processing times can be impressively shortened. Moreover, with the possibility of a flexible use of the same licence 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 Amazon EC2.

4.5.3 Impact and Business Benefits

Prysmian already moved to a design workflow for their products completely based on CAE simulations about 7 years ago. At the time it was estimated that this would provide 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 experiment 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 year framework as described later in more detail.

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

Cloud-based HPC simulation has been used to address the optimization of the design of high-voltage cables with a prototype service combining FEM simulations of a new design and CFD simulations of cables in the environmental conditions pertaining to their installation.

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 for software licences.

Considering the FEM component, with the current Prysmian IT infrastructure for design consisting of 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 an 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-year operational lifecycle, this amounts to a total of close to €300K. At the current average fees of HPC cloud 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, leading to estimates that using HPC cloud services instead of investing in an in-house system can provide savings of around 30%.

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

In summary, customer value and benefits from such HPC-cloud CFD simulation 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.

- Reduced cost, encompassing pay on demand capabilities, reduced depreciation costs, and more effective labour deployment. This has been estimated 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 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.

4.6 Experiment 406: Cloud-based HPC tools for urban planning

4.6.1 Motivation

Fortissimo experiment 406 is titled "Experiment 6: Cloud-based HPC tools for urban planning".

The participants in this experiment are IES, developers of one of the world's leading 3D building-simulation software suites, The Virtual Environment (VE), and The University of Edinburgh, 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 Relevance to Fortissimo

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 Impact and Business Benefits

To validate the experiment prior to launching the first version to customers tested the upload of building model simulations of various sizes ranging from small to extra-large was tested. 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 1: Simulation performance

Figure 2 illustrates the largest of these simulations. This sample model would have taken 2 weeks to run the SunCast simulation on the desktop but on the Fortissimo infrastructure (using the INDY HPC system at Edinburgh) we were able to reduce simulation time to only 2.2 days.

Very detailed model + shading fins all-round:

- **35,626 surfaces & 5,538 windows, = 3,385 CPU core-hours calculations**

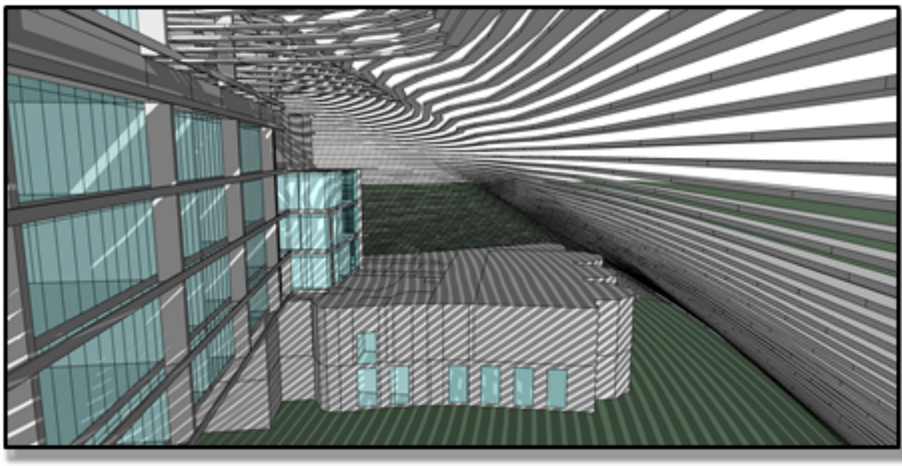

By far the most complex solar analysis we've ever looked into. **Not feasible before**

Figure 2: A previously infeasible simulation is now possible.

Based on these tests it was possible to determine timings for simulations and therefore produce a cost model 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 (calculated at experiment time).

The following 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.

4.7 Experiment 407: Cloud-based CFD simulation in automotive design

The Fortissimo experiment 407 concerns the provision of cloud-based supercomputing infrastructure to enable in-depth CFD simulations used within the context of automotive industry. OpenFOAM is utilised as the cloud-based CFD package to tackle some representative external aerodynamic problems widely used in automotive design.

4.7.1 Motivation

The automotive industry heavily relies on 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. However the initial investment in the necessary HPC hardware, combined with costly commercial software licensing models (often charging on a per core/CPU basis) can make in-house CFD solutions expensive and beyond the reach of SMEs.

Some open-source solutions, such as OpenFOAM and similar, can be credible alternatives to commercial CFD packages. However, open-source packages are often hard to learn, difficult to set up, and require specialised knowledge to use. A cloud-based infrastructure can make such solutions more accessible by providing end users with friendly interface and streamlined processes. At the same time, cloud-based HPC solutions, particularly those on pay-per-use charging model, can be significantly cheaper to use for SME users.

4.7.2 Relevance to Fortissimo

From the end users' perspective, it is important for a cloud-based service to be part of a larger market place like Fortissimo that end users can trust. From a service provider’s point of view, standard facilities offered by a central market place, such as user authentication, secure file storage, and programming interface can simplify the technical work to build a cloud-based service.

4.7.3 Impact and Business Benefits

Typical licence costs 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 The University of Edinburgh 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-user a competitive edge.

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 they run the simulations. However, Fortissimo marketplace can attract third party suppliers with such specialised skills. End users can then pay for these value-added services via the marketplace.

A detailed cost analysis comparing the cost of running simulations on the cloud or on local clusters was conducted as part of the experiment. The main conclusion of this analysis is that a cloud-based solution is likely to be more cost-effective than an in-house solution in most practical situations because local clusters are unlikely to be fully utilised.

4.8 Experiment 408: Main routing architecture optimisation research experiment in the Cloud

This report describes the results of the Fortissimo experiment 408 by the partners KE-works, Noesis, Fraunhofer SCAI and Gompute, named MORE.CLOUD. The use case covers the electrical design process of an Electrical Wiring Interconnect System (EWIS) in aircraft. It is supported by a simulation workflow (known as the MORE workflow) integrating three Knowledge Based Engineering (KBE) applications. The future state process has been defined in collaboration with the customer to improve the Operational Excellence.

The first task was to define or update the Main Routing Architecture (MRA). The MRA defines the topology of routes of electrical wiring through the aircraft. The second task was to route the signals over this MRA from and to systems in an aircraft. Then the last task was to select connectors to couple the signals to the systems. KE-works, acting as an end user in this experiment, aims to reduce weight and cost for wiring systems for complex products such as aircraft. Furthermore KE-works developed dedicated design applications for wiring systems for complex products such as an aircraft. Noesis, acting as the Software Vendor, provides an optimizer tool applicable for the experiment objective. KE-works, acts as a software vendor providing a web-based workflow management platform (KE-chain). Fraunhofer SCAI acts as a software vendor providing a licence manager tool.

The MORE.CLOUD (Main Routing Architecture Optimisation Research Experiment in the CLOUD) experiment is involved in the execution of different EWIS design 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 are enabled to use the flexible licence manager component available on 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, resulting in cost and weight savings. 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 Licence as a Service licensing provider. A two phase implementation was taken where phase 1 focuses on making the experiment able to run on the computing framework environment and phase 2 focuses on integrating the ‘computing’ infrastructure with the ‘provisioning’ framework, which includes the Engineering as a Service Portal and elasticLM licence service.

4.8.1 Motivation

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 current trend for Small and Medium Enterprises (SMEs) in need of computational resources is to use public cloud infrastructures. The main reasons being expected lower costs than those of maintaining their own infrastructure and an increased flexibility in response to changing computing demand. A closer look shows that the applications are often 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 that are designed for parallel execution in High Performance Computing (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 licences that in general are tailored for executing the simulation using local resources.

The experiment has the following objectives:

- Formalization and execution of business processes to simulation tasks (Engineering as a Service) of the MORE experiment;
- Use of licence-protected simulation applications in off-site cloud-based HPC environments using elasticLM to support a pay-as-you-go licensing schema;
- Solving the KE-works problem of integrating the elasticLM software and providing Engineering Services to manufacturing industries.
- Taking advantage of using a workflow-based cloud HPC environment without taking the burden to buy and maintain the entire computational infrastructure and without the need of in-depth knowledge of HPC platforms to run parallel computations.

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 licences 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 licences as needed when running the application using external the cloud-based HPC resources.

The expected business benefits for electrical wiring companies (and similar) are a reduction of lead-time for a single (non-recurring) electrical design of an Electrical Wiring and Interconnection System of about 90%, a cost reduction of computational hardware of about 8-10 times by using cloud HPC over an in-house HPC infrastructure. Using more advanced simulations with larger design freedom we expect to improve the quality of the design reducing the cost per ship set produced, the expected cost reduction is about 2.5% per product. In addition, overall cost reductions and large business benefits are expected due to perceived value by the customer of the electrical wiring manufacturer with respect to the possible quick response to changes and requests, the higher quality of the proposed design and reduced risk of infeasible design. These benefits are not possible to quantify at this point as almost no information is available and each project is different.

4.8.2 Relevance to Fortissimo

With regards to the supply chain introduced in this experiment, the present experiment will motivate companies with similar challenges to shift towards computationally 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 Impact and Business Benefits

A comparison of the current non-optimal process with the automated and optimized process in a cloud HPC environment shows a 2.5% reduction in cost and weight of the wiring system, which is a saving recurring per aircraft and also a lead-time reduction of 90%. This results in:

- Shorter design loops are possible with fast feedback to customers allowing for later incorporation of design changes or less cost of delay and thus more flexibility towards the customer resulting in higher customer satisfaction
- Possibility of including additional routing design variables (i.e. more complex analyses), resulting in even higher cost and weight savings in same amount of time.

Given the successful results achieved by the experiment there is a substantial exploitation potential of the Fortissimo solution into the current engineering industrial practices. This potential is mainly applicable in the field of design engineering, especially for SMEs and Independent Software Vendors (ISVs) as solution providers to larger industry. In particular, the following main key results will be exploited after the 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 on-demand licence management systems into Optimus for adoption of flexible business models
- Integration of on-demand licence management systems into KBE applications for adoption of flexible business models
- Knowledge Based Engineering simulation and workflow setup on a cloud HPC platform via Optimus
- HPC cloud computing integration into Optimus
- Linux based Knowledge Based Engineering (KBE) applications previously available only on Windows machines

The main barrier encountered during the execution of the experiment is regarding the security requirements (permissions on the cluster) needed to deploy and correctly configure the different software packages in the experiment architecture on the HPC cluster. As mitigation a dedicated cluster was setup where these requirements could be met. The 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. Thanks to the extreme flexibility of the approach implemented in this experiment the core partners of experiment 408 on the base of their past experiences 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 for computational power and efficiency.

4.9 Experiment 409: Cloud-based correlation between simulation and tests for mechatronics

The objective of the experiment was to showcase the use of cloud based HPC services to investigate a 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-year-old SME working in test engineering for automotive and aeronautical industries.

Xcos is a block diagram editor used for modelling and simulating dynamical systems. It is a re-implementation of a tool called Scicos (Scilab Connected Object Simulator) with a well-defined semantic. It can be used to interconnect continuous time functions with discrete time ones in the same environment. Events and discrete-timed clocks are managed as well. It is shipped with Scilab, the free and open-source numerical computation library.

Demonstration of a new mode of use for simulation model sharing on a cloud architecture with HPC capabilities - like the Fortissimo marketplace - would open new markets to develop.

The software provider in this experiment was Scilab Enterprises, the official developer and publisher of Scilab and Xcos software. The experiment opens new perspectives for future versions of Scilab/Xcos that will be used in a cloud environment for industrial processes.

The experiment addressed porting Scilab and testing it on the Fortissimo marketplace. Bull Extreme Factory cloud-based HPC infrastructure was the Fortissimo HPC resource used for this project. Scilab Enterprises and Bull worked 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 was carried out, and possible business models were proposed.

4.9.1 Motivation

The experiment's focus was project and confidentiality management when using cloud-based HPC simulation.

Whether it is for fundamental research or new product development, the data from computer simulations is among the most critical for companies and laboratories alike, as it goes right to the heart of their investments, strategy and know-how. For Bull, security is a fundamental component of HPC, and it is even more critically important when using cloud operations: so Bull's extreme factory solutions natively integrate very high levels of security to protect data and exchanges. Physical access to the data-centre hosting the supercomputer is restricted to badged administrators, interconnections are covered by a secure protocol, IP addresses are filtered (only the client's own domain name is allowed) and matched to login credentials, and users can only access a virtual machine, never the physical machine, to ensure total data partitioning. For security reasons, all application installations are carried out by Bull teams and no customer has any way to know who else is working on the cluster and which data or what applications are being used.

This traditional approach to ensure confidentiality on HPC servers is therefore to isolate the actors. However SDI, as subcontractor that builds test benches for big industrial players benefits from having a shared space with their customers. In the test case we studied in the experiment, SDI built a test bench for the brakes of a car manufacturer. The confidentiality issue both for SDI and for the car manufacturer is to be able to give enough data to one another without having to expose their intellectual property. One ff

approach would be to prepare a non-disclosure agreement and give the whole model, however this is seldom done because this would require a heavy enforcement of the agreement that SMEs are not ready to make or risks to give away sensitive data for the car manufacturer. The other approach, mostly used in the case of SDI, is for the car manufacturer to give minimal data to SDI for scenarios they have prepared. This is time consuming for both as SDI then needs to create a car model themselves that fits these data in order to prepare their test benches. This confidentiality issue could vanish if both SDI and the car manufacturer had a shared space and software solutions to be able to simulate the interactions between the brakes test benches and consumer's car model directly without having their intellectual property exposed. The approach selected in the experiment was to use the Fortissimo resource as a neutral place where actors could perform a simulation with all the benefits from the HPC capabilities without worries about their intellectual property.

The two different problematic scenarios SDI has to take into account are the following. In the first scenario, by far the most common, the customer gives SDI only an excerpt of simulation results so as not to divulge their intellectual property. SDI has thus to rewrite the customer model to perform the simulation and check all the parameters of the test case. This is costly, slow and imprecise. In the second problematic case, the customer agrees to share the model, but it is a very complex one, and SDI does not have the computing power in-house to use it. Therefore SDI has to simplify the model or provide the customer with their own simulation model for integration, and this creates confidentiality issues. The issue is similar for the customer once the study is complete.

Both scenarios create an overhead that could be avoided if a neutral place for co-simulation and software that could enforce confidentiality management were available. This requirement has its benefits in all the phases of a study case for SDI:

- In the feasibility phase to answer the call for quotes, SDI could use their existing brake models with the client's car model to find rapidly the potential difficulty that will arise during the realisation phase. Time is an issue here, as call for quotes last only one month.
- During the realisation of the study, the computation time needed to establish the correct parameters for the response of the brake can take 1hour in the easiest cases to 30 hours in the most complex ones in SDI's standard workstations, with 4hours being their current mean computation time. Having access to more computing power helps SDI get more accurate result, by having more simulations, and get them faster, by running the simulation in parallel. This also limits the number of physical test runs during the study validation phase.

The possibility to launch simulations via the Fortissimo HPC cloud resource helped increase the numbers of scenarios one is able 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 companies 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.

By renting the computational time, SDI can compete on more complex calls for quotes. By having a co-simulation space where confidentiality is guaranteed they can build trust in their customers, prove their know-how and avoid the costs of redesigning the customer's model without anyone having to expose their intellectual property.

Scilab Enterprises saw in Fortissimo the opportunity to adapt Scilab HPC capabilities and increase its visibility in the HPC ecosystem. Xcos could benefit from the confidentiality management and from an industrial example of its use on a cloud configuration.

4.9.2 Relevance to Fortissimo

Fortissimo Marketplace creates a neutral, cloud-based environment on which co-simulation at large scale can take place. All the benefits from the high performance computing without the high entrance cost to new players on the domain gives an opportunity for innovative SMEs to boost their business.

The experiment 409 addresses the problem of co-simulation on a HPC Cloud platform with the requirements of confidentiality as its main axis. In the field of mechatronics, a big industrial player can use the marketplace as a neutral place to simulate its system with solutions from subcontractors without any actors having to expose its intellectual property.

The ability to select the number of computation nodes depending on the complexity of the problems to address and the possibility to scale the simulation depending on the current need of each player gives a strong incentive to rent HPC technologies to small players not yet familiar with the HPC world.

4.9.3 Impact and 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 Functional Mock-up Units (FMU).

This approach interested FAW, third largest Chinese car manufacturer and a client of SDI, willing to make sure their intellectual property was protected. On the other hand, 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 are also attracted by the idea. As Scilab is a free open-source software, Scilab Enterprises benefits from the service generated with new innovative solutions as training and specific development.

The Fortissimo experiment gives Scilab Enterprises an example of a successful industrial use of the HPC capabilities of Scilab in its version 5.5.0 (released during the Fortissimo project). The upcoming version 6 of Scilab will give an even greater focus to HPC capabilities with the ability to treat huge volumes of data.

An HPC-Cloud marketplace such as Fortissimo could widen the number of clients for the HPC technologies. Scilab Enterprises estimates the gains to be made with the marketplace as follows:

	2017	**2018**	**2019**
Margin expected	+€200K	+€300K	+€500K

Table 3 Increase in margin expected with the Fortissimo marketplace for Scilab Enterprises

4.10 Experiment 410: Cloud-based simulation of air-quality over cities

4.10.1 Motivation

The European regulations on air-quality protection plans require more and more the need to test and evaluate adaptation/reduction scenarios. Whereas in the past the number of scenarios to be evaluated was 2 or 3, consulting companies and regional air quality agencies are now required to evaluate tens of scenarios. This requires a large increase in the computing capacity for calculation, beyond that which is available in-house to a typical end-user which are public entities (air-quality regional agency or environmental service of a city) or some consultancy companies. Moreover, it is also required to quantify uncertainties associated to air-quality simulations. An increased 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.

Three partners were involved in this experiment:

CERC is a world-leading SME specialising in environmental software development and consultancy, having developed a number of widely-used commercial and bespoke dispersion models, and was the ISV for this experiment. CERC develop and distribute the ADMS-Urban model, which was used in this experiment, based on a classical software licence scheme (one physical licence per client's server). CERC also carry out consultancy projects and produce air-quality forecasts using ADMS-Urban.

NUMTECH is a leader in France on the weather and atmospheric dispersion modelling, working mainly with large companies, local and regional authorities and organizations / research institutes. NUMTECH is a distributor of the air-quality model for France, they realize also consultancy studies with the model and manage operational systems allowing to survey air-quality over cities (final clients which acquire the model are not allowed to integrate it in such operational systems). NUMTECH can be considered as an advanced end-user since it acts like a final client when it performs consultancy. The purpose of this experiment is then to see what can be gained by using cloud-based HPC (for example to reduce IT cost, reduce time to produce results and allow the provision of additional services).

Bull is the HPC platform provider for this experiment. It designs, implements and runs internationally renowned solutions for public and private sector organizations, being both an IT manufacturer and services provider. Its HPC-on-Demand offer extreme factory is the platform that is being used in Fortissimo Experiments and in the upcoming market place.

The experiment allowed CERC and NUMTECH to evaluate the business benefits of cloud-based HPC services for their activities, including consultancy studies as distributor of the model. This evaluation required performing tests and adaptation (including porting of the code, management of the data transfer and developments to automate tasks). Fortissimo was an accelerator to initiate such analysis and interaction between CERC/NUMTECH and the HPC cloud provider (Bull) has helped facilitate such evaluation in a framework of partnership.

In this experiment, exploitation of cloud-based HPC focused mainly on access to CPU power with an associated reduction of cost compared to usage of internal servers as is done currently.

4.10.2 Relevance to Fortissimo

The experiment has provided feedback to Fortissimo on the requirements of an SME software vendor whose code has been in continuous development for a long time and is currently offered as a traditional desktop application. These requirements primarily focus on the features of an API, but also cover licence protection and monetisation models.

In addition, as a result of the work done in this experiment, ADMS-Urban has the potential to become a service to users via the Fortissimo marketplace.

4.10.3 Impact and Business Benefits

For NUMTECH (as an end-user), the impact and business benefits to exploit cloud-based HPC are the following ones:

- To increase the number of studies that can be launched per year due to having access to more CPU power over in-house servers for a similar net cost. This cost includes the purchase of the internal servers and also the cost to maintain them (electricity, IT people etc.). Compared to the use of internal servers, twice as many cores could be used in the cloud-based scenario at the same cost.
- To be able to respond to a specific customer request to produce a study in a very short time (especially when such a study is unplanned). It can be a commercial advantage compared to the current situation, but it is difficult to evaluate the impact in terms of turnover or increased market share.
- For the same price to the customer as present, to be able to carry out uncertainty evaluation or to evaluate more emission scenarios (based on the extra CPU power which can be used on cloud compared to internal servers). Once again, it is difficult to evaluate at this time the impact in terms of turnover or market share.

We can also note that some of the developments carried out during the project (Run Manager tool and improvements in the ADMS-Urban model) have increased profit margins by 5% (due to savings in human time) for large urban studies using internal servers.

For CERC (software provider) and its distributors (like NUMTECH), they can now in principle offer the ADMS-Urban software as a cloud service on a pay-per-use basis rather than requiring a customer to purchase an annual licence and install and run the software locally on their PCs. This enables attractive pricing options for customers needing to use the model only infrequently. However, there are still costing decisions to be made before this service could be offered in practice. To further enhance the user-experience, CERC's Run Manager software has been extended for use with Bull's Extreme Factory allowing users to submit, monitor and retrieve jobs without having to use the Extreme Factory's interface directly.

The ADMS-Urban model has been improved in a number of ways and a new 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. When the model is run the input is automatically reduced 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.

For Bull, a business opportunity has been created with ADMS-Urban from CERC and NUMTECH now available on extreme factory. This enables CERC as well as NUMTECH and other ADMS-Urban distributors to propose air-quality services to new or existing customers with adequate short-term or long-term HPC capacity available immediately on an operational-expenditure basis.

4.11 Experiment 411: Global analysis of risers, moorings and flowlines

4.11.1 Motivation

Principia develops software for the design of systems in the offshore industry (see the picture below). The general trend is to reduce the uncertainty in design by increasing the load cases (number of simulations) in order to capture more precisely the environment and loads experienced by the system or structure over twenty to forty years. In order to perform those calculations and provide solutions to its customers, Principia is striving to improve the efficiency of its software. One of the options is to offer an HPC solution. In this experiment, it was decided to test this option with DeepLines, a finite-element software package dedicated to mooring line, riser and pipeline calculations.

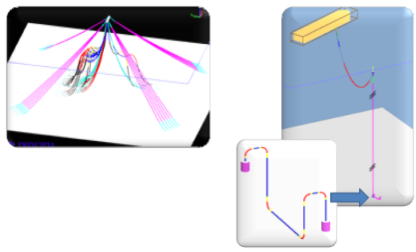

4.11.2 Relevance to Fortissimo

For Principia, this was a first test of HPC to understand:

1. The adaptations needed to the software,
2. The potential gains in terms of schedule on a complete study,
3. The interest of the customer in this type of approach.

Therefore, the Fortissimo project and the collaboration with Bull was a very good opportunity to gain valuable learning experience.

4.11.3 Impact and Business Benefits

The main conclusions are:

- It is possible to propose a first HPC offer for the software DeepLines, but further investment will be necessary to have the same options on HPC as exists right now for a conventional study
- Currently HPC has been adapted to large-scale studies with pre-defined output data, meaning that a first step is required before using HPC on a project to prepare the simulations
- There is definitely an interest from customers to have access at some point in their projects to an HPC solution.

4.12 Experiment 412: Cloud-based simulation environment for CO_2 emission prediction for automotive engines

4.12.1 Motivation

The goal of system simulation in the car industry is to optimize vehicle design in order to improve fuel efficiency and emission, while preserving or improving driving performance at reduced development and production costs.

The majority of projects in the area of vehicle optimization involve large-scale parameter and component variation studies on a limited palette of base vehicle models which requires high CPU power on-demand. Even larger companies struggle with providing sufficient computational resources necessary to accomplish optimization tasks in an acceptable time-frame. The other obstacle in the development process is the traditional node-locked annual-lease licensing mode which requires a large number of software licences for running simulation jobs in parallel.

A distributed cloud computing approach with innovative web-based shared licensing seems to be a promising solution to both of the problems described above.

4.12.2 Relevance to Fortissimo

Independent software vendors have to focus on improving simulation code capabilities and its performance while they usually have a lack of resources and/or expertise for developing scalable computing environments which would utilize the software in an optimal manner. The biggest advantage of the Fortissimo project is bringing ISVs, potential end-users and HPC providers together and providing a development environment in which every participating partner can focus on their own area of expertise and developing interfaces to each other.

Fortissimo provides a effective platform for exchanging ideas and best practice among the stakeholders not necessarily involved in the same projects and experiments. The learning potential and synergy effects will continue to grow in the exploitation phase with the planned introduction of Fortissimo Marketplace as a hub for bringing cloud-computing solutions from different areas closer to the end-users.

4.12.3 Impact and Business Benefits

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 used here is the commercial software tool AVL CRUISE while the result analysis is performed with the software tool AVL CONCERTO.

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

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

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

Cloud-based Industrial Process Improvements

Opposed to that traditional approach, the solution applied in this experiment utilizes cloud-based 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.

Cost Benefits of Using an HPC Cloud

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

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 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 annual software licences and an appropriate computing infrastructure and engineering to run a large scale parameter study, it would cause an investment measured by millions of Euros. This is the point when most of the potential projects are dropped and replaced with cheaper, but less efficient solutions, due to high primary investment costs.

Using cloud-based solutions on the other hand, even taking into account all additional cloud overhead, 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 have run costs reduced by up to 90%. This is the cost range where it can be expected that the companies belonging to the SME segment could participate in the projects which required high CPU power for short bursts.

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.

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

HPC as Opposed to Standard Cluster Offerings

In contrast with standard computing clusters, cloud-based HPC, in this case hosted by HLRS at the University of 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.

Note: this experiment was granted a time extension and is still in progress at the time of writing. Therefore this section describes the latest status at M20 of the Fortissimo Project.

4.13 Experiment 413: Multi-physics simulation in the Cloud

4.13.1 Motivation

Market Background: There is a general trend in the CAE market: more and more software vendors are providing multiphysics simulation capabilities. On the other side, a fast growing number of process engineers for manufacturing and production applications accept and know that they should enhance their currently used simulation environment by such 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 valuable feedback from their tests after a very short time (i.e. few days).

Project Goal: The idea of experiment 413 '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 provide access to various simulation codes, a vendor-neutral code-coupling environment [MpCCI], a valuable and growing number of tutorials, and enough CPU resources to run even bigger jobs. Software licences can be provided from the HPC centre (the HPC centre has a contract with the ISV which allows the use of software by third parties) or by the end-users (the HPC centre re-routes its licence request to an external licence server CPU). Experts in the use of the software tools are also available for support.

Target End-Users and Applications for WP413: The targeted end-users in 413 are manufacturers of electrical components such as transformers, electric motors and electromagnetic-controlled valves and injectors. 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 the complex structure of components and missing (or too expensive) sensor hardware. A combined simulation of the electromagnetic effects as well as the fluid effects can help to simulate these effects in one coupled multiphysics approach. In our experiment we use the following 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 licences 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 licences needed to be organized separately in bi-lateral negotiations between Fraunhofer SCAI and Ansys. During the work for 413 several formal and licence issues had to be solved before the technical setup could be used properly.

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

<u>SMEs</u>

- Only have limited simulation capabilities and experience (licence, hardware, experts)
- Design is driven by experience and basic calculation methods or simplified models
- Not much time to investigate new analysis and design methods

- SMEs need a fully functional solution which has been customised to their specific applications and needs

Large Companies

- Simulation experience and capabilities (i.e. Fluent, other EMAG codes) are 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 the later purchase of new software

The industrial setting

Multiphysics simulation is a general trend in engineering. Multiphysics models allow the examination of effects coming from different physical disciplines by coupling two (or more) simulation codes in a single, but complex application setup. A growing number of engineers are deeply interested in the evaluation of such multiphysics solutions. However a very important hurdle for such software tests is often the installation and configuration of the software base itself:

- 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 often just trivial IT issues, just ‘stupid IT issues’ -- for a CAE engineer they often may lead to a termination of the evaluation if the in-house IT environment is not well prepared for such software tests. .

The improvements in the industrial process

The idea of experiment 413 was 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 monophysics simulations) to run more realistic cases. Finally the HPC centre offers resources for large-scale production jobs.

The overall approach of this experiment thus was not to provide 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 the work was on the evaluation of new cloud or web-based services which give multiphysics newcomers easier access to coupled software on appropriate hardware. Experiment 413 is a multi-CAE application, i.e. coupling and adjusting three software/licences (1 proprietary, 1 open source, and 1 commercial) available via an HPC Cloud, including the development of cloud interface, to model, simulate and visualize a complex problem.

There are only a few steps for an engineer to get all information about the demonstration 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 an evaluation
3. A selection of predefined tutorial cases which can be used online
4. Upload mechanisms to provide own models to the experts at SCAI and Powersys
5. Define values for simulation model parameters
6. Retrieve simulation results and view by MpCCI Visualizer (available via download)

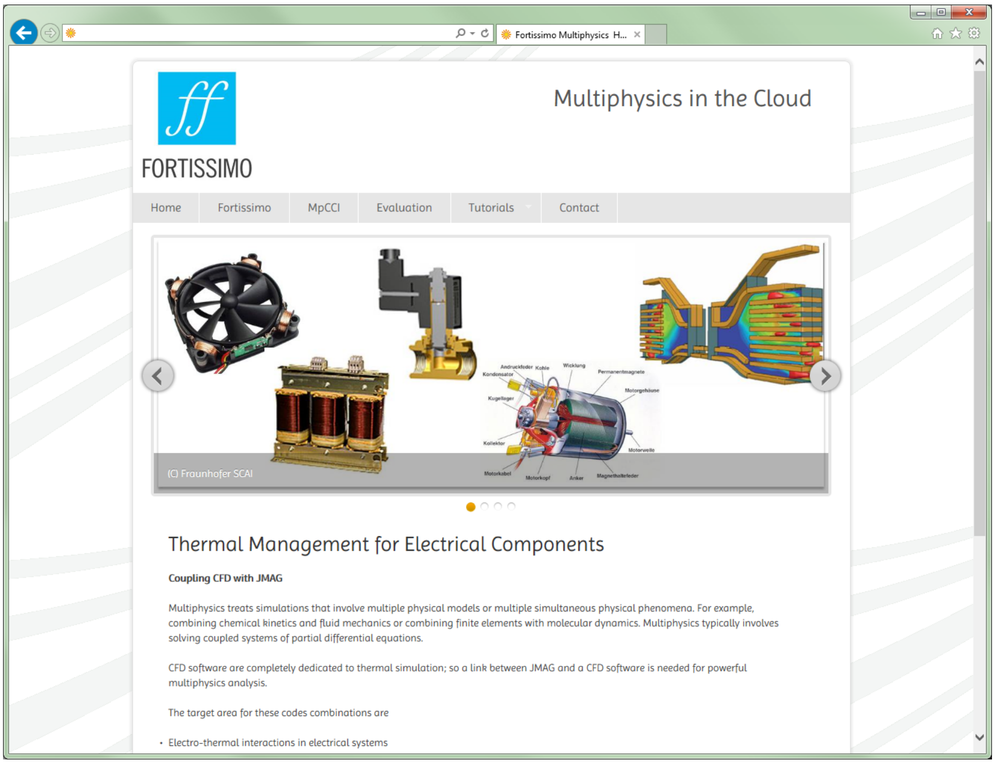

Figure 2: Web-Interface to Wp413

4.13.2 Relevance to Fortissimo

This experiment has demonstrated the need for cloud-based HPC simulations. The major advantage for the multi-physics newcomer will be the ability to test a new simulation environment through a cloud service, instead of installing new and complex software packages locally just for an evaluation run

4.13.3 Impact and Business Benefits

The solution concept of experiment 413 gives SMEs initial access to a novel multiphysics technology. The user can define his problem and the application experts will setup and provide a customized solution environment.

The benefits of such a cloud-HPC based solution are:

- the user does not need to purchase or run their own HPC-hardware

- without any efforts on the engineer's side a pre-installed simulation environment at the HPC centre can be used
- the HPC experts and application service providers can help to handle licence issues
- the engineer can run ready-to-use and customized simulation models

Compared to a local installation the user can save several days of preparation and installation work. The major impact to the CAE market is that such pre-installed CAE solutions can help to push new simulation technologies to existing or even new engineering areas.

4.14 Experiment 414: Cloud-based molecular simulation for industrial chemical engineering applications

4.14.1 Motivation

The objective of this experiment was to showcase the use of cloud-based HPC services to predict the thermodynamic properties of hazardous fluids. Based on molecular force-field methods, powerful predictive methods now 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 e.g. 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 experiment has focused on *ms2,* a molecular simulation tool developed by University of Paderborn that targets the generation of thermophysical properties of condensed fluids). At the moment, the efficient usage of *ms*2 requires expert knowledge and ideally access to HPC systems. The user has to be familiar with the Linux terminal and the HPC-environment. Furthermore, there is no graphical user interface present. Many potential customers (like SMEs) are avoiding the simulation approach in general, because they are deterred by the required expertise and the cost for High Performance Computing equipment and software, even though it can clearly make them more competitive.

4.14.2 Relevance to Fortissimo

Cloud-based molecular simulation will lower the entry barrier for prospective users because they are far more familiar with web interfaces than with Linux-based terminal commands. 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 SMEs in the chemical industry resulting in new business opportunities for the technology supplier.

Due to the volume and complexity of the calculations with the *ms*2 simulation code, users typically rely on high performance computing (HPC) infrastructure. While a single calculation is relatively time intensive, the calculation of a correlation of the simulation data (**E**quation **o**f **S**tate, 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 2 computing nodes consisting of 16 cores in total[1]. In order to acquire enough data to generate the full information set as needed for the accurate prediction of the EOS, around 200 state points need to be evaluated. This means that even with exclusive access to small compute cluster, the calculation would take up to 4000

[1] 16 cores is a good trade-off between the quality of the results and cost.

hours (close to six months). Access to multiple HPC systems through a cloud-based approach would therefore be very attractive.

4.14.3 Impact and Business Benefits

In order to interpolate and extrapolate the simulation data, the cloud-based molecular simulation is able to create equations of state (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.

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 €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.

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 €50,000, 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 molecular simulation study for the data generation is around €80,000 cheaper than an experimental study. The total building cost of a distillation column for example is around €1.5m. Therefore, a saving of €80,000 would result in a 5.4% reduction of the total cost. Note that, depending on the competitive situation, around 3-5 projects of this size are carried out per year.

Alongside with the monetary benefits, the time-to-market will be decreased due to the faster gathering of fluid properties for distillation and absorption significantly.

Experimental study			**Molecular Simulation study**					
Cost per day	*#days*	*total cost*		*Cost per CPU hour*	*#sim*	*Duration (h per sim*	*#Cores per sim*	*Total Cost*
€1,000	100	€100,000	HPC time	€0.15	200	20	16	€9,600
				Cost per day	*#days*			
			Personnel	€550	4			€2,200
			Validation	€1,000	5			€5,000
			Support and Consulting					Subject to Negotiation (X)
SUM		**€100,000**	**SUM**					**~€16,800 + X**

Figure 4 Calculation of the savings for data generation

4.15 Experiment 415: Cloud-based simulation of low-pressure die-casting of copper alloys

4.15.1 Motivation

The objective of this experiment was to demonstrate the cost-effectiveness of using cloud-based HPC for simulating metal-casting processes.

IMR produces foundry equipment for brass alloys and bronze, such as:

- Low-pressure die-casting machines,
- Gravity die-casting machines,
- Core-shooting machines.

The company also offers dedicated consulting services regarding the mould design and manufacturing processes to its customers.

The plants (used in 80% of the cases for the production of taps and valves and in 20% for other high performance artefacts) are not provided with a mould if not requested by the clients themselves. The moulds mounted on the casting machines 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. In the following images are some examples of the resulting artefacts.

The success of the entire system depends on the quality of the piece obtained. Therefore, it is 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 mainly 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 avoids changes of the mould prototype obtained empirically by trial and not with the objective data of the simulation. Using the simulation with HPC reduces by 20% the time of development of the mould and saves 20% of the cost of testing before mass production.

The average cost for a set of moulds for a product designed and tested, is currently about €41,000. The use of HPC simulation thus saves about €8,000 per set and 3 weeks of testing and modifications.

Cost source:	**Current cost:**	**Savings:**
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 on test time)
Drawing modification after testing	€2,000	€1,000 (saving on drawing time)

Mould/core modification	€8,000	€4,000 (saving on modification time)
Final test and sample	€5,000	€2,000 (saving on final test time)
TOTAL	**€41,000**	**€8,000**

- Current cost with the standard development of a prototype mould are €41,000.
- The savings generated by the use of casting simulation with HPC are €8,000.
- The cost for each set of mould done with HPC simulation. (HPC rental cost and personnel involved) is €2,000.
- IMR has about 8 sets of moulds per year to develop (save €8,000 x 8 = €64,000)
- Total saving per year €64,000 by using casting flow simulation.

4.15.2 Relevance to Fortissimo

Before Fortissimo, IMR had only limited knowledge of HPC simulation. Because of that, when simulations were required by the complexity of the mould design in order to aid the mould design, the company outsourced them to an external partner who had the hardware, the software and necessary expertise to perform the simulations.

With the aid of the HPC providers and experts within the Experiment Consortium, IMR has been able to explore and validate and learn about the key aspects and opportunities of HPC analysis and simulation.

4.15.3 Impact and Business Benefits

The licence fees for commercial software like Magma or Procast for use on HPC infrastructures are estimated to be well above €100,000, an investment not affordable by a SME like IMR. Adopting open-source software like Elmer the costs are reduced to those of the implementation on the HPC system and its renting, only for the time necessary for the development of particularly complicated moulds.

Furthermore, renting an HPC infrastructure on an on-demand basis and for short periods, allows the use of hardware and software always up to date at an affordable cost.

Benefits:

The rental of HPC resources and use of open source simulation software:

- Reduces the development time of the moulds
- Reduces the cost and the number of tests before going into production
- Improves the competitiveness of the SME and generates for it an image of high-tech company and enhances its service offered to the customer
- Improves collaboration between those who designed and built the mould and foundry experts.

SMEs offering simulation service will have many other opportunities to offer turnkey systems not only to their own niche industry, but also to new markets and / or products.

The simulation service offered can improve the annual sales of mould sets. (Estimated 4 x €40,000 = €160,000).

New company profile, with high-tech image can improve sales of turnkey plant. (Estimated 2 x €400,000 = €800,000)

4.16 Experiment 416: Cloud-based simulation of structural crash tests

4.16.1 Motivation

The objective of the experiment was to showcase the use of cloud based HPC services to investigate the behaviour of automotive parts and components involving 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 designs, materials and manufacturing processes. The challenge for the supply chain often SMEs, 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. 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 in a future business model.

The experiment addressed porting the commercial software and testing it on the Fortissimo cloud-based HPC infrastructure. ESI-Group, Bull and GENCI worked 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 was carried out, and possible business models were investigated.

The Green Vehicle Initiative is a major influence for the automotive industry to answer environmental requirements and regulations. This challenge calls for disruptive innovation in vehicle design and materials. 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 for the emergence of new industry leaders like Gestamp, who have developed a new grade of steel using an innovative manufacturing processes known as press hardening. It is one of the major innovations introduced to the market capable of reducing vehicle weight with recognized world-wide advantages.

4.16.2 Relevance to Fortissimo

However, the industrialization and the customization of this innovation also required a new design and development process in order to enable engineering and 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 requires 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 the support of a collaborative decision making process during the product development phase; including engineering stakeholders from 3 different disciplines and regulations requirements: crash & safety, NVH (noise, vibration and harshness) and static stiffness. Figure 1 illustrates the drawback of the silo approach and the single core model advantages.

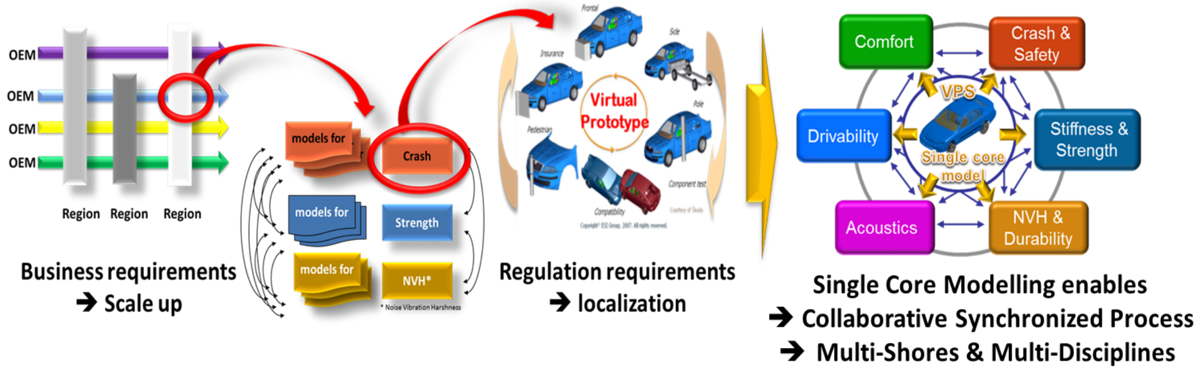

Figure 5: Silo Approach and Single Core Model advantages

Single Core Modelling will reduce the CAD-CAE synchronization efforts, and eliminate the conversion & consolidation between different CAE-CAE Models for Design reviews and trade-offs. Today's approach is a very consuming task and a 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 realise this objective. Cloud-based HPC, as provided by Fortissimo, 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, shorten the time to market, improve quality at competitive pricing.

4.16.3 Impact and Business Benefits

The benefits will be measured in 2 distinct ways: First in terms of time saving thanks to the leverage of HPC resources which enable the introduction of the single core modelling as new collaborative engineering approach to support decision making process. A second natural benefit will consist of overall financial gains which can be obtained by using the HPC Cloud for only the portion of time that was necessary compared to provisioning hardware, personnel and licences for the full year around. Quantifying this financial benefit is very dependent on overall simulation usage at the customer site , and utilization ratios of its HPC infrastructures.

Single core modelling enables the handling of only one body in white numerical model for several applications (four applications in our example). Thus, when the design, thickness or material of the 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. The time spent on setting up and afterwards updating the model is then divided per 3.5 to 4 (see the 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-effective 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

Table 4: Detailed benefits of Single Core Approach

The use of HPC cloud services will also be a huge benefit for multi-site companies. Indeed, Gestamp R&D Centres located in Paris (France), Luleå (Sweden) and Barcelona (Spain) have increasing needs in handling projects over several locations. Therefore, the HPC cloud approach enables them 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.

4.17 Experiment 417: Cloud-based Computational Fluid Dynamics Simulation

Experiment 417 focuses on Computational Fluid Dynamics (CFD) and Computer Aided Engineering (CAE) which are vital areas 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.

Cineca, iconCFD, NTUA and Koenigsegg have collaborated on the experiment which has focused on setting up a cloud-based HPC service for Computational Fluid Dynamics Simulation. The specific case refers to the use of accurate, and consequently very computer intensive, transient detached-eddy simulation solvers for drag and lift prediction of supercars. However the experiment has explored and demonstrated the general potential of offering cloud-based CFD simulation services, providing access to demanding applications for SMEs which until now could not afford, in terms of acquiring the required computational resources, to incorporate such applications in their business.

By providing CFD simulation services running on a cloud infrastructure making use of High Performance Computing systems, experiment 417 has contributed in the scope of Fortissimo project to make advanced simulation accessible to SMEs through the realization of a marketplace where hardware, expertise, applications, and tools are easily available and affordable on a pay-per-use basis.

4.17.1 Motivation

The CFD market can be characterised as having customers with two distinctly different drivers: heavy CFD users with high budgets (typically Original Equipment Manufacturers, OEMs) who are moving into more complex physics with improved accuracy and 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 licences limited for budget reasons in term of concurrent sessions or parallel threads capability, often not adequate to the science they attempt to simulate. They are therefore often obliged to “dumb-down” the level of simulation refinement they may achieve or need for real breakthroughs.

Both customer groups share 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 not only increases the cost effectiveness but accelerates the uptake and application of more complex CFD, which in many cases is required to capture the science necessary to influence product development.

Experiment 417 has used the Fortissimo infrastructure to port existing simulation tools over HPC cloud infrastructures to enable end users to have instant access to a “one-stop-shop” which offers simulation services on demand, while users are able to obtain the results of their simulations in less time, and with a pre-defined cost per simulation (based on various criteria like CPU-hours, simulation time etc.). More specifically, experiment 417 has pursued 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 has brought together Cineca, a leading HPC centre, iconCFD, an ISV/business owner/simulation code support provider, NTUA, a research organisation with expertise on system integration design, interoperability, decision-support systems and business modelling and Koenigsegg, a niche manufacturer of high-performance sports cars, which requires greater access to affordable, scalable, HPC resources for its CFD requirements, as iconCFD's end user.

4.17.2 Relevance to Fortissimo

The equations underpinning 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 CPUs working on the same overall problem by distributing equations and parts of models into 'sub' calculations). As hardware costs have plummeted over the past ten years, the real cost of CFD simulations might have been expected to decrease. However, although it is cheaper than ever to procure multiple CPU systems (e.g. Linux clusters), CFD licence costs did not decrease. Thus, as a company adds extra CPUs to its computer resources, the software vendors have priced their solvers on a per-parallel thread, or groups of parallel threads, 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. However, cloud-based HPC has become a cost-effective alternative to large clusters, offering small companies the prospect of being able to exploit CFD simulations to the full.

The need for cloud-based HPC simulation is of utmost importance in this experiment, as by moving simulation of CFD models onto a cloud-based HPC environment, processing times can be impressively shortened. 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.

The experiment started from the design need of the industrial end-user, Koenigsegg, building a CFD/CAE process workflow, based upon iconCFD FOAMpro CFD software, to be integrated with the Fortissimo infrastructure, so as to give end-users a seamless access to Cloud HPC resources for their CFD-based problem solving activity.

4.17.3 Impact and Business Benefits

The provision of a Cloud-based HPC solution can bring 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 solvers) than ever before. In addition, iconCFD has already demonstrated the maturity of automatic design optimisation technologies, that when combined with open-source CFD solver technology, (industrially supported by iconCFD) can offer significant cost advantages over current competition and existing manual CFD simulation methods.

The benefits that can be gained 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.

The customer value and benefits from HPC-cloud CFD simulation 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.
- 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 SMEs approaching CFD.

4.18 Experiment 418: Cloud-based CFD turbo-machinery simulation

4.18.1 Motivation

Experiment 418 used the Fortissimo infrastructure to port an existing commercial software (the ANSYS CFD suite) over HPC cloud infrastructures and to build on top of that a customized solution to enable end users with limited knowledge of Computational Fluid Dynamics (CFD) and HPC to explore an appropriate parameter space of the design of turbomachinery (centrifugal pumps) to achieve the desired optimization level. The solution may be implemented on the Fortissimo marketplace as a service, to complement Enginsoft's service portfolio. The experiment brought together Cineca, the Italian High Performance Computing (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.

More specifically, experiment 4.18 pursued 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.

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 performance. In these terms, 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.

4.18.2 Relevance to Fortissimo

While CFD is already a common tool in many SMEs dealing with centrifugal pumps design, with Reynolds-Averaged Navier-Stokes (RANS) simulations run on a workstation, today the designer has to cope with two key challenges to stay on the market:

- competitive market sale target
- strict energy efficiency regulations

The designer is thus driven to "raise the bar" and knows that a few percent points gained in performance often make the difference in a sale scenario. An accurate prediction tool coupled to a computing system easily accessible, like cloud computing, may allow SMEs to gain a competitive edge.

However, detailed numerical simulation of centrifugal pumps is not easy due to some difficulties: complex geometry, unsteadiness, turbulence, secondary flows, separation, boundary layer, etc. These aspects are not satisfied by fast RANS models, but 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 software with high parallel efficiency and 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.

This approach implies investments in computing infrastructure and specific competences that are 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.

4.18.3 Impact and Business Benefits

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, Enginsoft expects to increase its market share by at least 1% with a total profit of €100K representing a positive return on investment in only six months.

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 summarized as follows:

- Increased speed, as a result of obtaining simulation results within hours instead of weeks.
- Increased flexibility, occurring mainly through service 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 SMEs 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 are shown in the table below

	Year 1	Year 1+2
Personnel	€50,000	€100,000
Training	€10,000	€10,000
Hardware and maintenance	€20,000	€30,000
Software licences	€50,000	€100,000
Total investment	**€130,000**	**€240,000**

We may add then an estimate cost for the following years of about €120K 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 a EnginSoft know-how engineer to set up the case and a HPC cloud service fee (per core_hour consumed) the cost would be:

	Per project
Token know - how engineer	€15,000
75000 HPC core_hours	€15,000
25% cost of internal design Engineer	€10,000
Total investment	**€40,000**

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.

Hence a break-even point between internal activity and outsourcing to cloud is reached after 7 year in the worst case, and it is more likely, given the volatility of workforce and the lifetime of an HPC system, that cloud-based HPC will always be more cost-effective.

4.19 Experiment 419: CFD simulations of eolian snow transport for civil engineering applications

4.19.1 Motivation

This experiment demonstrated the business case for using cloud based HPC for assessing the loading due to snow on buildings, with a view to improving building safety in areas where heavy snowfall is likely.

Roof collapses due to accumulated and drifting snow (see Figure 6) 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 build-up of snow and calculate the corresponding 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.

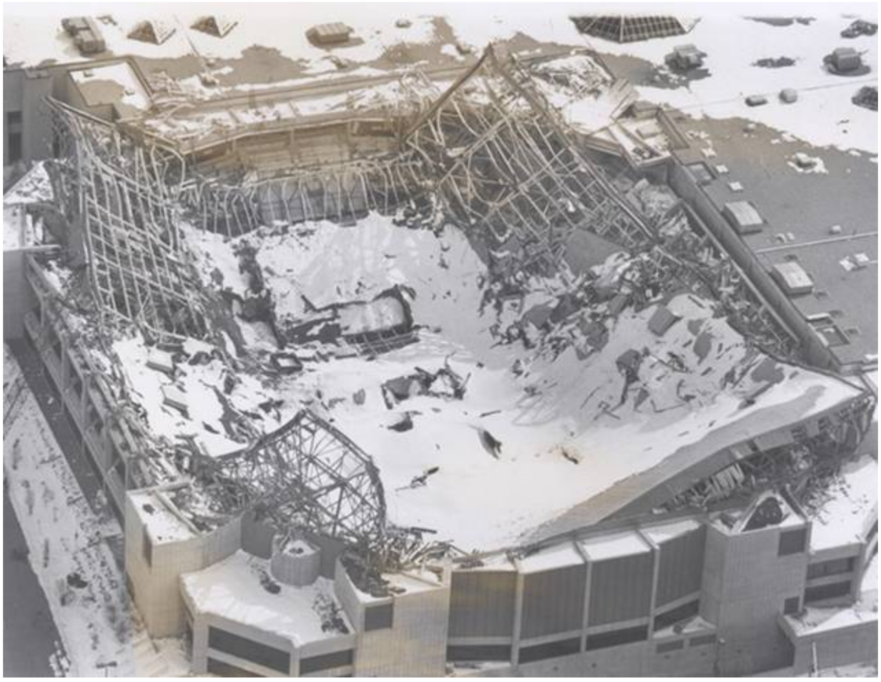

Figure 6 Roof collapse in 1978 in Hartford, Connecticut

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 number of computing cycles for unstructured meshes is an order of magnitude larger than for structured meshes, hence the need for even larger HPC systems.

4.19.2 Relevance to Fortissimo

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 costs of analysis and computational that is typically €10K-50K. Moreover, drifting snow analysis can be completed within a few weeks using an HPC system. 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. A pay-per-use billing model of Cloud based computing is ideal for cost accounting in consulting projects.

4.19.3 Impact and Business Benefits

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 demand 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.

An extensive test programme and success criteria were elaborated in order to properly evaluate the different computational strategies available so the cost savings can be optimized.

Results/conclusions/benefits:

- The commercial feasibility was proven for a snow-drift simulation service on a HPC-on-demand infrastructure, to be used by an expert consultant.
- Effective work flow strategies and typical solver and platform parameters were investigated to gain understanding of the model complexity, so that a realistic estimate could be made of the cost in resource and time.
- Cost estimates can be guaranteed, thus minimizing the risk of budget overrun in the computational resource bin, and the risk of project delay by faulty planning.
- The interface to the on-demand infrastructure was further developed to pave the way towards a flexible and easy-to-use high-performance simulation platform, including flexible access through a web browser, remote visualisation of large data sets and dedicated pre and post-processing capabilities.
- The marketplace allows the provision of a consulting service for an SME-type expert user, but it also opens up the way to a web portal for end users of the snow drift simulation package with support of the SME consulting company.
- 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.

4.20 Experiment 420: Cloud based access to commercial computational chemistry codes

4.20.1 Motivation

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).

However, the *material 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 sulphide clusters, as used *e.g.* in catalyst-enhanced HydroDesulphurisation (HDS) 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 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, analyse the results within the remote environment, as they become available.

This web front-end to cloud-based computing, independent from the Fortissimo marketplace, will allow industrial modellers to perform demanding simulations on cutting-edge platforms, without the need for deep expertise or investment in expensive computing resources. The availability of easy-to-access modelling software in 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.

4.20.2 Relevance to Fortissimo

HPC clouds as offered by the Fortissimo partners present a number of key advantages with respect to generic cluster offerings:

- A cloud service such as Amazon's EC2 may be cost-efficient when running web services, but the kind of intense, peak computing power demanded by HPC applications quickly becomes comparatively expensive. Dedicated HPC providers offer better pricing options for heavy "number crunching", where often a discount is given when large number of compute cycles are bought.
- Another limitation is of technical nature: scientific applications such as ours require fast interconnects, high bandwidth and low latency, while standard cloud provider usually offer slow communications between nodes, and therefore worse parallel scalability.
- Another advantage of HPC platforms *vs* standard ones lies in the level of HPC expertise of the support staff. Users of scientific computational software have special requirements and can outsource comparatively more expertise to HPC providers, e.g. installing the software for instance. Customers of standard cloud 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". Allowing for more aggressive optimization and therefore shorter run-times.

Fortissimo offered an invaluable opportunity to collaborate with a not-for-profit HPC provider, porting our software to their platform and kick starting the process of drafting a cloud-based business model.

4.20.3 Impact and Business Benefits

The benefits associated with using a HPC cloud are foremost 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 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 cost around €60,000, and a commercial licence for a typical materials modelling package can go well over €80,000. On the other hand, commercial HPC providers can charge around. €40,000 per year for the exclusive usage of 128 cores, clearly a more cost-efficient option. In general, end-users of cloud services avoid hardware capital or leasing costs and hardware maintenance and upgrade costs, and have reduced up-front software licence costs.

Cloud-computing vendors can manage all of their customers on a single instance of the software, which allows the amortisation of 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 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 managing it (maintenance).
- 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 NVIDIA Tesla K40m. The price for fully dedicated CPUs includes an important discount with respect to a PAYG model.

<table>
<tr><th></th><th>In-house</th><th>HPC-CLOUD</th></tr>
<tr><td>Manpower/year</td><td>€50,000</td><td rowspan="2">€20,000</td></tr>
<tr><td>Hardware/year</td><td>€12,000</td></tr>
<tr><td>Software licences/year</td><td>€16,000</td><td>€20,000</td></tr>
<tr><td>Total/year</td><td>€78,000</td><td>€40,000</td></tr>
</table>

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 licences). It is easy to see how this experiment has convinced Albemarle to allocate a sizeable budget for HPC cloud computing next year. The resources at SURFsara can be used for pre-competitive studies but is not available to industry afterwards. An up-and-running Fortissimo marketplace could fulfil their needs.

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 licence costs associated with each option.

All those benefits (technical and financial) have convinced this experiment's end-user, Albemarle Corporation, to allocate a sizeable budget for HPC-cloud computing next year.

5 Concluding Remarks

The initial tranche of experiments in Fortissimo resulted in a rich set of outputs with considerable exploitation potential.

The results of the experiments strongly support the case for cloud-based HPC and the Fortissimo approach. Significant business benefits have been identified by the experiments (in many cases these have been directly quantified) and useful suggestions for how the Fortissimo Marketplace should develop have been made. Furthermore many of the experiments have concrete plans to make business solutions available through the Marketplace.

The next stage of the project will be to develop the Marketplace according to new requirements generated by the first round of experiments, and to continue evaluating and enhancing the Fortissimo approach by means of two further sets of experiments selected by open calls.