Smart Manufacturing: CAE as a Service, in the Cloud

Wolfgang Gentzsch

The UberCloud Inc.

1 Introduction

The benefits for small and medium size enterprises (SMEs) of using high performance computing (HPC) technology within their design and development processes can be huge, such as enormous cost savings; reducing product failure during design, development, and production; develop optimized processes; achieve higher quality products to keep existing and gaining new customers; and shorten the time to market. Potentially, all this can lead to increased competitiveness and more innovation.
However, less than 5% of manufacturers are using HPC servers for computer simulations to design and develop their products, according to the two studies, 'Reflect' and 'Reveal', from the US Council of Competitiveness, [1]. The vast majority (about 95%) of the companies perform virtual prototyping or large-scale data modelling still on their desktop computers (workstations or laptops). But, 57% of these companies said that they have application problems that they can't solve with their existing desktop computers, because their desktops are too slow for the problems they want to solve, or because geometry or physics are too complex and need more memory than is available from their desktop. Therefore, most of these companies have a real need for high performance computing.
There are two realistic options today how to acquire additional HPC computing power beyond what is available from the desktop system. One option (which is widely proven to work) is buying an HPC server which is many times faster than what engineers currently have available on their desk. However, for many companies, especially SMEs, buying a large HPC server is often not a viable alternative. One reason for this is the high and complex Total Cost of Ownership, TCO, as demonstrated from IDC [2] in the following figure:

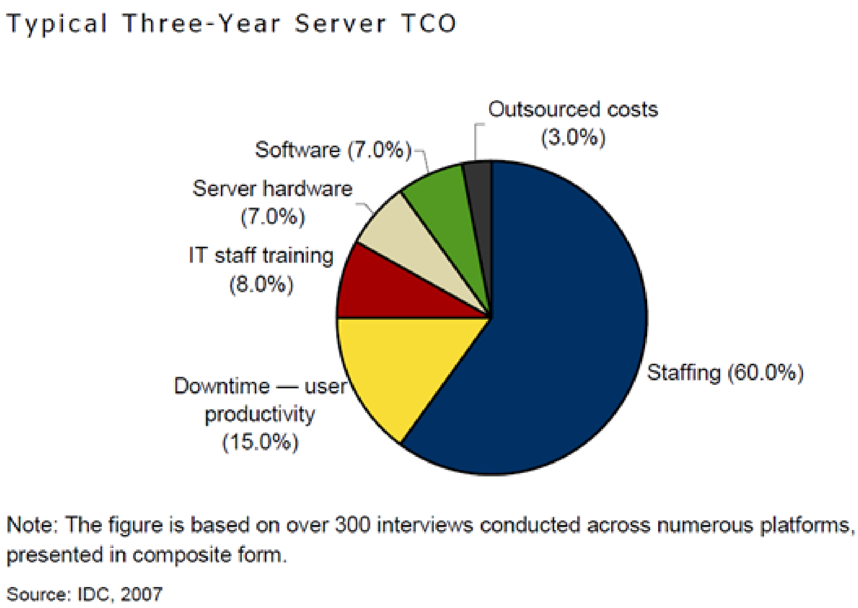

Fig.1: Total Cost of Ownership (TCO) of an HPC server over three years.

In addition to the high cost of expertise, equipment, maintenance, software, and training, there are often long and painful internal procurement and approval processes, and additional skills and manpower are needed to operate and maintain such a system. Still, buying an HPC server is one possible option for an SME, and several HPC vendors have developed a complete ecosystem of HPC products, solutions, and services, and developed the corresponding go-to-market strategy to sell HPC

servers to those SMEs who have decided to buy their own HPC system. Therefore, in this study, we will not discuss this option further, but just investigate the second option, in the following.
The second viable option for SMEs to experience the benefits from HPC without having to buy and operate their own HPC system themselves is recently offered by cloud computing. HPC in the Cloud allows the SME engineers to continue using their own desktop system for daily design and development work, and to submit (burst) the (sometimes much) larger, more complex, more time-consuming jobs into the cloud. Additional benefits of the HPC Cloud solution are on-demand access to ‘infinite’ resources, pay per use, reduced capital expenditure (CAPEX), greater business agility, higher-quality results to keep existing and gaining new customers, lower risk, lower product failure rate, and dynamically scaling resources up and down as needed.
In the following, we present an overview of the status and trend of HPC in the Cloud for the manufacturing SME market and look at the competitive landscape of current Cloud resource and software providers. We include an outlook on how the UberCloud Experiment can accelerate the process of go-to-market and customer acceptance.

2 The HPC Cloud

The architecture of an HPC Cloud is designed to run specific HPC workloads from various science and engineering application areas. Still standard cloud services can address a portion of the HPC market. Those applications are usually embarrassingly parallel (like parameter studies with varying parameter input) and have low I/O requirements (light I/O). Careful analysis of application requirements is needed in order to determine the effective HPC performance in standard cloud offerings. However, many applications developed and optimized for HPC systems require additional system features in the Cloud as well, such as

- large capacity and capability, application software choice, and physical or virtualized environment depending on performance needs. The use of high performance interconnects and dynamic provisioning can offer cloud features while maintaining HPC performance levels.
- high performance I/O is often necessary to ensure that many I/O-heavy HPC applications will run to their fullest potential. As an example, pNFS might provide a good plug-and-play interface for many of these applications. Back-end storage design, however, will be important in achieving acceptable performance.
- fast network connection between the high performance cloud resources and the end-user’s desktop system. Scientific and engineering simulation results are often in the range of many Gigabytes to a few Terabytes. Additional solutions here are remote visualization, data compression, or overnight sending a disk with the resulting data back to the end-user.

Based on these observations, performing HPC in the cloud is indeed possible. Some (mostly ‘embarrassingly’ parallel) applications perform well on existing standard cloud, but many HPC applications cannot be shoehorned into existing standard cloud solutions. Clouds that are specifically designed for HPC are needed, and indeed represent a viable solution for many organizations, especially if the number of cores required is moderate. In addition, other issues may need to be discussed before an HPC cloud can deliver low-cost and flexible HPC cycles.

3 HPC Cloud Adoption for Small and Medium Enterprises

In this chapter we concentrate on “Digital Manufacturing” which describes the use of advanced computing technologies for simulations to guide engineering and production in areas like structural analysis, aerodynamics, fluid flows, crash-worthiness, environment and stress testing, process engineering, manufacturability, to improve quality, achieve faster time-to-market, reduce costs.
Most of the market research in this area is still focusing on just HPC, the reason can be found in an Intersect360/NCMS study on Modeling & Simulation at 260 U.S. Manufacturers [3] which we believe is still valid today. This study shows that in 2010 61% of the companies with over 10,000 employees are using HPC, but only 8% of companies with fewer than 100 employees are using HPC systems for their product design cycle, for designing, analyzing, optimizing, and validating new products and prototypes. On the other hand, 72% of desktop CAE users see a competitive advantage in adopting more advanced computation. Obviously, there are a few strong roadblocks towards adopting HPC especially in the mid-market. These roadblocks are, as mentioned already above, mainly the high Total Cost of Ownership and the cumbersome procurement processes for HPC servers.
Because of low penetration of HPC in manufacturing, use of HPC in the Cloud in this market is still in its infancy. There are still today many publications in the technical press (Digital Manufacturing Report, Desktop Engineering, Bio-IT World, etc.) which repeatedly discuss the basics of cloud computing

(even its definition and categorization) and especially address the different roadblocks, although sufficiently known in the meantime. Today, there are several **important trends and initiatives which foster HPC in the Cloud** also for the manufacturing industry SME market, in the near future:

- The benefits for the design & development engineers and their companies from using technical (HPC) computing in the Cloud are undoubtedly huge (as explained above).
- The growing acceptance of general enterprise cloud computing services in the SME community will also generate a strong acceptance of HPC related services for engineering simulations in the cloud, with a certain time delay. Because cloud services are widely accepted and used today at the enterprise level (e.g. for ERP, CRM, administration), the acceptance and use of Cloud services in the companies R&D department seems to be straight-forward.
- Large manufacturing companies more and more expect their supply chain partners to perform high-quality end-to-end simulations on HPC systems.
- And, last but not least, US and international initiatives like the Missing Middle, the National Center for Manufacturing Sciences, and the UberCloud HPC Experiment, and electronic magazines like HPCwire, the Digital Manufacturing Report, Desktop Engineering, Bio-IT World, and the HPC Magazine, are currently creating strong awareness for the benefits of in-house HPC **and** HPC as a Service (HPCaaS) in the manufacturing industry SME market.

3.1 A Growing Number of HPC Cloud Service Providers

Over the past five years we saw hundreds of cloud services providers coming into market, offering hardware resources, software, and expertise related services. To just mention a few cloud resource providers which offer especially HPC services in the Cloud: Amazon AWS, CloudSigma, Fujitsu TC Cloud, GOMPUTE, MEGWARE, Microsoft Azure, Nimbix, OCF, Oxalya/OVH, Penguin Computing, Rescale, Sabalcore, Serviware/Bull, SICOS, TotalCAE, transtec, and many more providers here [4].
At the same time as commercial cloud service providers evolve, we see a strong interest of a growing number of supercomputing centers offering access to their HPC clusters to (local or regional) industry, on a pay-per-use basis. Examples are CESGA Supercomputing Centre in Spain, CILEA in Italy, FCSCL Spain, Georgia Tech in the US, GRNET in Greece, HSR in Switzerland, Monash University in Australia, NCSA in the US, Ohio Supercomputer Center in the US, Rutgers University, San Diego Supercomputer Center in the US, and SARA in The Netherlands.
Many ISVs are still quite concerned that the cloud computing paradigm could potentially disrupt their existing business model based on selling annual node-locked or floating licenses. Their concern is that their customers might turn to on-demand cloud-based pay-per-use software services **instead** of using the software on their own systems. However, the contrary is true: cloud computing has the potential to increases the revenue for ISVs simply because engineers turning to the Cloud will (have to) continue to use the application software license on their workstations, for business as usual, day-to-day work as demanded by the typical and standard R&D scenarios. The Cloud enables ADDITIONAL opportunities for the engineers to do more, faster, better simulations, see the example in [5]. For all these additional scenarios, the ISV sells additional licenses, on demand, pay per use, paid by the hour, or via subscription contract, resulting in additional business, on top of the workstation license business.
Even for those companies which already have an HPC cluster in their computing center, bursting into a Cloud offers higher efficiency and flexibility for the engineering team and for the computing center, and results in additional license-on-demand business for the ISV. And finally, ISVs offering application software on demand are able to attract new customers who start fresh with computer simulations, and who never would think of buying a license for just a few simulations.

3.2 HPC Cloud Software Licensing

HPC Cloud software licensing is still considered as one of the major roadblocks for the wider adoption of cloud computing especially for SME manufacturers, as stated by a poll during an UberCloud Webinar in June 2013: the ISVs' slow adoption of more flexible (on-demand) licensing models for the cloud was the major concern of 61% of respondents.
Beside those software providers who are slow in providing on-demand licenses, for example those who still are providing their keys via dongles, the majority of especially the larger software providers are either working on a professional on-demand licensing strategy, or they already offer Software as a Service. In the first group of those which are preparing software licensing on demand are ANSYS and Simulia, and in the second group we find Autodesk with Sim360 Pro and CD-adapco with Power on Demand. Autodesk Sim360 offers a subscription based on-demand licensing model, see [6], while CD-adapco has a token-based Cloud offering [7].
A large number of major software providers with a focus on digital manufacturing or HPC tools with cloud features are currently participating in the UberCloud HPC Experiment, among them: Acellera,

Adaptive, Advanced Cluster Systems, Advection Technologies, AMPS Technologies, ANSYS, Autodesk, BGI, BlackDog Endeavors, Bright Computing, CAELinux, CEI, Certara, CHAM, Ciespace, Cloudsoft, Cloudyn, eXact, CPUsage, Cycle Computing, Datadvance, ELEKS, Equalis, ESI Group, ESTECO, Expert Engineering Solutions, Fidesys, Flow Science, Foldyne Research, Friendship Systems, Gompute, GPU Systems, HCL Infoystems, HPC Solves, HPC Sphere, Kitware, Kuava, Landmark (Halliburton), MapR Technologies, micrOcost, migenius, MSC Software, Nice Software, Nimbus Informatics, Numerate, Open Source Research Institute, Ozen Engineering, Personal Peptides, Phenosystems, PlayPinion, Qtility Software, QuantConnect, Rescale, RMC Software, SimScale, Simulia, Stillwater Supercomputing, TECIC, TotalSim, TYCHO, Univa, and Visual Solutions.

3.3 Application Performance in the Cloud

On one end of the application spectrum there are so called massively (or 'pleasingly') parallel applications like e.g. parameter studies in digital manufacturing and computational biology (drug design) with the application code running many copies in parallel on many cores, with each parameter job running on **one** core; or there are applications running on many servers, each parameter job running on one server (and moderately parallel on the server's cores). All these applications could be well-suited even for standard enterprise cloud servers.
On the other end of the spectrum, there are parallel (distributed) applications with tightly-coupled communication needs among the parallel tasks and/or high scalability that would perfectly run on HPC servers. The third class of application codes has mid-size mid-scale parallel requirements; these codes can easily run on one server's parallel cores.
For a cloud datacenter to provide a single-point of access to this heterogeneous HPC cloud consisting of enterprise and HPC servers, three components are necessary: a user-friendly portal providing access to this cloud; an intelligent resource manager scheduling the different application jobs to suitable systems of the HPC cloud; and the different computing resources ideally preloaded with the requested application codes.

4 UberCloud Experiment Accelerating HPC in the Cloud

Cloud Computing and its emerging technologies, such as virtualization, web access platforms and their integrated toolboxes, solution stacks accessible on demand, automatic cloud bursting capabilities, and more, enable research and industry to use additional computing resources in an elastic and affordable way, on demand. Now, the UberCloud HPC Experiment provides a platform for researchers and engineers to explore, learn, and understand the end-to-end process of accessing and using HPC Clouds, to identify the concerns, and resolve the roadblocks, details see [8]. End-user, software provider, resource provider, and an HPC expert are collaborating in a team and are guided through a 23-step process, jointly solving the end-user's application problem.
Since July 2012, HPC Experiment has attracted 2200+ organizations from 72 countries (status end of April 2015). The organizers were able to build 172 of these teams, in CFD, FEM, and Computational Biology, and to publish more than 60 articles about the UberCloud initiative including many case studies reporting about the different applications and lessons learned. Then, the UberCloud University and a virtual Exhibition [9] have been added, and two Compendiums (sponsored by Intel) with together 40 case studies have been published in 2013 and 2014, [10].
Inspired by the results of the Magellan Report [11], the idea to this UberCloud HPC Experiment came up in May 2012 during a discussion about the acceptance of cloud computing for Enterprise versus for High Performance Computing: while adoption of cloud computing in the enterprise market is rapidly growing (according to Gartner, annual growth rate is 41.3% through 2016, [12]), the awareness and adoption of cloud computing in HPC and Digital Manufacturing is still very slow, mainly due to the roadblocks on the way to accessing and using HPC Cloud resources, such as inflexible software licensing and slow data transfer, security of data and applications, and lack of specific architectural features missing in HPC clouds (resulting in reduced performance). The idea of the UberCloud HPC Experiment was to find out more about the end-to-end process of bringing engineering applications to the cloud, and to learn more about the real roadblocks and how to overcome them. The UberCloud HPC Experiment then started in July 2012 with a Call for Participation in HPCwire.
To start, let's define what roles each stakeholder has to play to make service-based HPC come together. In this case, stakeholders consist of industrial end-users, resource providers, software providers, and high performance computing experts.
The industry end-user: A typical example is a small or medium size manufacturer in the process of designing, prototyping, and developing its next-generation product. These users are prime candidates for HPC-as-a-Service when in-house computation on workstations has become too lengthy a process,

but acquiring additional computing power in the form of an HPC server is too cumbersome or is not in line with IT budgets. HPC is not likely to be the core expertise of this group either.

The application software provider: This includes software owners of all stripes, including ISVs, public domain software organizations and individual developers. The experiment usually prefers rock-solid software, which has the potential to be used on a wider scale. For the purpose of this experiment, on-demand license usage will be tracked in order to determine the feasibility of using the service model as a revenue stream.

The HPC resource providers: This pertains to anyone who owns HPC resources, such as computers and storage, and is networked to the outside world. A classic HPC center would fall into this category, as well as a standard datacenter used to handle batch jobs, or a cluster-owning commercial entity that is willing to offer up cycles to run non-competitive workloads during periods of low CPU-utilization.

The HPC experts: This group includes individuals and companies with HPC expertise, especially in areas like cluster management. It also encompasses PhD-level domain specialists with in-depth application knowledge. In the experiment, experts will work as team leaders, with end-users, computer centers, and software providers, to help glue the pieces together.

For example, suppose the user is in need of additional compute resources to increase the quality of a product design or to speed up a product design cycle, say for simulating more sophisticated geometries or physics, or for running many more simulations for a higher quality result. That suggests a specific software stack, domain expertise, and even hardware configuration. The general idea is to look at the end-user's tasks and software, and select the appropriate resources and expertise that match certain requirements.

Then, with modest guidance from the UberCloud Experiment, the user, resource provider, and HPC expert will implement and run the task and deliver the results back to the end-user. The hardware and software providers will measure resource usage; the HPC expert will summarize the steps of analysis and implementation; the end-user will evaluate the quality of the process and of the results and the degree of user-friendliness this process provides. The experiment orchestrators will analyze the feedback received. Finally, the team will get together, extract lessons learned, and present further recommendations as input to their case study. A first Intel-sponsored Compendium with the 25 best case studies out of the first 60 teams appeared in June 2013, [13].

HPC Experiment Status: The current status, as of April 2014, is quite impressive. 172 teams have been formed around industry end-user applications running on remote HPC center and HPC Cloud resources. Over 2500 organizations from 72 countries are actively or passively participating. The end-to-end process of taking applications to the cloud, performing the computations, and bringing the resulting data back to the end-user has been partitioned into 23 single steps which the teams closely follow on the Basecamp collaboration environment. An UberCloud University has been founded providing regular educational lectures for the community. And the one-stop UberCloud Exhibit [12] offers an HPC Services catalogue where the community members can exhibit their cloud related services or select the services which they want to use for their team experiment or for their daily work. Many UberCloud Experiment teams publish their results widely, in the meantime, see e.g. article from Sam Zakrzewski and Wim Slagter from ANSYS about: On Cloud Nine [13]. Finally, for two years in a row, in November 2013 and 2014, The UberCloud received the HPCwire Readers Choice Award for the best HPC Cloud implementation, [14].

5 Use cases from the UberCloud HPC Experiment

As a glimpse into the wealth of practical use case results so far, here are three out of 172 experiment teams demonstrating the wide spectrum of CAE applications.

5.1 Team 2: Simulating new probe design for a medical device

The (anonymous) end user's corporation is one of the world's leading analytical instrumentation companies. They use computer-aided engineering for virtual prototyping and design optimization on sensors and antenna systems used in medical imaging devices. Other participants in this team were software provider Felix Wolfheimer, CST, and team expert Chris Dagdigian from the BioTeam. The team used Amazon AWS cloud resources.

Periodically, the end-user needs large compute capacity in order to simulate and refine potential product changes and improvements. The periodic nature of the computing requirements makes it difficult to justify capital expenditure for complex assets that may end up sitting idle for long periods of time. To date the company has invested in a modest amount of internal computing capacity sufficient to meet base requirements. Additional computing resources would allow the end user to greatly expand the sensitivity of current simulations and may enable new product & design initiatives previously written off as "untestable".

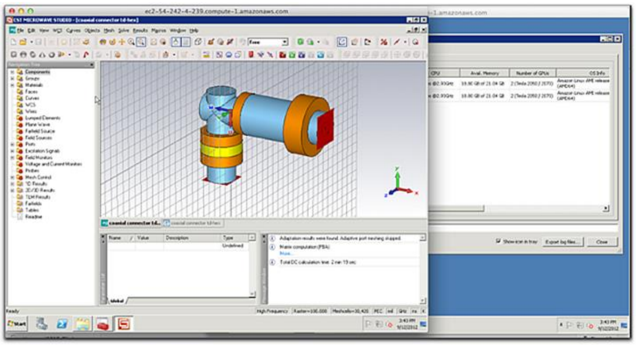

Fig. 2: Team 2 - Front-end showing a simulation model and solvers in action.

The hybrid cloud-bursting architecture allowed local computing resources residing at the end-user site to be utilized along with Amazon cloud-based resources. The project explored scaling limits of the Amazon EC2 instances and scaling runs designed to test computing task distribution via the Message Passing Interface (MPI). The use of MPI will allow us to leverage different EC2 instance type configurations. The team also tested the use of the Amazon EC2 Spot Market in which cloud-based assets can be obtained from an auction-like marketplace offering significant cost savings over traditional on-demand hourly prices.

5.2 Team 58: Wind tunnel flow around bicycle and rider

The team consists of end-user Mio Suzuki from Trek Bicycle, software provider and HPC expert Mihai Pruna from CADNexus, and resource provider Kevin Van Workum from Sabalcore Computing.
The CAPRI to OpenFOAM Connector and the Sabalcore HPC Computing Cloud infrastructure were used to analyze the airflow around bicycle design iterations from Trek Bicycle. The goal was to establish a great synergy among iterative CAD design, CFD analysis and HPC Cloud. Automating iterative design changes in CAD models coupled with CFD significantly enhanced the productivity of engineers and enabled them to make better decisions. Using a cloud-based solution to meet the HPC requirements of computationally intensive applications decreased turnaround time in iterative design scenarios, and reduced the overall cost of the design.

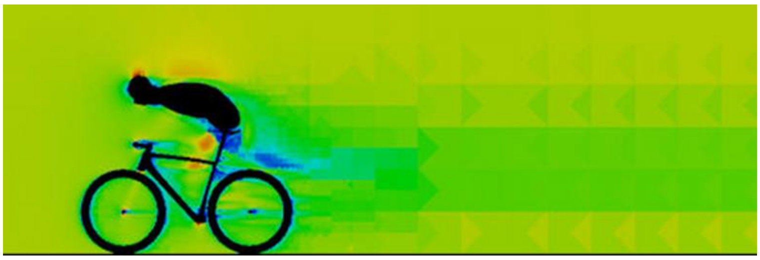

Fig. 3: Team 58 - Velocity color plot, generated with the CADNexus Visualizer Lightweight Postprocessor.

5.3 Team 170: Full car crash simulation with LS-Dyna in the cloud

This team consists of end-users Dan Soca and Stefan Castravete from Caelynx, Uli Göhner from software provider DYNAmore, and Alexander Heine from resource provider CPU24/7.
The end-user Caelynx [15] is an engineering services consulting company with expertise in advanced finite element analysis (FEA) and computational fluid dynamics (CFD) methods to address its customers' product performance requirements. Caelynx consults for leading innovators in the aerospace, automotive, consumer product, high-tech, energy, and life science sectors, transforming their product development processes. Caelynx helps improving product design, providing optimum performance, reducing prototype development costs, and minimizing time to market.
The cloud provider CPU 24/7 [16] operates cloud data centres in Berlin and Hamburg, and specializes in providing scalable HPC power 'as a service' for CAE. CPU 24/7 was founded by German engineers, most of them having their experience and expertise in the field of numerical simulations and the respective computing power. CPU 24/7 solutions always comprise HPC instances

on powerful bare metal servers with pre-installed CAE software, support and storage as well as data traffic.
The software provider DYNAmore [17] is dedicated to support engineers to solve non-linear mechanical problems numerically. DYNAmore's tools to model and solve the problems are the finite element software LS-DYNA as solver and LS-OPT for optimization. With 85 engineers in Europe DYNAmore sells, teaches, supports, and co-develops the software and provides engineering services.

A full Toyota Yaris Sedan car model is proposed to be analyzed via Finite Element Method using the commercial software LS-Dyna. The full car FE Dyna deck (without dummy and airbags) is available over the World Wide Web at no cost from NCAC, [18]. On this model the team will add the frontal airbag and driver dummy and the passengers dummies. The following analyses will be performed:

1. Frontal impact with rigid wall at different car speeds.
2. Side impact with a rigid pillar (pol test). The car will not have lateral airbags
3. Two vehicle impact.

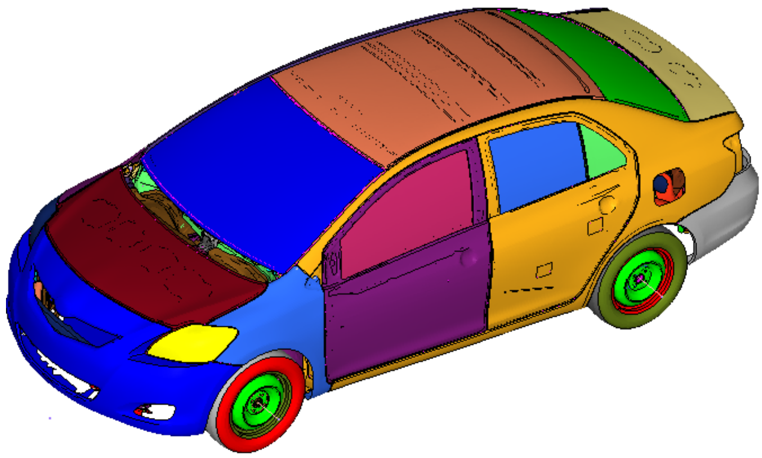

Fig. 4: Toyota Yaris Full Vehicle Model

The analyses will be performed in the CPU24/7 computing cloud where it will be scaled under different No of CPU cores. As output the team is reporting HIC values, energy dissipation, and deceleration at different locations inside vehicle.
At the time of writing this article the experiment team had just been formed and started its first simulations in the cloud. Results and experience with cloud resources will be presented at the European LS-Dyna Conference 2015 in Würzburg, Germany.

6 Summary

The benefits for manufacturers of using high performance technical computing (HPTC) within their design and development processes can be huge; such as better quality products; high return on investment; reducing product failure early in design; and shorten time to market. Potentially, this leads to increased competitiveness and innovation. Engineers and scientists today have mainly three options running their computer simulations: workstations on their desk, HPTC servers within their firewalls, or recently cloud computing. All three options have their benefits and limitations. This review lecture will examine the benefits for manufacturers and scientists of HPTC in the Cloud, and discuss challenges and their main solutions. Finally, the UberCloud CAE Experiment will be presented which provides a platform for engineers and scientists to discover, explore, and understand the end-to-end process of accessing and using Clouds, and to identify and resolve the roadblocks. The lecture will conclude with several real-life engineering case studies in the cloud, among other a full car crash simulation with LS-Dyna in CPU/24/7 cloud resources.

7 Literature

[1] Council of Competitiveness, 'Make', 'Reflect', 'Reveal', 'Compete' studies 2010. http://www.compete.org/
[2] Randy Perry and Al Gillen. Demonstrating business value: Selling to your C-level executives. IDC White Paper, April 2007.

[3] Intersect360/NCMS study on Modeling Simulation at 260 Manufacturers, July 2010. http://www.intersect360.com/industry/downloadsummary.php?id=35

[4] UberCloud Marketplace resource providers. http://www.theubercloud.com/category/high-performance-computing-resource-providers/

[5] Wolfgang Gentzsch. Why Are CAE Software Vendors Moving to HPC Cloud. http://www.hpcwire.com/hpccloud/2012-09-04/why_are_cae_software_vendors_moving_to_the_hpc_cloud_.html

[6] Autodesk Sim 360, http://www.autodesk.com/products/sim-360/overview

[7] CD-adapco Software Licenses. http://www.cd-adapco.com/products/power-licensing

[8] Wolfgang Gentzsch, Burak Yenier. The Uber-Cloud Experiment. In HPCwire, June 28, 2012. http://www.hpcwire.com/hpcwire/2012-06-28/the_uber-cloud_experiment.html

[9] The UberCloud Services Exhibit. http://www.exhibit.hpcexperiment.com/

[10] The UberCloud HPC Experiment: Compendium of Case Studies. Intel, 2013 and 2014. https://www.theubercloud.com/ubercloud-compendium-2013/ https://www.theubercloud.com/ubercloud-compendium-2014/

[11] Katherine Yelick, Susan Coghlan, Brent Draney, Richard Shane Canon. The Magellan Report on Cloud Computing for Science, Dec. 2011. http://science.energy.gov/~/media/ascr/pdf/program-documents/docs/Magellan_Final_Report.pdf

[12] Gartner Predicts Infrastructure Services Will Accelerate Cloud Computing Growth. Forbes Feb. 2013. http://www.forbes.com/sites/louiscolumbus/2013/02/19/gartner-predicts-infrastructure-services-will-accelerate-cloud-computing-growth/

[13] Sam Zakrzewski, Wim Slagter. On Cloud Nine. Digital Manufacturing Report. April 22, 2013. http://www.digitalmanufacturingreport.com/dmr/2013-04-22/on_cloud_nine.html

[14] The UberCloud Receives Top Honors in 2013 and 2014HPCwire Readers' Choice Award. http://www.hpcwire.com/off-the-wire/ubercloud-receives-top-honors-2013-hpcwire-readers-choice-awards/

[15] The end-user Caelynx, www.caelynx.ro

[16] The cloud provider CPU 24/7, http://www.cpu-24-7.com/

[17] The software provider DYNAmore, https://www.dynamore.de/en

[18] Toyota Yaris Sedan car model. NCAC, http://www.ncac.gwu.edu/vml/models.html.