EXCLUSIVE ! Win the Green 500 champion server from ASUS

HPC REVIEW

The reference media in high-performance IT solutions

www.hpcreview.com

HPC | BIG DATA | CLOUD | STORAGE | VISUALIZATION | VIRTUALIZATION | NETWORKS | WORKSTATIONS

LAB REVIEWS

DataCore SANsymphony v10

OCZ Saber 1000 SSD

#4 GLOBAL EDITION

BUILDING A MARKETPLACE FOR HPC IN THE CLOUD

For Engineers, Scientists and Service providers

HOW TO

Understanding Hadoop as a service

ISC 30th ANNIVERSARY SPECIAL

- Why Mercedes-Benz uses HPC
- Financial modeling on HPC Systems
- A quantum leap in computing
- European investments in Exascale Technologies
- Applications leveraging HPC Supercomputer systems

Welcome to this special ISC issue !

HPC is an ever-changing and fast evolving field of computing. A handful of events worldwide help grasp how fast the industry and research advance in their software and hardware definition to be able to address ever-increasing needs. ISC which celebrates ISC's 30th anniversary this year, is a prominent event of the HPC scene. We have therefore assembled a special feature in which focused half a dozen renowned speakers deliver their vision on diffe-rents parts of the HPC market. As an added bonus, they will all attend the show and be available to answer your questions throughout their keynotes.

Another extremely hot topic we cover in this issue is the future of the HPC ecosystem, as a marketplace in the cloud for engineers, scientists and service providers. As Wolfgang Gentzsch, one of its distinguished authors of this highly qualitative review puts it, designing better quality products, shorten time to market, reducing product failure early in design, and increasing return on investment can be achieved by making high-performance computing (HPC) simulations part of the early phase of the product life cycle. Along with several examples of how it can be accomplished.

Our exclusive summer contest in partnership with ASUS, AMD and Intel is still open! All you need to have a unique chance to win one of this year's winning Green500 server nodes is to register on our hpcreview.com website.

Happy reading!

CONTENTS

Small and Medium enterprises enter the HPC limelight

STEVE CONWAY
is IDC Research Vice President, HPC

A decade of close scrutiny has shed much more light on the technical computing needs of small and medium enterprises (SMEs), but they are still shrouded in partial darkness. That's hardly surprising for a diverse global group with millions of members ranging from automotive suppliers and shotgun genomics labs to corner newsstands and strip mall nail salons. Many SMEs presumably will never need HPC in their lifetimes, while many others already benefit from this game-changing technology or could do so. Some SMEs are performing breakthrough work, often by using HPC resources at large, national HPC centers.

WHAT DO WE KNOW?

Let's zero in on this vibrant group that experts credit as the principal growth engine in many economies. The manufacturing sector is replete with SMEs. Within that sector, the National Center for Manufacturing Sciences (NCMS) reports that, in the U.S. alone, there are 300,000 SMEs (defined as companies with 500 or fewer employees). These firms account for 12 million jobs and are indirectly responsible for another 18 million — more than twice the employment of all large manufacturers combined.

Research results submitted for this article by Jon Riley, NCMS vice president of digital manufacturing, indicate that SMEs can be roughly divided into three categories:

- Those that have already adopted [HPC] digital tools (~10 percent)
- Those that may be interested, but are unconvinced or unable to adopt (~75 percent)
- Those that are not interested in adopting (~15 percent)

More later on the important minority who already use HPC resources. Within the majority who don't, an interesting contingent consists of SMEs that perform technical computing on desktop systems but have not moved up to technical servers (HPC systems). Credit for initiating research on this group goes to Suzy Tichenor (now at ORNL) and Bob Graybill (now CEO of Nimbis Services). During their time at the Council on Competitiveness, they collaborated with IDC to investigate these "desktop-only" users. The key findings from those 2007-8 studies are largely applicable today:

- In many cases, the companies' desktop computers were not meeting their advanced requirements, resulting in reduced competitiveness. They responded to this dilemma by scaling down the problems to fit their desktop systems, ignoring the problems, or reverting to much slower, more expensive physical testing and prototyping.

Although the world's 500 biggest supercomputers steal much of the limelight, they represent less than one-half of one percent of the approximately 110,000 HPC systems sold around the world each year.

- 40 percent of the organizations were already planning or actively considering the move to HPC servers.
- The chief barriers to HPC adoption were inadequate understanding, the perceived lack of "strategic fit" software and budget constraints. Some SMEs we interviewed said moving to servers would boost annual ISV software licensing costs from under $10,000 to nearly $50,000.

THE BOTTOM 100,000 BUY A LOT

A decade or so ago, a mega-IT company announced plans to enter the HPC market and target not the top 500 supercomputer sites, but "the bottom 250,000." The plan did not succeed, but it made some sense. Although the world's 500 biggest supercomputers steal much of the limelight, they represent less than one-half of one percent of the approximately 110,000 HPC systems sold around the world each year.

In 2013, HPC systems sold for less than $250,000 each accounted for $5 billion in revenue, about half (49 percent) of the $10.3 billion in worldwide HPC server revenue. At an average price of $40,400, the sub-$250,000 systems consumed 1.7 million processor parts, 52 percent of 3.3 million processors shipped last year, according to IDC research. The bottom half of the HPC market was hammered during the global economic recession and started coming on strong again last year. SMEs are not the only ones that buy sub-$250,000 HPC systems, but that's where most SME buyers of HPC gear reside.

ALTERNATIVES TO PURCHASING

Purchasing systems is not the only option open to SMEs wanting to do HPC, of course. Public cloud computing can be an attractive alternative for SMEs that haven't invested in on-premise HPC resources for their cloud-friendly workloads. We know of SMEs that are doing very well, thank you, by relying entirely on public clouds. Still other SMEs turn to large, national HPC data centers for more powerful computing resources and expertise.

Among the world's premier HPC data centers, none has deeper experience with industrial firms of all sizes than High Performance Computing Center Stuttgart (HLRS). HLRS Director Michael Resch explained why SMEs work with his center, which is situated in the heart of Germany's auto industry:

"We see an increasing demand for HPC from SMEs. Especially in our region, SMEs serve as technology solution providers for larger companies. Increasingly, these large clients require a validation of their technology through simulation. In certain fields, simulation can play a crucial role but is not well-known inside large-scale companies. Very small companies with very special knowledge in modelling and simulation make a living in these small market niches, but they need access to large-scale systems for the computational part of their portfolio."

SME SUCCESS STORIES

Michael Resch provided two live examples:

- RECOM Services started as experts in combustion modelling for large-scale power plants. Initially, their focus was on the local energy provider. About 10 years ago, HLRS set up a contractual framework that allows RECOM Services to use its large-scale resources and support. RECOM Services has 12 employees who are all highly skilled engineers.

Industrial partnership programs have been proliferating in programs at DOE's national laboratories, Europe's PRACE initiative and at leading HPC organizations including NCSA, Teratec, Hartree Centre, HPC Wales, KISTI, the Shanghai Supercomputer Center and others too numerous to mention here.

They could never afford to buy even a small HPC system. Through HLRS, they have access to a petascale system and have expanded their market into the U.S. and Asia continuously in recent years. Their annual costs are about one percent of what they would need to buy their own system — and then such a system would be much smaller than HLRS systems.

• M.A.R.K. 13 is a small company with a focus on media solutions in the field of handling and designing movies. With it 45-person staff, M.A.R.K. 13 is in a business that requires swift reaction time and extremely high quality. When approached to do work for the Australian-German animated movie “Maya the Bee,” M.A.R.K. 13 entered into a collaboration with HLRS to guarantee both high quality and in-time production of 3-D pictures for the 79-minute movie. Although the movie only required about one percent of HLRS resources, it could not have been done without high investment costs by such a small company. So, using HLRS resources not only helped to speed up work whenever necessary by adding more nodes, but also substantially reduced the financial risk for the customer.

The HPC achievements of these two businesses are being echoed around the world by SMEs such as Children's Mercy Hospitals (Kansas City), Intelligent Light, Ping, SmartTruck Systems and hundreds of others.

WHAT'S NEXT?

HLRS is not the only large center that works closely with SMEs and other industrial firms, of course. Germany's two other national HPC centers, the Leibniz Supercomputing Center and the Juelich Research Center, also provide services to industry — as do national HPC centers and labs in North America, EMEA and the Asia Pacific region. Across the globe, government leaders increasingly recognize that HPC can boost industrial and economic competitiveness and SMEs can play a strong part in this transformation.

As a result, industrial partnership programs have been proliferating in programs at DOE's national laboratories, Europe's PRACE initiative and at leading HPC organizations including NCSA, Teratec, Hartree Centre, HPC Wales, KISTI, the Shanghai Supercomputer Center and others too numerous to mention here. A fair amount of this activity is directed at SMEs, which are the most numerous and, in some geographies, the largest enterprises.

SMEs will continue to get more attention, including additional research to illuminate some of the remaining dark corners of this market segment. IDC, for example, recently began what we believe will be the most ambitious effort yet to identify and characterize large numbers of these sites. And organizations, such as the Council on Competitiveness, the National Center for Manufacturing Services, NCSA's Industrial Partners Program, PRACE, HLRS, Teratec and others are hard at work easing the transition to HPC for SME organizations.

ASUS®
IN SEARCH OF INCREDIBLE
intel
RS400-E8-PS2-F
Most Compact 1U Dual Xeon Flexible I/O Server
http://www.asus.com/Commercial_Servers_Workstations/

ISC : 30 years of HPC Community Leadership

The ISC trade show celebrates this year its 30th anniversary. Here are some of the highlights of this year's keynotes and events.

Now in its 30th year, ISC High Performance is the world's oldest and Europe's most important conference and networking event for the HPC community.

Now in its 30th year, ISC High Performance is the world's oldest and Europe's most important conference and networking event for the HPC community. It offers a strong five-day technical program focusing on HPC technological development and its application in scientific fields, as well as its adoption in commercial environments.

Over 400 hand-picked expert speakers and 160 exhibitors, consisting of leading research centers and vendors, will greet attendees at ISC High Performance. A number of events complement the technical program including Tutorials, the TOP500 Announcement, Research Paper Sessions, Birds-of-a-Feather (BoF) Sessions, the Poster Sessions, Exhibitor Forums, and Workshops.

ISC High Performance is open to engineers, IT specialists, system developers, vendors, scientists, researchers, students, journalists, and other members of the HPC global community. The exhibition attracts decision-makers from automotive, finance, defense, aeronautical, gas & oil, banking, pharmaceutical and other industries, as well those providing hardware, software and services for the HPC community. Attendees will learn firsthand about new products and applications, in addition to the latest technological advances in the HPC industry.

PROGRAM FOCUS

ln 2015, we will be highlighting technologies and applications most relevant ta our attendees. Both current and future HPC technologies will be covered, and some of the most challenging applications areas will receive special attention.

- **Future hpc-system design concepts**
- **Highly scalable operating systems**
- **From data to knowledge**
- **Large scale computing in bioengineering & biomedicine**
- **Hpc in the financial sector**
- **Highlights from europe's horizon 2020 ict**
- **Hpc processor & node architecture trends**
- **Hpc and public health sector**
- **Advanced display & visualization technology**
- **Emerging/advanced interconnect architectures**
- **Human brain project-presentations**
- **Top10 exascale research topics**
- **Memory systems for hpc**
- **What's the impact of quantum computing**
- **Hpc in latin america**

ISC EXCHIBITION

The ISC exhibition features the largest collection of HPC vendors, universities, and research organizations annually assembled in Europe. Together, they represent the level of innovation, diversity and creativity that are the hallmarks of the global high performance computing community. Having them all available on the same exhibition floor presents a unique opportunity for users to survey the HPC landscape and for vendors to display their latest and greatest wares.

WHO ATTENDS ISC

ISC attracts top managers, researchers, engineers and sales and marketing staff from universities, national laboratories, automotive, defense, aeronautical, gas & oil, banking and other industries that rely on HPC to help fulfill their orqanizations mission. According to our latest survey, over 97 percent of last year's [20141 attendees are planning to return to the conference as they were highly satisfied with the program and our choice of speakers.

WHO EXHIBITS AT ISC

Exhibitors come from every facet of the HPC community, including system vendors, server manufacturers, chip and memory makers, ISVs, storage vendors, networking solution providers, hardware and software distributors, service providers, data centers, universities, research labs and special projects.

MERCEDES-BENZ CARS HIGH PERFORMANCE COMPUTING FOR A HIGHLY EFFICIENT DEVELOPMENT

Jürgen Kohler Senior Manager, NVH CAE & Vehicle Concepts, Mercedes-Benz Cars Development, Daimler

CAE has become an indispensable part of modern development processes and an integral component alongside design, testing and trials. The challenge of designing ever more complex vehicles efficiently, to a high degree of maturity over an acceptable period of time and without the need for expensive trial phases can be achieved only through the intensive use of simulation tools. The digital prototype has been integral to the work of Mercedes-Benz Cars Development for more than ten years.

Since the launch of the digital prototype, each vehicle has been digitally predesigned and each component specifically optimized by computer in terms of functional requirements such as active and passive safety as well as driving, acoustic and thermal comfort, to name just a few. Such aspects frequently entail trade-offs not only among each other, but also with other aspects such as weight, cost and manufacturability. In the quest for the optimum “brand-appropriate” compromise, the digital prototype has proven to be a major boon. This is evident in the high number of computing jobs that are performed not only to assess a wide range of design variants, but also to make vehicles immune to manufacturing variations. Further methodical and numerical refinement has seen a constant expansion of the application area of simulation tools, a process that remains ongoing.

Various computation models and simulation methods are used depending on the type and complexity of the task. To achieve the greatest possible informative value, state-of-the-art computing tools are needed - in some cases, multiple tools that interact with each other in co-simulation. In addition, computing models are constantly getting bigger because they now also map localized details geometrically and physically.

To ensure that this does not impact the turnaround times in day-to-day development work, rapid numerical algorithms and cost-optimized high-performance computers are deployed, that can keep pace with this trend and are designed for high-capacity computing. This is why Mercedes-Benz Cars Development uses locally placed high-capacity computer clusters. Their efficient and failsafe operation ensures a reliable vehicle development process.

ABOUT THE AUTHOR *Jürgen Kohler studied Aerospace Engineering at the University of Stuttgart. In 1992 he started his career at the Mercedes Benz AG and became Manager of Crash-Simulation in 1997. From 2001 to 2005 he was Senior Manager for CAE Passive Safety, Durability and NVH, from 2006 to 2010 for CAE Passive Safety, Durability and Stiffness CAE and Test. In 2011 he took over responsibility for*

NVH CAE and Vehicle Concepts in the Mercedes-Benz Cars Development. Jürgen Kohler is also a member of the Board of Automotive Simulation Centre Stuttgart (ASCS e. V.) and the Chairman of the PRACE Industrial Advisory Committee.

INTERVIEW
WHY MERCEDES-BENZ USES HPC

How long has Mercedes-Benz been employing high-performance computing for designing vehicles? And what motivated Daimler to turn to CAE and HPC-based simulations?
Kohler: Mercedes-Benz first started with its simulation activities in research and development back in the 1960s. At first, programs for the industrial application of the finite element method were self-developed in-house to support, for example, the elastic-static design of the vehicle structure by computational means. Later on the company switched to the commercial FE program, NASTRAN when it was first offered on the market. There are various versions available today and we use NX NASTRAN. So with the use of more detailed models to consider all relevant design details, the access to HPC became compulsory for us. It was also equally necessary to get the analyses results back in an "acceptable" amount of time. In the past, we had to wait several days; today we are able to obtain results overnight.

ISC: How has the use of simulations changed the development and manufacturing of automobiles?
Kohler: Our so called "Digital Prototype" has been integral to the work of Mercedes-Benz Cars Development for over a decade now. The company realized early on that complex vehicle design can only be achieved through the intensive use of simulation tools in closest interaction with precise and complementary testing. Today CAE has become an indispensable part of modern development processes for us and an integral component alongside design, testing, and trials. This was undoubtedly triggered by the challenge of designing ever-more complex vehicles to a high degree of maturity over an acceptable period of time, and without the need for expensive trial phases.

ISC: How powerful are the compute clusters that Mercedes-Benz uses for vehicle development?
Kohler: The computing models are constantly getting more complex and bigger because they now also map localized details geometrically and physically. To ensure that this does not impact the turnaround times in day-to-day development work, cost-optimized high-performance computers are deployed, which can keep pace with these demands. This is why Mercedes-Benz Cars Development uses high-capacity compute clusters, consisting of a number of multi-CPU machines with sufficient RAM and capacity. For example, when simulating NVH – noise, vibration, and harshness – we can run more than 400 demanding full-vehicle jobs per day, where one run involves calculating 8,000 modes from a 30 MDOF matrix.

AN INSIDER'S VIEW OF FINANCIAL MODELING ON HPC SYSTEMS

By Erik Vynckier, CIO of AllianceBernstein

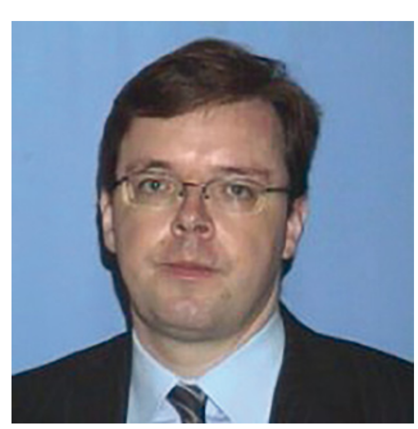

Back in 2008, the global economy came crashing down sending many organizations and individuals into a state of financial ruin. Three major banks in Iceland collapsed, forcing the country into a deep recession. Fingers were pointed at the banking institutions — bank officers and mortgage lenders were blamed for abusing their fiduciary duties toward their customers by putting indulgence and manipulative greed above the stability of the society.

Since then regulations have been strengthened and the financial industry has adopted sophisticated mathematical models and high performance computing systems to study and assess all types of risks in the attempt to minimize them. However, is this practice sufficient enough to assure that our money is safe with the banks and their investment strategies are sound?

Erik Vynckier, CIO of AllianceBernstein, will be speaking about "High Performance Computing in the Financial Industry: Problems, Methods & Solutions," at the upcoming ISC conference in Frankfurt, Germany in July, which focuses on supercomputing technology in research and enterprise settings.

Prior to entering investment banking and later joining the insurance sector, you had also spent a considerable amount of time in the petrochemical industry. What were some of the connections you make there now with respect to complex modeling and the systems required to do it well?

Vynckier: Mathematical modeling and technology are great unifiers of knowledge across sectors. You can change application domains, and by catching up on the knowledge base, quickly re-establish yourself in a new area.

The mathematics and the quantitative modeling expertise, as well as the development & implementation of numerical programs aren't necessarily all that different in the financial sector from the industrial sector.

There is one danger to watch out for however: simple one-for-one porting of models from one context to another, without paying mind to the actual mechanisms at work in the application is very naive and risky. The devil hides in the details!

In fact, scientific models poorly ported to the financial arena led to some grave mistakes and even catastrophic failures. Ill-adapted models that didn't fit the financial markets were crudely implemented, often without questioning or investigating the key assumptions that made them successful in science. In this way, poor modeling contributed to and aggravated the credit crisis.

So how did this all translate into the financial modeling realm?

Vynckier: At previous companies I have implemented a high performance computing platform for real-time, dynamic cross-asset hedging of guaranteed life assurance policies. I also developed a scenario tool for the projection and stress testing of derivative overlays commonly used in liability driven investment strategies. Accurate valuation, accurate hedging, optimal collateral planning and confident product development resulted – all on the same platform. Sharing an integrated platform across different functions and departments limits development costs and increases the speed of developing and bringing to market new financial products.

THE SURVIVAL MACHINE AND SURVIVING THE MACHINES

By Jesus Labarta, full professor on Computer Architecture at the Technical University of Catalonia (UPC) since 1990

Have you ever listened to a neuroscientist talking about the human brain? It is absolutely fascinating. I will always remember a talk in which the brain was referred to as a "survival machine," or SM. The idea is that the SM reacts properly to and hence survives in changing environments. But how is it supposed to know how to react? It makes use of predictive models that are based on data assimilation.

SURVIVAL OF THE MOST ADAPTABLE

Thinking about this capability of adapting to changing environments, I couldn't help but think about HPC – or more specific about the relation between programming models and their interaction with changing hardware environments: In HPC we tend to be extremely confident about our mental models of how a certain computer and program will behave. But this often does not stand the reality check. In fact, there is usually a huge divergence between our mental model and the actual application behavior on a given system. Hitting the "power wall" made us not only go multicore but also broke with what in the past was a clean interface (the ISA) between architecture and software. On our way to an Exascale era, systems will become even more heterogeneous and extremely variable. The only way to cope with such a "hostile" environment is by developing programming models and runtimes capable of adapting dynamically to the underlying hardware environment.

BREAKING WITH OUTDATED APPROACHES TO PROGRAMMING ENVIRONMENTS

In this context, we advocate for decoupling again the programs from the machine. Programs should be written documents to convey algorithms and ideas between humans rather than thinking of them as ways to tell machines what to do or how. I firmly believe that we need to rely on runtime systems, and see them as survival machines. They can be capable of dynamically adapting fairly abstract computational and data access requirements of the application as expressed in the program, to the available resource in the platform at hand. A programming model should provide the programmer with an interface by which they can express their computational and data access requirements in an abstract way. A mechanism to convey additional information may be conveyed so that it is possible for the runtime to efficiently exploit the system. But this should always be in the form of abstract, high level hints and can never be a requirement for a functional description of the ideas. The challenge for a programming model is to provide such capabilities in a way that can be perceived as incremental by current programmers, even if this is a long-term underlying cultural shift in the way programs are written.

THE WAY FORWARD: THE OMPSS PROGRAMMING MODEL

We consider that our work on the OmpSs programming model and NANOS runtime that supports it provides practical evidence that these types of ideas are viable and can be deployed in an incremental way. The fundamental concepts in OmpSs are tasks identified as pragma annotations in front of regular sequential code and directionality clauses that describe the data access by the tasks. From such information, the runtime is able to detect dependences and potential concurrency between tasks as well as perform the data movements that might be required on a specific underlying hardware.

OMPSS WITHIN THE DEEP PROJECT

Within the DEEP project we have extended the original OmpSs environment with MPI offloading capabilities. In it, one or several processes in an MPI application can collectively spawn other MPI computations that can be dynamically mapped to different parts of a heterogeneous cluster based on appropriateness of the cores or the interconnect, for example.

The approach has shown extremely useful not only to offload computations as initially required by the DEEP project objectives, but also to do multilevel nested spawns or reverse offloading back to the original nodes with minimal and concise annotations. Offloading I/O tasks has proven to be an interesting functionality where some of the computational nodes do not have such capability. In many cases, this I/O offloading resulted in a transparent overlap of computation and I/O without any explicit programming effort. It is interesting to see how many unexpected execution orders or overlaps do show up in OmpSs codes that are beneficial for performance. Previous explicit parallelization efforts never dared to explore these opportunities because they implied a complex code reworking or simply the programmer never thought about them.

I understand that people might be hesitant to hand over the control they think they have to an autonomous run-time. But ongoing evidence shows huge potentials. Our vision is that a clean program has to focus only on specifying computational and data access requirements. This specification will then be mapped to the hardware by a dynamic, adaptive and otherwise intelligent runtime, as this is the only way to survive the diversity, heterogeneity and variability of machines of the future.

Come meet us at the joint European Exascale Projects booth #634. We are looking forward to lively discussions with you on this and further topics.

ISC EVENT SCHEDULE *You are interested in learning more about research on Exascale in Europe? Then register for our workshop: "Is Europe Ready For Exascale? A Summary of Four Years of European Exascale Research" held on Thursday, July 16 2015. For more information, please visit http://bit.ly/1Clnlg4*

ABOUT THE AUTHOR *Jesus Labarta is full professor on Computer Architecture at the Technical University of Catalonia (UPC) since 1990. Since 1981 he has been lecturing on computer architecture, operating systems, computer networks and performance evaluation. His research interest has been centred on parallel computing, covering areas from multiprocessor architecture, memory hierarchy, programming models, parallelizing compilers, operating systems, parallelization of numerical kernels, performance analysis and prediction tools. Since 2005 he is responsible of the Computer Science Research Department within the Barcelona Supercomputing Center (BSC). He has been involved in research cooperation with many leading companies on HPC related topics. His major directions of current work relate to performance analysis tools, programming models and resource management. His team distributes the Open Source BSC tools (Paraver and Dimemas) and performs research on increasing the intelligence embedded in the performance analysis tools. He is involved in the development of the OmpSs programming model and its different implementations for SMP, GPUs and cluster platforms. He has been involved in Exascale activities such as IESP and EESI where he has been responsible of the Runtime and Programming model sections of the respective Roadmaps. He leads the programming models and resource management activities in the HPC subproject of the Human Brain Project.*

TO PACK OR NOT TO PACK ?

By Dr. Christian Simmendinger, senior HPC consultant for T-Systems - Solutions for Research

ISC 2015 is late this year. At this time of the year families usually are already packing for their summer holidays. Though this blog post is not related to holidays, but to Supercomputing, there is a direct relation between the two topics: packing. When my kids were younger, packing stuff for holiday was always a problem: While we managed to bring along the essentials (cuddly toy and tooth brush), we inevitably forgot items the kids considered indispensable (like cuddly toy number 2,3,4,5 and 6). Assembling all the required stuff took ages. Things were never where they should have been - rather they were spread across the entire flat. Over the years, packing time has constantly improved. The kids now not only know where the stuff is (still spread across the flat), but they now pack it themselves and then bundle it into their trunk in our old campervan.

HOW DOES ALL THAT TRANSLATE TO SUPERCOMPUTING?

The answer is that in programming models – and especially in hybrid programming models – we frequently have a similar situation. Packing data for communication can take a long time, if you do it single threaded. If every thread assembles the data it owns (cuddly toys and tooth brush) and copies it into an assigned slot in the linear communication buffer, you are ready for departure much earlier.

As an example, let's assume, we run an HPC application on an Intel Xeon Phi Processor with 240 threads. Let's further assume that we have 3 communication neighbors (3 other Phi Processors). For a small problem size we then might have (per message buffer) up to 80 threads participating in the assembly of a single linear message buffer.

Similar to going on holidays with kids, the fastest departure here is neither achieved with the kids all traveling by themselves (splitting the communication buffer in 80 different messages) nor with a single thread (you) packing all the stuff. Instead all threads ideally pack their stuff into that linear buffer. Whenever the last contribution has arrived - you are ready to go.

There are some consequences to this simple scheme. One of these consequences is related to MPI data types. MPI data types today are fusing the packing and the send. In a highly multithreaded environment that concept becomes questionable. This is especially true, if access to (cache-coherently shared) memory becomes expensive and the last contributing thread needs to perform a single-threaded assembly (finding tooth brush, plus cuddly toys number 1-6 for 80 kids) of the linear communication buffer .

In the context of the European EXA2CT FP7 project we have implemented the various approaches to packing for the use case of a ghost cell exchange for unstructured meshes on Xeon Phi.

The difference in elapsed runtime between optimal multi-threaded packing and single-threaded packing becomes substantial very fast: For a small mesh, with just 4 Xeon Phi cards we observe that an optimal multi-threaded packing improves parallel efficiency already by a factor of up to two in terms of total application runtime.

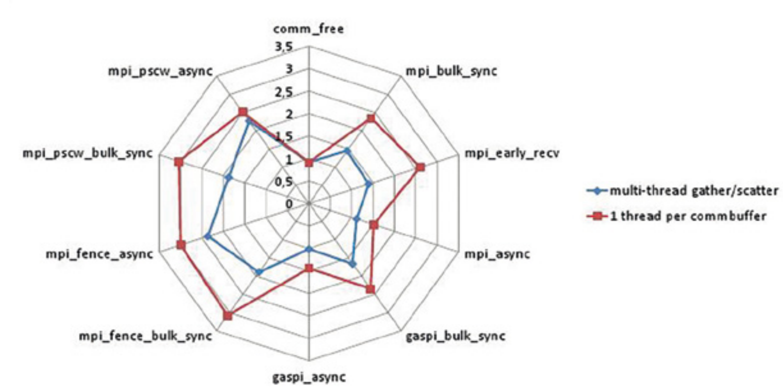

Multithreaded gather/scatter (pack/unpack) vs single threaded pack/unpack. 4x240 threads on Xeon Phi. Total application runtime, normalized to a (theoretical) communication free model.

EVENT SCHEDULE *If you want to learn more about optimal multithreaded packing, the EXA2CT FP7 project, or other European Exascale projects, come and meet us at the joint European Exascale Projects booth #634. We are looking forward to lively discussions with you.*

ABOUT THE AUTHOR *Dr. Christian Simmendinger is a senior HPC consultant for T-Systems - Solutions for Research. He currently works for the European EXA2CT FP7 project and also was the project leader of the (BMBF funded) GASPI project. Both projects aim at establishing a novel PGAS API (GASPI) for next-generation Exascale supercomputers. In 2013 – together with Rui Machado and Carsten Lojewski - he was awarded the Joseph-von-Fraunhofer Preis (Award) for his contributions to GASPI/GPI.*

EUROPEAN INVESTMENT IN EXASCALE TECHNOLOGIES

By Dr Lorna Smith, group manager at EPCC at the University of Edinburgh

It hardly seems any time since the European Commission announced a multi-million euro investment in the development of European exascale technologies. At that time the HPC community was only just beginning to consider and understand the challenges of developing these technologies, and the IESP published its software roadmap. Against this backdrop the EC announced that three flagship projects had been funded: CRESTA, DEEP and Mont-Blanc.

Fast-forward three years and the CRESTA project is now complete, having successfully delivered all of its objectives and outputs. As one of the first real projects funded to prepare and develop software, tools and applications for future exascale technologies, the experiences of the project are particularly important for the community. But as we continue to ensure CRESTA's software is fully exploited, how has the European exascale landscape changed and how did CRESTA contribute?

Of the three flagship projects, CRESTA was the only purely software focused project, with DEEP and Mont-Blanc developing prototype hardware. The project was also unusual due to the scale of investment in a software project – 11.3M euros in total. This allowed the 13 project partners from around Europe to explore and deliver a range of innovative and disruptive products. If we have learnt anything from exascale software development in CRESTA, it is that it is rarely an incremental development process. Significant investment is required to allow developers to explore truly disruptive solutions. Many developers struggle to find direct funding for software development, with funding often directed at the science enabled by the software. Hence the EC's investment in software has been critical and is perhaps a factor behind Europe's recognition as an expert in HPC software.

Another unusual aspect to CRESTA has been the use of a co-design process. While co-design is widely used in the development of embedded systems, its use on HPC projects is less common. However the significant change anticipated in exascale hardware argues strongly that the traditional disjoint between application scientists and hardware and software developers is no longer viable. As one of the first projects to use a co-design process in practice, CRESTA has acted as an important demonstrator. From implementing an appropriate management methodology, through motivating and engaging partners, to measuring and assessing the impact of co-design on outcomes, the process has touched on all aspects of the project. The end result has clearly been a success, and some outcomes simply could not have been achieved without this approach. The use of co-design in HPC development is growing and a number of the European Exascale Projects are also now pursuing this approach. As many of you involved in HPC software pro-

Fast-forward three years and the CRESTA project is now complete, having successfully delivered all of its objectives and outputs. As one of the first real projects funded to prepare and develop software, tools and applications for future exascale technologies, the experiences of the project are particularly important for the community.

jects will be aware, the EC is looking to ensure projects such as these deliver economic growth, European research leadership and tackle major societal challenges, and rightly so. However articulating and quantifying the impact of software development projects can be challenging. Across the community many developers are looking for ways to achieve this and CRESTA was no exception. Ensuring the project outputs were available and utilised beyond the end of the project became key – in addition to the obvious open source, on-line repository, CRESTA established a series of pilot projects which were successful in ensuring the exploitation of the software. The applications all had important societal outputs, and a series of case studies showcased these.

So as ISC 2015 approaches, how has the European exascale landscape changed? Since the original three projects were funded in 2011, the EC has invested further funding in exascale technologies, resulting in three additional projects: EPiGRAM, NUMEXAS and EXA2CT. These six projects have collaborated closely, holding joint workshops, Birds-of-a-feather and exhibition booths. More recently, the EC has announced 21 new projects in this area. Of these, EPiGRAM, NextGenIO, INTERTWinE and ExaFLOW all incorporate outcomes from CRESTA. This investment is allowing European exascale technology investment to go from strength to strength.

ISC EVENT SCHEDULE *The European Exascale Projects are holding a full day Workshop at ISC 2015 on Thursday July 16 titled "Is Europe Ready For Exascale? A Summary of Four Years of European Exascale Research". Please come along and learn more - or drop by the EPCC booth (no 1203) or the European Exascale Projects Booth (no 634).*

ABOUT THE AUTHOR *Dr Lorna Smith is a group manager at EPCC at the University of Edinburgh. She is an experienced project manager, having managed both European, international and UK projects. She was involved in the establishment and management of the Exascale Technology Centre at EPCC, was the project manager for the FP7 funded Exascale Project CRESTA and is the local coordinator for the Exascale EPiGRAM project. She is the deputy directory of the ARCHER CSE service, providing training, user support and programme management for the UK's National HPC service. She has a wide range of interests in supercomputing including the investigation of new languages and the challenges of utilising future Exascale platforms.*

A QUANTUM LEAP IN COMPUTING, MAYBE

By Thomas Sterling, executive director of the Center for Research in Extreme Scale Technologies, and professor of informatics and computing at Indiana University

Quantum Computing has been a concept since the 1980's that has remained outside the domain of real-world high performance computing. Through the era of Moore's Law and exponential progress in feature size, clock rates, and resulting performance, the need for alternative paradigms and technologies has attracted little attention or interest. But there has remained a curiosity among a limited community, primarily but not only in academia, that has driven slow but persistent advances in associated ideas including theory, technologies, algorithms, and most recently commercially available proof-of-concept systems. The driving motivation is that at least in principle some problems could be solved by a quantum computer that could never in practical terms be performed by a conventional supercomputer in a human lifetime. The poster child for this is factoring of products of prime numbers for cryptology. But there are many obstacles between this extraordinary ambition and the reality of today's technological capabilities; obstacles that have persistently deferred interest in what has been perceived as unlikely speculation. Nonetheless, progress through incremental steps by the international research community has reversed the conventional viewpoint. Quantum community can no longer be ignored. Although still at its most inchoate phase, early results are beginning to suggest that quantum computers may one day be a real contribution to the domain of high performance computing; maybe.

Of course those of us with a practical Newtonian upbringing; enhanced with a bit of Einstein equations (special relativity can be derived with high school algebra and trigonometry); we all quote that E = mc2, understand the many von Neumann derivatives of Vector, SIMD, and MPP (including my personal favorite: Beowulf Linux clusters) comprising the last four decades of supercomputer history. Quantum computing is something completely different. Based on quantum mechanics, quantum computing is as weird compared to conventional practices as quantum mechanics is from Newtonian physics. Trying to explain quantum computing is like teaching computer science at Hogwarts. But the remarkable properties of quantum computing are due to the equally counter-intuitive mechanisms of quantum mechanics. Among these are superposition and entanglement; just terms to most of us but central to the unique behaviors possible.

Superposition is the ability to store many different states simultaneously in a single set of memory units. Conventionally, the memory units are bits. Each bit stores either a 1 or 0 (if you didn't know that, you probably shouldn't be reading this article). For a handful of bits, the numeric range increases by 2n for n bits. But it is still a single value. Superposition allows multiple values to be stored in a single number of quantum computer bits know as "qubits". Fancy diagrams of spheres with vectors from the center pointing somewhere out to the surface are used to illustrate the idea that multiple values are stored at the same time in a set of qubits. The number of such values is exponential with the number of qubits. Yes, you should already be confused. But the key concept is that quantum computers should not only store this multiplicity of values but be able to process them simultaneously as well. Thus a quantum computer, at least in principle, is an extraordinary parallel computer, but not by doing the operations separately; rather doing them in the same logic at the same time.

There is an important, actually dominant, caveat. You can't "look" at value(s) in a qubit without the many-value value collapsing to a singular value. Yes, the act of looking actually affects the state of the quantum machine. This is where the property of "entanglement" comes in. If superposition appears very strange, this is nothing compared to entanglement in quantum mechanics. If two particles are directly associated (like near each other), they can be coupled such that if they are then separated they are still related. In fact if you look at one of the two particles, you know what the state of the other distant particle is. And if that is not weird enough, this is done instantaneously; not just very fast but truly instant. Yes, it appears to defy the speed of light. Albert Einstein rejected the idea back in the 1930s referring to it as "spooky action at a distance." His reaction was entirely reasonable and he was among the earliest to truly to understand this implication. But the father of "relativity" was wrong. All tests performed today to prove or disprove this theory has demonstrated its reality. One of the consequences of entanglement is that, again in principle, viewing one entangled particle allows the other pair-wise par-

Much of the engineering going into the realization of a conceptual quantum computer is dedicated to resolving or mitigating the problem of noise. These challenges have been understood for the last three decades and yet quantum computing has been ignored by the HPC community.

ticle to be read without disturbing the multi quantum state of its partnering particle.

Another problem imposes a barrier between theory and realistic implementation; quantum noise. Keeping qubits stable is extremely hard, perhaps impossible for an indefinite period. To understate the case, operation of quantum devices is statistical. Getting the right answer is not guaranteed and there is a probability associated with each system's computation dependent on the sensitivity to sources of noise. One can imagine making many runs of the same problem and generating a probability distribution function of answers. Much of the engineering going into the realization of a conceptual quantum computer is dedicated to resolving or at least mitigating the problem of noise. These challenges have been largely understood for the last three decades and yet quantum computing has been ignored by the HPC community, and for good reason – there was essentially no reason to believe they would ever become reality, at least in the lifetime of current practitioners. That may have changed.

Like conventional digital computing, many different technologies are considered as possible enablers for future quantum computers. Experiments by a number of research institutions have explored this space and are responsible for slowly emerging confidence in the possible likelihood of alternative technologies. In 2008, a new company, D-wave, was established in Canada to create the world's first commercial quantum computers. Based on earlier technology prototypes, D-wave announced its first commercial offering in 2011, D-Wave One, incorporating 128 qubits and in 2013 a 512 qubit system, D-Wave Two. These systems employ a technology based on superconducting integrated circuits with pair-wise flux qubits. Although superconducting circuits require temperatures of about 4 Kelvins to work, to minimize the quantum noise, the D-Wave machines maintain a thermal environment more on the order of 20 milli-Kelvins. Even with this success, there is a debate as to whether or not it is a "real" quantum computer (whatever that means). Rather, it is an adiabatic annealing machine emphasizing optimization problems. The question of actual quantum entanglement is still unresolved and only recently has there been indirect evidence that such phenomena is actually happening. Nonetheless, there is real interest now among such giants as Lockheed, Google, and others even in these early systems as proof of concept. The engineering achievements so far are very impressive.

Demonstration of a number of algorithms has already been achieved on these and other laboratory examples of quantum processing. The field is wide open and the future while uncertain (the centerpiece of the quantum computing technology itself) is very promising. HPC can no longer ignore, nor should it want to, then paradigm shift that may provide part of the solution to addressing the end of Moore's Law.

ABOUT THE AUTHOR *Dr. Thomas Sterling is executive director of the Center for Research in Extreme Scale Technologies, and professor of informatics and computing at Indiana Uni-*

versity. He also is an adjunct professor at the Louisiana State University Center for Computation and Technology, and CSRI Fellow at Sandia National Laboratories. Sterling's current research focuses on the ParalleX advanced execution model for extreme scale computing, and devising a new model of computation to guide the development of future generation Exascale computing systems. This research is conducted through several projects sponsored separately by DOE, NSF, DARPA, Army Core of Engineers, and NASA. Since receiving his PhD from MIT, Sterling has engaged in applied research in fields associated with parallel computing system structures, semantics, and operation in industry, government labs, and higher education.

APPLICATIONS LEVERAGING SUPERCOMPUTER SYSTEMS

By Dr Yutong Lu, Director, System Software Laboratory, School of Computer Science, National University of Defense Technology

Scientific discovery and life science research need more and more supercomputing and data processing power, but the gap between the capability of the supercomputers and the requirements of the most demanding applications is becoming wider. The peak performance of supercomputers, such as Tianhe-2, is increased dramatically by using the latest heterogeneous architecture and high-speed interconnects, however the efficiency of real applications is still limited by conventional models, algorithms, and IO processing. Therefore, we need to improve application scalability and efficiency to adapt to the capability of emerging supercomputers, exploring the relationships between the CPU, accelerator, interconnect and the storage subsystem. This talk analyzes the key issues of application scalability for advanced HPC systems in the post-petascale era. The capabilities of CPUs, accelerators, interconnect and I/O storage systems will be discussed and how they fit together to create a whole HPC-System. Included in this discussion will be how large-scale applications running on Tianhe-2 are designed and supported by balancing the system's computation, communication, data management, and fault tolerance. Some large applications from different domains will be presented, including CFD, genomics, and cosmology, as well as big data types of workloads. A co-design approach will be highlighted for the research and development activities, to deliver a whole system for scalable computing, and to support the large-scale domain applications efficiently.

ABOUT THE AUTHOR *Prof. Dr. Yutong Lu is the Director of the System Software Laboratory, School of Computer Science, National University of Defense Technology (NUDT), Changsha, China. She is also a professor in the State Key Laboratory of High Performance Computing, China. She got her B.S, M.S, and PhD degrees from the NUDT. Her extensive research and development experience has spanned several generations of domestic supercomputers in China. During this period, Prof. Lu was the Director Designer for the Tianhe-1A and Tianhe-2 systems. Her continuing research interests include parallel operating systems (OS), high-speed communication, global file system, and advanced programming environment.*

AMD FIREPRO

AMD FirePro™ S9150 Server GPU

The GPU of choice for high-performance computing

The most compute-intensive workloads in data analytics or scientific computing are no challenge for the AMD FirePro™ S9150 server GPU. With support for OpenCL™ 2.0, 16 GB GDDR5 memory, and up to 2.53 TFLOPS of peak double-precision and up to 10.8 GFLOPS-per-watt peak double-precision performance, the choice is clear.

The AMD FirePro™ S9150 features:

- Up to 5.07 TFLOPS peak single-precision floating point performance
- Up to 2.53 TFLOPS peak double-precision floating point performance
- 16 GB of GDDR5 memory
- Up to 320 GB/s memory bandwidth
- AMD Graphics Core Next (GCN) Architecture
- Full Rate Double Precision
- AMD STREAM Technology
- AMD PowerTune Technology
- OpenCL™ 2.0 support

THE GREEN 500

#1 on the November, 2014 Green500 list[1]

Please visit ***www.fireprographics.com/s-series*** *to find out where you can get the AMD FirePro S9150.*

1. AMD FirePro™ S9150 server GPU powers the #1 supercomputer on the November, 2014 Green500 list. For more details, please visit http://www.green500.org/news/green500-list-november-2014

2. AMD FirePro™ S9150 max power is 235W and delivers up to 2.53 TFLOPS peak double and up to 5.07 peak single precision floating point performance. Nvidia's highest performing single-GPU server card in the market as of March 2015 is the Tesla K40, max power of 235W, with up to 1.43 TFLOPS peak double and up to 4.29 peak single-precision compute performance. FP-97

BOOKS

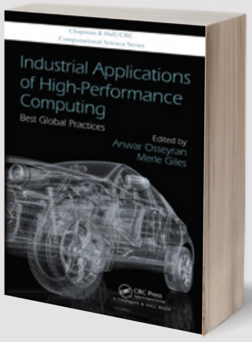

Industrial Applications of High-Performance Computing: Best Global Practices

Industrial Applications of High-Performance Computing: Best Global Practices offers a global overview of high-performance computing (HPC) for industrial applications, along with a discussion of software challenges, business models, access models (e.g., cloud computing), public-private partnerships, simulation and modeling, visualization, big data analysis, and governmental and industrial influence. Featuring the contributions of leading experts from 11 different countries, this authoritative book provides a brief history of the development of the supercomputer; describes the supercomputing environments of various government entities in terms of policy and service models. Includes a case study section that addresses more subtle and technical aspects of industrial supercomputing; shows how access to supercomputing matters, and how supercomputing can be used to solve large-scale and complex science and engineering problems; emphasizes the need for collaboration between companies, political organizations, government agencies, and entire nations. **Anwar Osseyran, Merle Giles, Chapman-Hall / CRC Press, 410 pages (£54.39)**

The End of Error

Written by one of the foremost experts in high-performance computing and the inventor of Gustafson's Law, The End of Error: Unum Computing explains a new approach to computer arithmetic: the universal number (unum). The unum encompasses all IEEE floating-point formats as well as fixed-point and exact integer arithmetic. This new number type obtains more accurate answers than floating-point arithmetic yet uses fewer bits in many cases, saving memory, bandwidth, energy, and power. Richly illustrated in color, this groundbreaking book represents a fundamental change in how to perform calculations automatically. It illustrates how this novel approach can solve problems that have vexed engineers and scientists for decades, including problems that have been historically limited to serial processing. The book is accessible to anyone who uses computers for technical calculations, with much of the book only requiring high school math. John L. Gustafson, Chapman-Hall / CRC Press, 416 pages (£33.14)

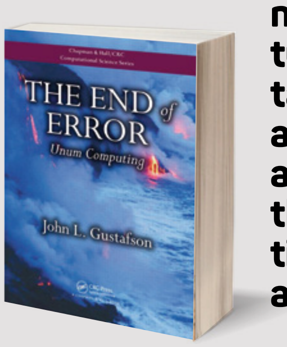

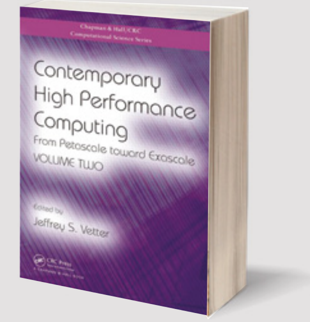

Contemporary High Performance Computing: From Petascale toward Exascale, Volume Two

Jeffrey S. Vetter, Chapman-Hall / CRC Press, 255 pages (£54.39)

HPC is used to solve a number of complex questions in computational and data-intensive sciences. These questions include the simulation and modeling of physical phenomena, such as climate change, energy production, drug design, global security, and materials design; the analysis of large data sets such as those in genome sequencing, astronomical observation, and cybersecurity; and the intricate design of engineered products, such as airplanes and automobiles. This book documents international HPC ecosystems, including the sponsors and sites that host them. Each chapter presents highlights of applications, workloads, and benchmarks; describes hardware architectures, system software, and programming systems ; explores storage, visualization, and analytics. Examines the data center/facility as well as system statistics; featuring pictures of buildings and systems in production, floorplans, and many block diagrams and charts to illustrate system design and performance, Contemporary High Performance Computing.

MOOCS

Big Data in Education

Education is increasingly occurring online or in educational software, resulting in an explosion of data that can be used to improve educational effectiveness and support basic research on learning. In this course, you will learn how and when to use key methods for educational data mining and learning analytics on this data. You will learn about the methods being developed by researchers in the educational data mining, learning analytics, learning at scale, student modeling, and artificial intelligence in education communities, as well as standard data mining methods frequently applied to educational data. You will learn how to apply these methods, and when to apply them, as well as their strengths and weaknesses for different applications. The course will discuss how to use each method to answer education research questions and to drive intervention and improvement in educational software and systems. Methods will be covered both at a theoretical level, and in terms of how to apply and execute them using software tools like RapidMiner. We will also discuss validity and generalizability; towards establishing how trustworthy and applicable the results of an analysis are. Some knowledge of either statistics, data mining, mathematical modeling, or algorithms is recommended. Experience with programming is not required.This course is comparable in difficulty to the first course in the Masters in Learning Analytics at Teachers College Columbia University, though it does not go into the same depth as that course. **Free.**

Starts July 1, 2015
Level: Advanced
Length: 8 weeks
Effort: 6-12 hours/week
Subject: Statistics & Data Analysis
Institution: ColumbiaX
Languages: English
Video Transcripts: English
https://www.edx.org/course/big-data-education-columbiax-bde1x#!

Linux System Administration Essentials

Linux professionals are in high demand with >50K open jobs today. 97% of hiring managers want Linux talent because Linux is the #1 Internet, Cloud, mobile and CE OS. Show employers you have the skills and the knowledge to be a top sysadmin with the Linux Foundation Certified System Administrator (LFCS). This Linux course is for those just starting their career in IT as well as professionals with experience on other operating systems who want to add Linux to their portfolio. This sysadmin course will help you develop the fundamental skills that employers look for in entry level Linux System Administrators by focusing heavily on hands-on practice via nearly 100 lab exercises. Entry level certification is the best way for candidates who are new to Linux to show their capability. After finishing the course on edX, learners will be able to register for LFCS Exam for free. You will then have 12 months to complete the exam from registration date. You can also take advantage of a free exam retake, if needed.

Level: Intermediate
Subject: Computer Science
Institution: LinuxFoundationX
Languages: English
Video Transcripts: English
Type: Professional Education, Self-paced
Price: $499
https://www.edx.org/course/linux-system-administration-essentials-linuxfoundationx-lfs201x#!

THE HPC OBSERVATORY

$44 Billion

This is the projection of the worldwide turnover of HPC in 2020. Market Research Media research firm expects the area of supercomputing will grow an average of 8.3% annually to reach $ 44 billion in 2020 . This sector will also generate 220 billion dollars in cumulative sales over the period 2015-2020. *Source :* *http://www.marketresearchmedia.com/*

THE TOP 3 OF THETOP 500

1 Tianhe-2
National Supercomputing Center in Canton: 33863/54902 TFlops Manufacturer: NUDT**, architecture** Xeon E5-2692 Xeon Phi 31S1P +, TH Express-2

2 Titan
OOak Ridge National Laboratory, USA: 17590/27113 TFlops Manufacturer: Cray XK7**, architecture** Opteron 6274 + Nvidia Tesla K20X Cray Gemini Interconnect

3 Sequoia
Lawrence Livermore National Laboratory, USA: 17173/20133 TFlops Manufacturer: IBM Blue Gene / Q **architecture** PowerPC A2

The TOP500 ranks every six months the 500 most powerful supercomputers in the world. The two selected values, RMAX and RPEAK represent the computing power and maximum theoretical Linpack.

THE TOP 3 GREEN 500

1 5271,81 MFlops/W
GSI Heimholtz Center (Germany)
57.15 kilowatts consumption

2 4945,63 MFlops/W
High Energy Accelerator KEK (Japan)
37.83 kilowatts consumption

3 447,58 MFlops/W
GSIC Center, Tokyo Institute of Technology (Japan)
35.39 kilowatts consumption

Green 500 rankis the most energy efficient supercomputers in the world. Energy efficiency is assessed by measuring performance per watt. The unit is here MFLOP / Watt.

BUILDING A MARKETPLACE FOR HPC IN THE CLOUD

For Engineers, Scientists and Service providers

We Always Want Better, Faster, Cheaper

Designing better quality products, shorten time to market, reducing product failure early in design, and increasing return on investment can be achieved by making high-performance computing (HPC) simulations part of the early phase of the product life cycle. Today engineers and scientists have access to three major computing tools for their simulations: workstations on their desk, in-house computing servers, and HPC cloud resources. All three options come with benefits and challenges. We will examine these benefits and challenges for manufacturers and scientists using HPC in the Cloud which complement in-house workstations and servers. We will also present UberCloud's approach to take HPC to the masses – a highly sophisticated claim which sounded completely unrealistic when we started with UberCloud three years ago. But look and see how close we got in the meantime; today, it's indeed better, faster, cheaper, with cloud computing.

YOUR CHOICE: WORKSTATION, SERVER, CLOUDS

Why then are many engineers running simulations still just on their workstations, although - according to the US Council of Competitiveness, [1] – 57% are regularly dissatisfied with performance limitations of their workstations? The main reason is that the alternatives – in-house HPC servers and public HPC clouds - are still coming with several challenges, at least until recently.

57% of engineers are dissatisfied with the performance of their workstations. They need more computing resources.

Alternative one, buying a powerful HPC server, comes with high "Total Cost of Ownership" (TCO) demonstrated by IDC in 2007 [2]: in addition to server cost, the expenses for staffing, training, software, downtime, and maintenance easily sum up over three years to the ten-fold of the original server cost. And, in addition, there are often long and painful internal procurement and approval processes. Finally, for many, "Return on Investment" (ROI) is not easy to describe, although it is expected to be huge according to a recent IDC study on ROI in HPC, [3].

HPC CLOUD – WHAT'S IN IT FOR ME

The second alternative for SMBs to experience the benefits from HPC is recently offered by cloud computing. For engineers and scientists: HPC in the Cloud enables engineers and scientists to continue using their workstation for daily design and development, and to submit larger, complex, and time-consuming jobs or many jobs in parallel to the cloud. Major benefits of the HPC Cloud solution are on-demand access to 'infinite' resources especially for SMBs which cannot afford to buy larger more expensive HPC servers, pay per use or short-term subscription, reduced capital expenditure, greater business agility, and dynamically scaling resources up and down as needed.

Cloud users: Submit large, complex, time-consuming jobs to the cloud.

For independent software vendors (ISVs): For ISVs the Cloud is a new delivery platform, and it is an entirely new route to market, and possibly a channel to new customers, that will provide its software more flexibility while also requiring a change in strategies related to pricing, reseller relationships, licensing and revenue streams. The Cloud enables ISVs to instantly respond to their customers' needs for more and short-term software licenses when peak loads occur on the customers systems and business needs to react quickly. New software container technology enables ISVs to install, evaluate, and test their software once and then it runs anywhere, without complex re-installation, on any resource. As already demonstrated in the enterprise IT world the economies of scale that come with the cloud combined with the innovation and agility of fully SaaS (Software as a Service) ISVs has led to a world of primarily SaaS-based applications which now influences the technical software community too.

For ISVs the Cloud is a new delivery platform, a new route to market, and possibly a channel to new customers.

For cloud resource providers: That's a tough one. In case you are not Amazon, Google, or Microsoft you might face the danger of getting squeezed between the big guys and their price battles. The only way to escape is to add value on top of your cloud infrastructure resources. A common way today is to closely work with ISVs to bundling software and hardware and make it accessible 'all in one', together with an excellent customer service and open information policy. However, a drawback for the customer can then occur because it is often a challenge to find out pricing for the bundle, a situation which might scare away many customers.

Cloud providers should work with ISVs to bundling software & hardware and make it accessible 'all in one'.

For consultants and freelancers: This group is very lucky! In a sudden consultants have a completely new business opportunity by adding software and computing resources in the cloud to their existing consulting services without paying upfront, and at the end of the project they just charge it to their customer's account.

Cloud opportunity for consultants: adding software and cloud services to their business.

The HPC Cloud Market

Global spending on Infrastructure as a Service (IaaS) cloud computing is expected to reach almost US$16.5 billion in 2015, an increase of 32.8 percent from 2014, with a compound annual growth rate (CAGR) from 2014 to 2019 forecast at 29.1 percent, according to Gartner's latest forecast [4]. Other forecasts [5] look at Software as a Service (SaaS) and find that the SaaS market is expected to grow steadily into the future, estimating that by 2015, SaaS revenue will have climbed as high as 27.94 billion U.S. dollars up from 13.46 billion in 2011.

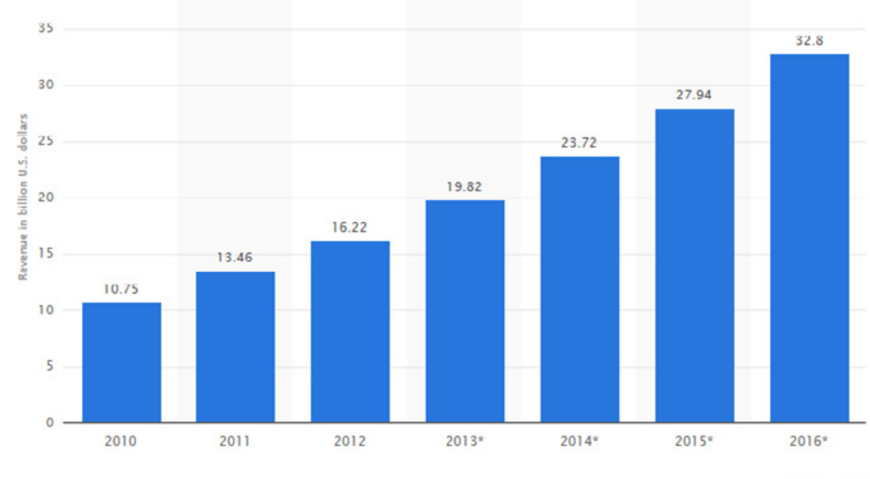

The Software as a Service cloud market 2010 to 2016 [5]

However, the HPC in the Cloud market is still far behind the enterprise SaaS market, for all the good reasons we are mentioning in this article. The only market data currently available for HPC Cloud are from an IDC research involving 905 HPC sites of various sizes in government, academia and industry, [6], where 23.5% of these sites used cloud computing of some type in 2013, up substantially from 13.8% in 2011. Of the HPC sites that use cloud computing, half use private clouds and the other half uses public clouds – our guess: in an early stage. Early adopters include government, manufacturers, bio-life sciences, oil and gas, digital content creation, financial services and other high-end analytics, online consumer services, and health care firms.

One important market segment and potential cloud opportunity is not considered in the IDC study, the digital manufacturing market of engineers and scientists using mainly workstations. According to a study of the Council of Competitiveness [1], about 97% of the companies perform virtual prototyping or large-scale data modeling just on their desktop computers, making them more likely prospects for HPC beyond their desktops. 57% of these companies said they have problems that they can't solve with their existing desktop computers, and thus they have a real need for more computing power. According to Jon Peddie Research [7], workstations currently are a $7 Bio revenue market, worldwide, with about 4 million units shipped last year. That's about 20 million users using mainly workstations and PCs for their daily design and development simulations, making them potential candidates for HPC Clouds.

20 million workstation users are potential candidates for HPC in the Cloud.

HPC CLOUD ADOPTION TRENDS

Currently we are observing several important trends and initiatives fostering HPC in the Cloud:

• **Benefits:** The benefits for design & development engineers and their companies from using computing in the Cloud are undoubtedly huge (as explained above) and will drive adoption as soon as this is recognized by engineers and scientists.

• **Software Vendors:** There is a growing cloud acceptance by ISVs, particularly ANSYS announcing ANSYS Enterprise Cloud on Amazon AWS. There are other ISVs like CD-adapco with token-based cloud-licensing, or most elegantly COMSOL allowing free cloud usage for customers who have a network license.

• **Consumer & Enterprise Clouds:** Because cloud services are widely accepted in the consumer and enterprise communities (e.g. for ERP, CRM, administration), the acceptance of Cloud services in companies' R&D departments seems to be a natural next step and will happen with a certain time delay.

Cloud market accelerators

- **Cloud benefits are obvious**
- **ISV adoption is growing**
- **Consumer & enterprise clouds**
- **Supply chain is under pressure**
- **Big data demands big compute**
- **3D printing needs simulation**
- **Digital natives live in clouds**
- **Government funded initiatives**

• **Supply Chain Pressure:** Large manufacturing companies more and more expect their supply chain partners to perform high-quality end-to-end simulations on HPC systems.

• **Big Data Analytics:** Big data is a big candidate for clouds. Collecting data en-mass is one thing, but extracting insight and knowledge out of tera- or even peta-bytes of data delivers much higher value. This analytics process is a natural cloud candidate.

• **3D Printing:** Experts expect 30 million people to enter the 3D printing field by 2025. Many of them will work with CAD and CAE, but not everybody would buy their own HPC server then.

• **Digital Natives:** All those kids born after 1995 are called 'digital natives' and they are growing up with very different understanding for privacy and security. Most of their digital activities already happen in the cloud, of course.

• **Cloud-related Initiatives:** And, last but not least, US and international initiatives like the Missing Middle strongly supported by companies like Intel, the National Center for Manufacturing Sciences (NCMS), the UberCloud HPC Experiment, Cloud service providers like Amazon AWS, Microsoft Azure, Nimbix, Sabalcore and many others, and electronic magazines like HPCwire, Desktop Engineering, Bio-IT World, and HPC Today are currently creating strong awareness for the benefits of in-house HPC and HPC as a Service in the wider SME market.

In summary we foresee the HPC in the Cloud market growing especially towards the long tail [8], with tens of millions of very small to medium size enterprises (VSMEs, including the one-man shows) which can't (or which won't) afford to buy expensive HPC servers for their different simulations. They all are or will become potential candidates for HPC Cloud services.

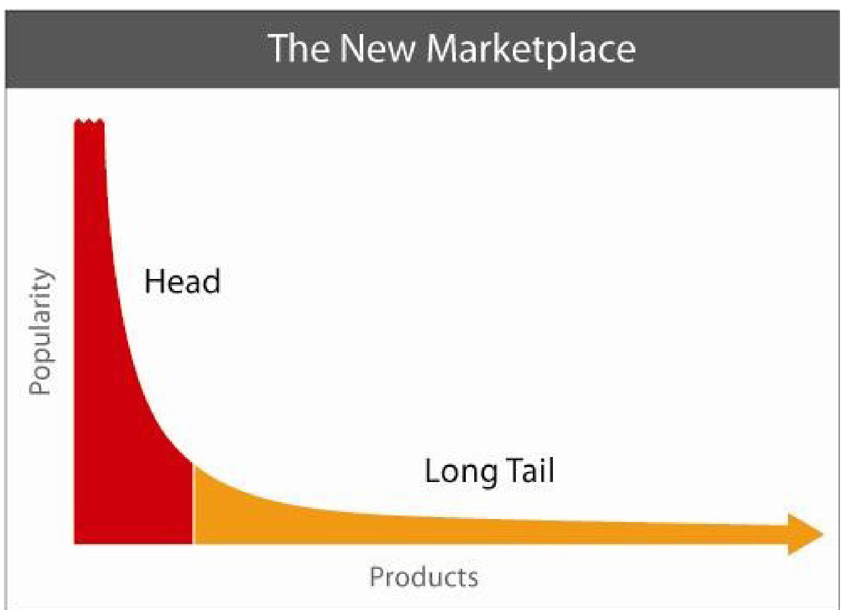

The Long Tail of very small to medium businesses considering HPC in the Cloud [8]

And for existing service providers – be it ISVs or Cloud resource providers – these individual businesses are simply too small to be able to feed their expensive sales people, in contrast to a fully automated online marketplace which doesn't have these high overhead cost. Therefore, naturally, software and cloud service providers are concentrating on the larger accounts, while the only way the tens of millions of VSME are and will be served is an online marketplace for engineering and scientific applications in the Cloud.

For existing service providers small businesses are simply too small to be able to feed their expensive sales force.

The HPC Cloud Marketplace – Field of Dreams ?

When we started UberCloud in 2012 it was the time in Silicon Valley when every entrepreneur and every venture capitalist was talking about startups like Uber and Airbnb. Although there were many online marketplaces before, these two shooting stars made it to every front page of every news magazine because of their ingenious and disruptive approaches shaking up the long-established taxi and hotel businesses. Influenced by this enormous hype it was quite obvious for an HPC veteran and a Cloud practitioner (that's us) to give an online marketplace for HPC as a Service some thoughts. But then we found that there was no solid information about end-user experience nor reliable market analyses in HPC Cloud; nothing to build on. So we decided to abandon the Build It He Will Come approach which worked in this great movie Field of Dreams but won't work in HPC; engineers and scientists tend to be more realists …

To explore HPC in the Cloud we chose a crowd-sourcing approach and invited our engineering and scientific community to participate in UberCloud experiments.

Instead, from the very beginning, we selected a crowd-sourcing approach and invited our whole engineering and scientific community to participate in our experiments, using our experimentation platform and technology, and providing feedback. That way, successively, we developed the UberCloud community platform for HPC Cloud information, the HPC experiment (sponsored by Intel) with 175 free voluntary cloud experiments, the HPC container technology based on Docker, and finally the UberCloud Marketplace with currently 30 service stores and about 100 cloud service products.

The UberCloud Ecosystem
Community information platform, 175 free and voluntary HPC cloud experiments, HPC containers based on Docker, and the Marketplace with currently 30 service stores and 100 cloud products.

BEHIND THE SCENES OF THE UBERCLOUD

In articles and market research in mid-2012 we mostly found educated guesses about cloud adoption being so slow in the scientific and engineering community while in many enterprises cloud applications were widely in use. So we decided to start a few real-life experiments together with engineers taking their applications to the cloud. The first Call for Participation in June 2012 attracted an amazing 160 companies from over 30 countries, and we were able to perform 25 UberCloud experiments.

UberCloud Experiment:
Online platform for engineers, scientists, service providers to discover, explore, and understand HPC Cloud roadblocks.

The UberCloud Experiment provides an online platform for engineers, scientists, and their service providers to discover, explore, and understand the end-to-end process of accessing and using HPC in the Clouds, and to identify and resolve the roadblocks. End-users, software providers, resource providers, and HPC experts are collaborating in teams, jointly solving the end-user's application in the cloud. That way we were able to conduct 13 successful experiments. The other 12 teams failed, because of different reasons, such as: the ISV was not able to support the experiment; the end-user didn't get approval from management; the end-user dropped out because of a sudden in-house project; and more.

Until today (July 2015), the UberCloud Experiment has attracted more than 3000 organizations from 72 countries, allowing us to build 175 teams so far. Recently, the UberCloud Marketplace [9] has been added, and Intel sponsored 2 Compendiums (and a third one is in preparation) which contain more than 60 cloud case studies. SMBs who are aiming at developing better products faster are invited to join the free voluntary UberCloud Experiment to explore HPC as a Service, in the cloud, [10].

UBERCLOUD'S HOLY GRAIL – SOFTWARE CONTAINERS

In early 2013, after about 50 cloud experiments, and an in-depth analysis of the 22 experiments which stumbled across the cloud roadblocks and failed, we believe we found the solution: Linux containers. Timing was perfect: instead of starting from scratch we turned to Docker in March 2013 and attended one of their very first developer workshops in San Francisco. Since then we are developing more and more HPC features into Docker which after two years recently resulted in our first UberCloud application containers for engineering and scientific applications.

What is Docker?
Open platform for building, shipping and running distributed applications. It gives programmers, development teams and operation engineers the common toolbox to take advantage of the distributed and networked nature of modern applications.

UberCloud Containers are ready-to-execute packages of software. These packages are designed to deliver the tools that an engineer needs to complete his task in hand. The ISV or Open Source tools are pre-installed, configured, and tested, and are running on bare metal, without loss of performance. They are ready to execute, literally in an instant with no need to install software, deal with complex OS commands, or configure. The UberCloud Container technology allows wide variety and selection for the engineers because they are portable from server to server, Cloud to Cloud. The Cloud operators or IT departments no longer need to limit the variety, since they no longer have to install, tune and maintain the underlying software. They can rely on the UberCloud Containers to cut through this complexity. This technology also provides hardware abstraction, because the container is not tightly coupled with the server (the container and the software inside isn't installed on the server in the traditional sense). Abstraction between the hardware and software stacks provides the ease of use and agility that bare metal environments lack.

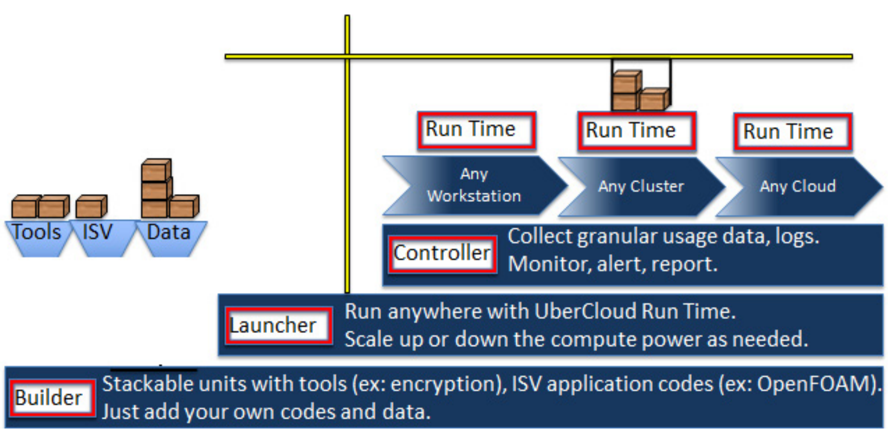

UberCloud Containers for applications in the HPC Cloud

The Cloud Challenges

However, HPC in the Cloud comes with challenges too: it's a new business and working paradigm; security, privacy, and trust in service providers can be an issue; conventional software licensing is slowly including the pay-per-use or subscription model; Internet bandwidth often depends on the 'last mile' unable to accommodate heavy data transfer; unpredictable costs of cloud computing can be a problem in securing a budget for a given project; and there is often a lack of easy, intuitive self-service access and use of cloud resources. But in the following we will describe how we were able to reduce or even remove most of these cloud challenges.

Cloud Challenges

- **Security, privacy, trust**
- **New business paradigm**
- **Software licensing**
- **Slow Internet bandwidth**
- **Unpredictable cost**
- **Lack of easy access**

IN-HOUSE VERSUS IN-CLOUD, THAT'S THE QUESTION

An UberCloud survey on LinkedIn about engineers' concern about cloud computing revealed that many engineers tend to compare the benefits of workstations versus servers versus clouds in a somewhat misinformed way. In fact, most of them compare the positive aspects of their workstation with some of the (more apparent) roadblocks of the cloud. Example: data transfer; there is obviously no data transfer necessary when you compute your task in your workstation, and yes, there often is heavy data transfer from the cloud back to your workstation (often in the order of several GBytes of result data). But this is simply due to the larger computations in the cloud which are impossible to do on our workstation. And the transfer speed highly depends on our last mile connection. This we can change…

First, we should answer this question: is my workstation big enough and fast enough for the kind of problems I want to solve? If your answer is YES, it is, then you don't need an HPC Server at all, and you don't need HPC in the Cloud, full stop.

But if your answer is NO, it isn't, my workstation is not big enough and fast enough for the kind of problems I want to solve, then a reasonable way to look for viable alternatives is to compare the two options and check which one is most reasonable for you: in-house HPC Ser-

Feature / Functionality	In-house HPC Server*	Remote HPC Cloud**
Procurement	can take months	depends on provider (days)
Budget	CAPEX, inflexible	OPEX, flexible
System operations	Complex, costly	none
Flexibility, agility	low	high
Reliability	single point of failure	redundant
Average utilization	20% – 80%	100%
Security	medium	high
Technology	ages quickly	frequently updated
Data transfer speed	high	depends on your last mile
Full control	high	medium to high**
Software licensing	expensive (e.g. annual)	pay per use or subscription
Access	mastering many scripts	seamless
Job wait time	depends on job queue	none

Comparing challenges for in-house versus cloud servers (*on average; **and this depends on your cloud provider and cloud technology you use).

ver versus remote HPC Cloud, and NOT versus your own workstation which already proved to be useless for your more complex, more challenging tasks. Servers against Clouds! We are certainly fully aware that such a comparison has to come along with some generalizations and simplifications.

Procurement: Bound to approval from upper management; lengthy process: purchase planning, standards determination, specifications development, supplier research and selection, value analysis, financing, price negotiation, making the purchase, supply contract administration, inventory control and stores, waiting in line, accepting delivery, installing and certification testing the hardware, training people, and other related functions. On the other hand, cloud services are short-term on demand or on reservation available for comparatively low cost.

Budget: CAPEX related assets have to be approved often by upper management, while OPEX usually falls into the responsibility of mid-management or even the employee.

Operations, maintenance: To run a compute server, one needs specially trained people; regularly upgrade system and application software; handle and fine-tune the system, workload and resource management; deal with power consumption, cooling, and room temperature; take care of downtime and user productivity. In contrast to cloud where none of these efforts apply.

Flexibility: You bought the system, and you stick with it for three years, even if for some applications your system might not be optimal. Completely different with clouds: there is flexibility in the choice of hardware, software, related tools, timing, pricing, utilization, and so on.

Agility: Users can self-service against a large and flexible service catalog. They can pick up whatever they need, whenever they need it. And, long wait queues can kill engineering inspiration.

Reliability: One single system result in one single point of failure. With clouds reliability can easily be achieved by working with several cloud providers.

Average utilization: Especially in small and medium enterprises, server utilization is unpredictable, because of different project deadlines, the engineers' vacations or business trips, and weekends where these servers are often almost completely 'jobless', resulting on low server utilization. In clouds, in contrary, prices per core per hour are calibrated with high utilization assumed.

Security: Please see our next paragraph about cloud security.

Technology: New technologies and products are coming to market at fast pace. To make up for this we have to regularly upgrade our existing equipment and thus invest even more money. Then we have to stick with our existing systems for at least throughout the depreciation phase. Clouds: to stay competitive providers are regularly refreshing their infrastructure. Therefore, in the cloud, we can shop around for the fastest and best suited hardware and services.

Data transfer: Here we should differentiate between intermediate results and the final dataset: Intermediate data can often be stored in Dropbox or Box.com which have fast connections to the clouds. For final datasets there are transfer technologies which compress and encrypt the data, can send it in parallel chunks, stream it back to the user, or applying data reduction from VCollab.

Full control over your assets: With the advent of the UberCloud containers, additional functionalities like collecting granular usage data, logs, monitoring, alerting, reporting, and more are bringing back the control a user wants.

Software licensing: Due to customer demand and competitive forces, more and more ISVs are now adding flexible cloud-based licensing models, e.g. for monthly, weekly, daily, or even hourly usage.

Access: In the meantime access to many clouds can be considered seamless, and is included in the bill.

Wait time: With your own compute server during peak loads when you need your server the most, your jobs are sitting in the wait queue for hours at end. Clouds change that, simply because Clouds offer “infinite” resources; and if the resources of one cloud provider are not “infinite” enough, you can move on to the next cloud provider. Clouds inherently have very short or no wait time at all.

An Honest Word About Security in the Cloud...

Here we quote a recent article from Mike Kavis, Vice President & Principal Cloud Architect at Cloud Technology Partners [11] about security of SaaS solutions, who refers to a study from Alert Logic [12]. What this report proves is that the security threats are the same, regardless of where the data resides. What is even more interesting is that the success rate of penetrations from outside threats was much higher in enterprise data centers than in external cloud environments. Based on this information, skeptics should dismiss the notion that data cannot be as secure in the cloud as it is behind the corporate firewall. Mike Kavis call this the “hypocrisy” of enterprise IT. It is almost comical when people declare SaaS to be unsecure while their company transmits unencrypted email, staff members have company information on personal unsecured mobile devices, and a number of systems run on unsupported or unpatched versions of software on premises.

The Hypocrisy of Enterprise IT:
People declare SaaS to be unsecure while their company transmits unencrypted email, staff members have company information on unsecured mobile devices, a number of systems run on unsupported unpatched software versions on premises.

... AND ABOUT 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 SMEs, as stated by a poll during a recent UberCloud Webinar: 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 ISVs like Simulia, and in the second group we find ANSYS with their recently announced Enterprise Cloud, Autodesk with CFD Flex, CD-adapco with Power on Demand tokens, and COMSOL with their cloud-ready network license.

Software Licensing
Companies like ANSYS, Autodesk, CADFEM, CD-adapco, CFDsupport, COMSOL, DYNAmore, Matlab, NICE, and Numeca are all providing advanced licensing for the cloud. Laggards might lose market share.

HOW COST EFFICIENT IS HPC IN THE CLOUD?

An IDC study from 2007 [02] found that only about 7% of the total cost of acquiring and operating an HPC system is coming from the hardware (additional analysis shows that this situation hasn't changed since then). The much bigger portion of the pie is coming from the high cost of expertise (staffing), equipment, maintenance, software, training, i.e. the Total Cost of Ownership (TCO). A recent cost analysis for an in-house HPC system versus an HPC Cloud service [13] shows that overall cost depends on the average utilization of the in-house sever. The following table shows the total cost of one CPU core hour for a 16-node (256-core) in-house HPC server depending on average percentage of utilization (or 'number of busy nodes') of this server.

For this example, with an average utilization of say 25%, the real cost for one core hour

Busy compute nodes	1	2	3	4	5	6	8	12	16
Utilization	6.3%	12.5%	18.8%	25.0%	31.3%	37.5%	50.0%	75.0%	100%
Cost per core per h	$2.37	$1.19	$0.79	$0.60	$0.48	$0.40	$0.30	$0.20	$0.15

Cost per core hour of a 16-node server depending on server utilization.

is $0.60, i.e. 4 times higher than for a 100% utilized cluster. For 50% utilization, which is not uncommon in HPC, the cost of one core hour is still $0.30. For comparison, the actual cost of one core hour of Microsoft's Azure most powerful HPC Cloud instances D and G vary between $0.09 - $0.30, [14].

One-year cost of HPC for in-house versus cloud versus hybrid computing.

According to Table 1, the cost of the cloud (say $0.20 per core hour) will exceed the one for the in-house solution only at about 75% (and higher) average utilization of the in-house HPC cluster, a situation which we only heard of from academic or big-industry Supercomputing Centers serving hundreds or even thousands of users. If we extrapolate these numbers to one year and include a hybrid (in-house / in-cloud) scenario, we'll get the results shown in Figure 4 (see details in [13]).

UberCloud Experiments Accelerating HPC in the Cloud

The voluntary free UberCloud Experiment provides a crowd-sourcing 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 [10]. 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 3000+ organizations from 72 countries (status July 2015). The organizers were able to build 175 of these teams, in CFD, FEM, and Computational Biology, and to publish many case studies in two Compendiums. The 2015 Compendium of case studies will be published in August, again sponsored by Intel, [15].

In an experiment, the team participants are the industrial end-user, a suitable resource provider, the software provider, and an HPC expert. 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. Finally, the team will get together, extract lessons learned, and present further recommendations as input to their case study.

As a glimpse into the wealth of practical use case results so far, here is a selection of Cloud experiments (selected from 175 so far) demonstrating the wide spectrum of CAE applications in the Cloud.

TEAM 94: GAS-LIQUID TWO-PHASE FLOW APPLICATION

End-User in this team was Kyoji Ishikawa from Chiyoda Corporation in Japan, together with Wim Slagter from ANSYS and Hiroyuki Kanazawa from resource provider Fujitsu Ltd. which provides HPC services including the virtual desktop function under the name of Technical Computing (TC) Cloud. The use case was the evaluation of the rate of gas entrainment due to liquid flow in the liquid storage facility for an energy plant (see Figure 5). Gas-liquid two-phase flow was simulated using ANSYS Fluent software.

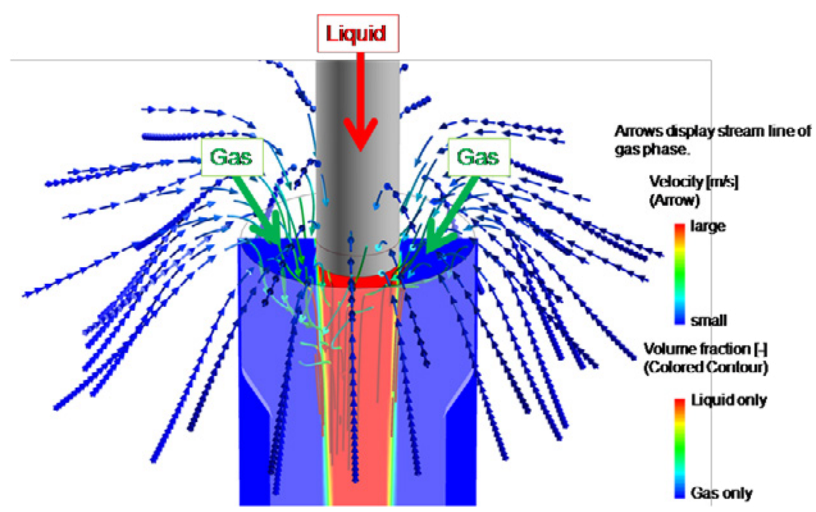

Visualization showing flow path and volume fraction of liquid and gas.

Because the end-user sometimes experiences a computational resource shortage for urgent, unexpected or unplanned projects, he decided to explore the use of remote cloud computing. Two simulation cases were carried out with total mesh numbers 300,000 and 2 million. Computation time on the in-house system for case 1 was 20 hours and for case 2 required 5 days. Cloud computing resources were 32 parallel cores (16 cores × 2 nodes) and Gigabit Ethernet or InfiniBand (IB). Pre- and post-processing has been done with remote visualization. That way the end-user was able to reduce the internal run-time from up to 5 days down to less than a day.

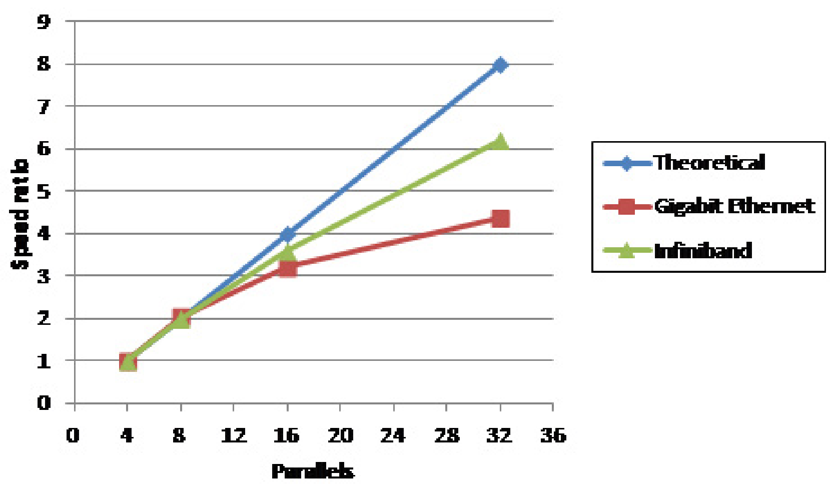

Cloud speed-ups for parallel Fluent simulations on 4, 16, and 32 cores.

TEAM 99: NORTH SEA ASSET LIFE EXTENSION – ASSESSING IMPACTS ON HELICOPTER OPERATIONS IN THE CLOUD

This team consisted of the end-user Dan Hamilton from Atkins Energy, software provider James Britton, CD-adapco, resource provider Jerry Dixon, OCF, and team mentor Dennis Nagy from BeyondCAE. The team tested the feasibility of using cloud services for the simulation of airflow over an offshore platform using STAR-CCM+ from CD-adapco, to determine the change in conditions within the helideck landing area as a result of geometrical changes stemming from a life extension project on an existing North Sea asset.

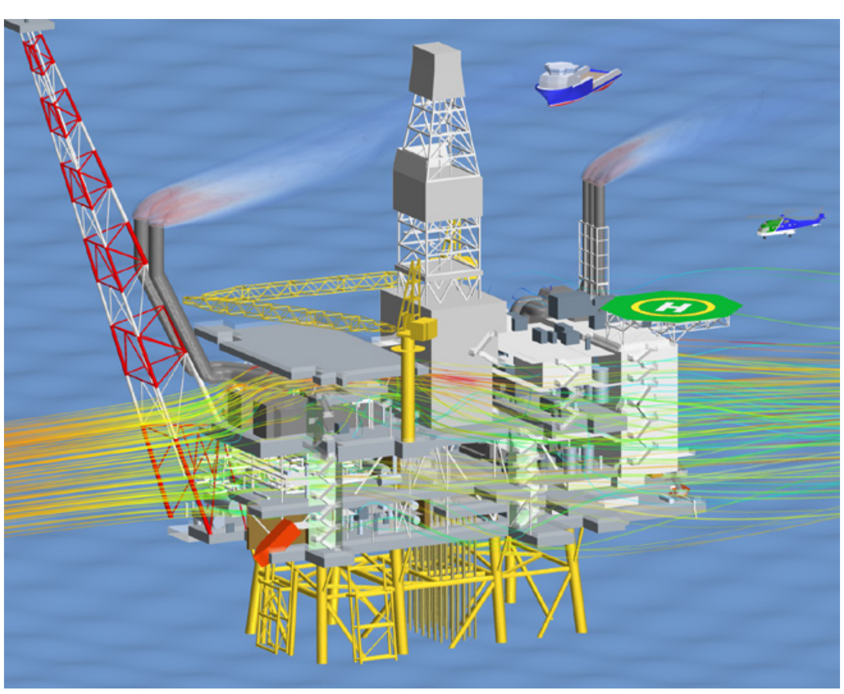

North Sea Asset Life Extension – Assessing Impacts on Helicopter Operations in the Cloud

For OCF, the availability of the STAR-CCM+ "Power-on-Demand" licensing is a perfect fit for their enCORE service. The user experience of using STAR-CCM+ in batch on enCORE was identical to that on in-house hardware. As the simulation files can be large, enCORE's policy of not charging for bandwidth usage is appealing. Having a resource like enCORE available for use allows to bid for and propose work requiring computational resources, which exceeds what is available in-house.

TEAM 118: COUPLING IN-HOUSE FE CODE WITH ANSYS FLUENT

This team's end user was Hubert Dengg from Rolls-Royce, software providers were Wim Slagter and René Kapa from ANSYS, resource providers and team experts were Thomas Gropp and Alexander Heine from CPU 24/7, and Marius Swoboda from Rolls-Royce Deutschland acted as HPC/CAE expert.

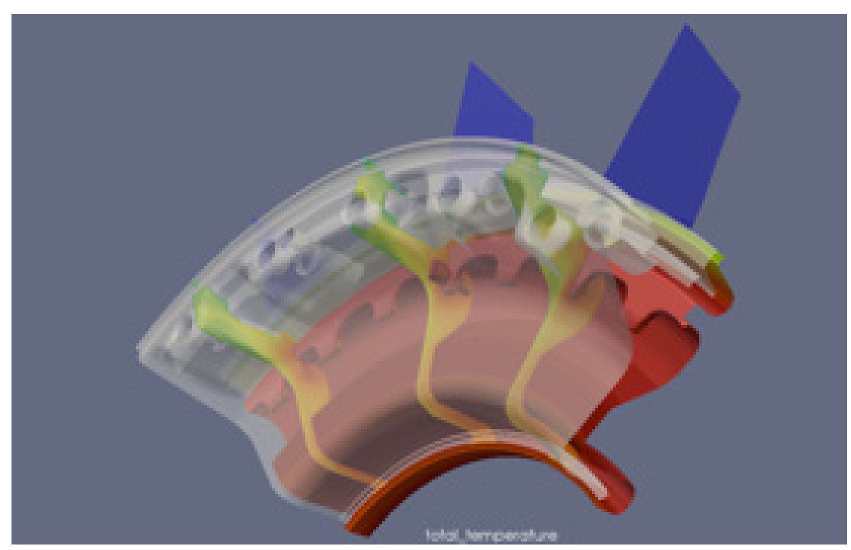

Contours of total temperature for a jet engine component.

A transient aerothermal analysis of a jet engine high pressure compressor assembly has been performed using FEA/CFD coupling technique. The aim of this cloud experiment was to link ANSYS Fluent with an in-house FEA code. The conjugate heat transfer process is very consuming in terms of computing power, especially when 3D CFD models with more than 10 million cells are required. As a consequence, using cloud resources was of great benefit. The computation was performed on a 32-core 2-node system. The calculation was done in cycles in which the FE code and Fluent CFD ran alternating, exchanging their results.

Outsourcing of the computational workload to an external cluster allowed the end user to distribute computing power in an efficient way – especially when the in-house computing resources were already at their limit. Bigger models usually give more detailed insights into the physical behavior of the system. In addition, the end user benefited from the HPC provider's knowledge of how to setup a cluster, run applications in parallel based on MPI, create a host file, handle licenses, and prepare everything needed for turn-key access to the cluster.

TEAM 142: VIRTUAL TESTING OF SEVERE SERVICE CONTROL VALVE

The end-user of this was Mark A. Lobo from Lobo Engineering. Autodesk provided CFD Flex and the supporting cloud infrastructure. And the application experts were Jon den Hartog and Heath Houghton from Autodesk.

Flow control valve specifications include performance ratings in order for a valve to be properly applied in fluid management systems. Control systems sort out input parameters, disturbances and specifications of each piping system component. System response is a function of the accuracy of control valves that respond to signals from the control system. Valve performance ratings provide information to the system designer that can be used to optimize control system response.

The premise of this project was not to explore virtual valve testing and to evaluate the practical and efficient use of CFD by the non-specialist design engineer. As a benchmark, the end user had no prior experience with the Autodesk software no formal training in the software, and he was depended on the included tutorials, help utility, thorough documentation to produce good results and good data.

In this project, over 200 simulations were run in the cloud. Given the runtimes involved and allowing for data download upon completion of the runs, it is possible to be solved within one day. For an engineer with 1 simulation license on a single workstation this would have required 800 hours (approximately 30 days) to complete if the simulations were running nonstop one after another.

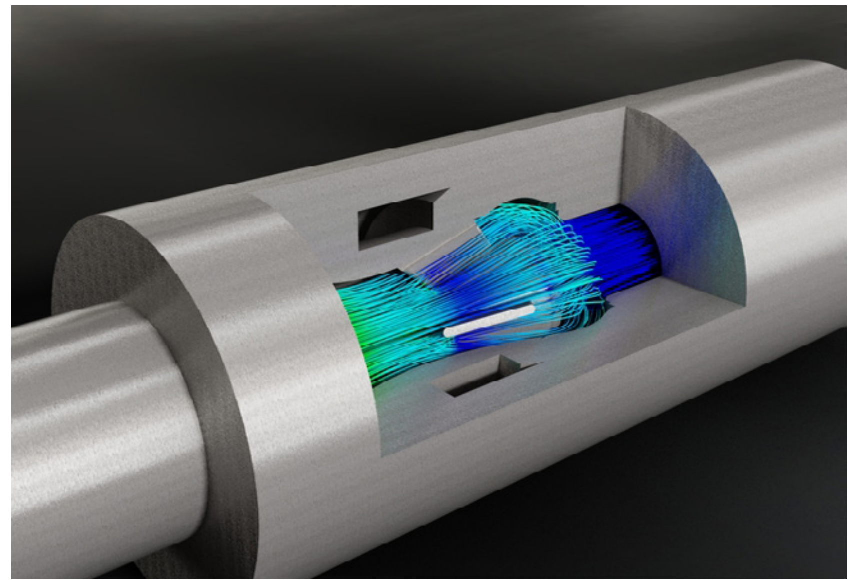

Control valve model with idealized flow path to minimize effects of a complex body cavity and trim design.

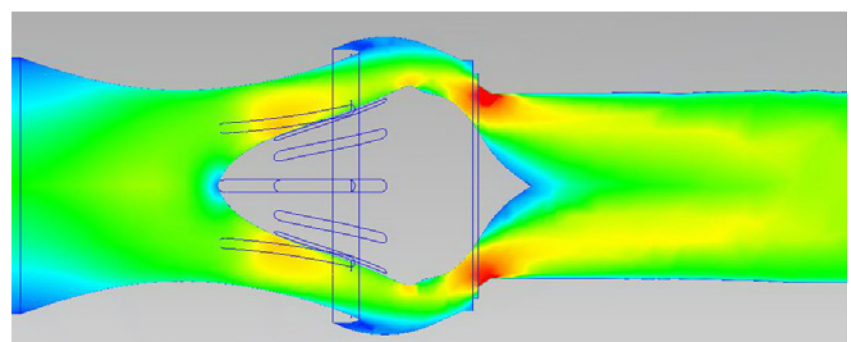

The control valve restriction components or "trim" reduce the annular area as the cavity profile on the right moves to the left.

TEAM 156: PULSATILE FLOW IN A RIGHT CORONARY ARTERY TREE

This team consisted of the end-user Prahlad G. Menon, from Carnegie Mellon University, cloud resource providers Amazon AWS and Nephoscale, and Burak Yenier from UberCloud who containerized OpenFOAM and post-processing tool ParaView. In this study, blood flow inside a patient-specific right coronary artery tree (see Figure 11) including one inlet and a total of 7 outflow branches was studied under realistic unsteady flow conditions after segmentation from tomographic MRI slice images obtained across the whole human body of a male volunteer.

The finite volume mesh consisting of 334,652 tetrahedral cells was passed over as input to the icoFoam solver in OpenFOAM to solve the blood flow, motion and equilibrium under the action of external forces.

The UberCloud OpenFOAM container was found to be far more user friendly – in terms of both simplicity and ease of access (via SSH) – than the supercomputing resources use by the end-user on a daily basis. Parallel scalability of the simulated problem good on the remote computing resource; a 2-3 fold speed improvement was noted for 16 core parallelism on the remote UberCloud container in contrast with an equivalent local simulation run on just 4 parallel cores of a local quad-core machine.

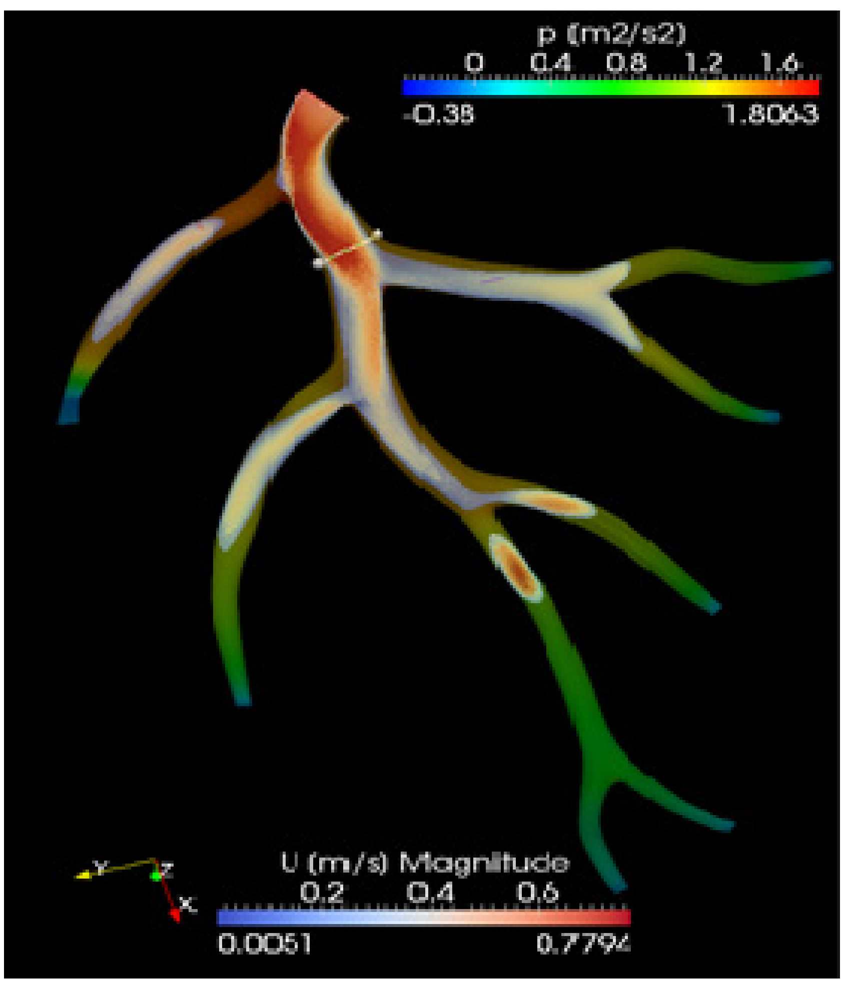

Wind Turbine Aerodynamics Study

The end-user and CFD expert of this team was Praveen Bhat, a technology consultant from India INIA, the software Provider was Metin Ozen, an ANSYS reseller in California, and the resource provider was ProfitBricks.

The team evaluated wind turbine performance using ANSYS CFX on a 62 core HPC cloud server with 240 GB RAM from ProfitBricks. The ANSYS software was running in UberCloud's new application container. The CFD simulation is performed to calculate the pressure distribution and velocity profiles around the wind turbine blades with average wind speed of 7 to 8 m/min. Figure 12 highlights the velocity distribution around the wind turbine blades.

The computation requirement for a fine mesh of 2.5 million cells is high and makes it impossible to run on a normal workstation. The HPC cloud allowed for solving very fine mesh models and for drastically reduced solution times of about1.5 hours.

The UberCloud ANSYS container enabled easy access and use of the Cloud server, and the regular UberCloud auto-update module through email provided the huge advantage of continuous monitoring job progress without any requirement to log-in to the server and check the status.

TEAM 165: WIND TURBINE AERODYNAMICS STUDY

The end-user and CFD expert of this team was again Praveen Bhat, a technology consultant from India INIA, the software Provider was D-adapco with software STAR-CCM+, and the resource provider was NephoScale in California.

The aerodynamics study was performed using STAR-CCM+ simulation software on a 40 (virtual) core HPC cloud with Star-CCM+ 9.06 running in the new UberCloud application container. The major aim of this experiment was to demonstrate the dramatic improvements concerning ease of access and use of cloud resources.

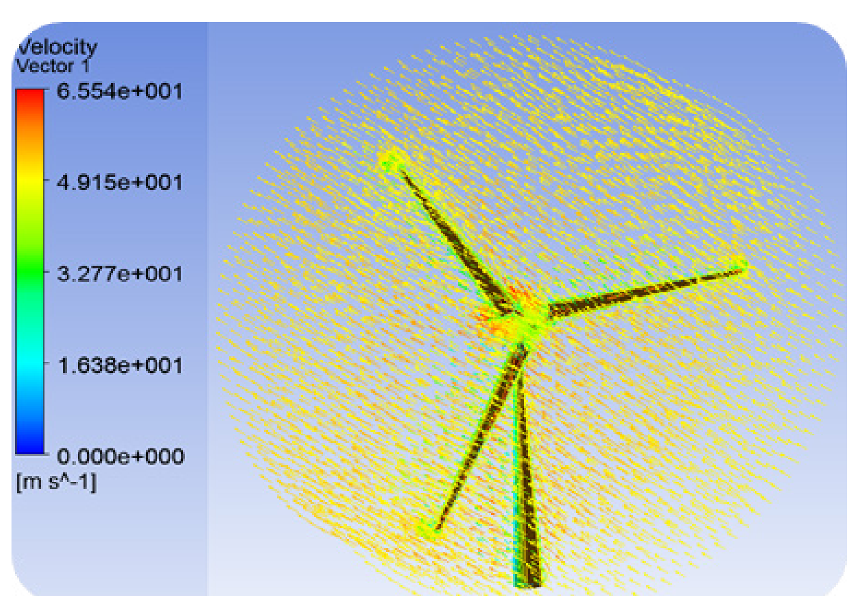

Vector plot of velocity profiles around the wind turbine blades

The effort for this experiment was 30 hours from the end-user for simulation setup, technical support, reporting and overall management of the project; 2 hours for UberCloud for regular system management to build the Star-CCM+ image, transfer it on to the server, launch the container with the image, monitoring the server etc. Computing resources were used for 600 core-hours in total. The combination of NephoScale Cloud, UberCloud Containers, and CD-adapco Star-CCM+ helped in speeding up the simulations, completing the project within the stipulated short time frame, and demonstrated the extreme user friendliness of the access and use of the cloud resources.

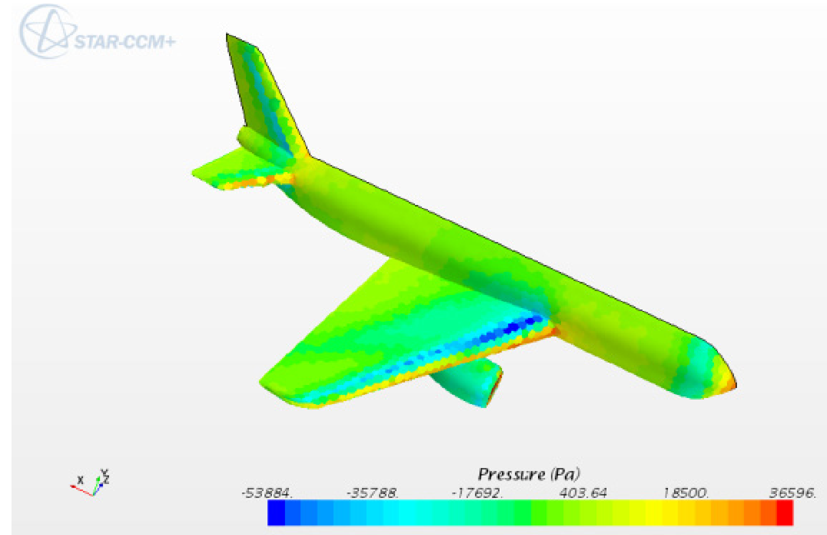

Plot of pressure distribution on the aircraft

TEAM 166: CFD STUDY ON FLIGHT AERODYNAMIC

The end-user of this team was Stefan Castravete from Caelynx in Romania, software provider was Ulli Göhner from DYNAmore, and resource providers and HPC experts were Thomas Gropp, Alexander Heine, and Christian Unger from CPU 24/7. The model is a full Toyota Yaris Sedan with frontal airbag which is analyzed via Finite Element Method and consists of 1.5 million elements using LS-DYNA. The effort invested was: for the end user 10 hours simulation setup, technical support, reporting and overall management of the project; UberCloud support 3 hours monitoring & administration of host servers; and 120 core-hours used.

Car behaviour at frontal impact with a rigid wall at 56 km/h

[1] Council of Competitiveness, 'Make', 'Reflect', 'Reveal', & 'Compete' studies 2010. http://www.compete.org/ [2] Randy Perry, Al Gillen. Demonstrating business value: Selling to your C-level executives. IDC 2007. [3] Earl Joseph, Chirag Decate, Steve Conway, IDC, Creating economic models showing the relationship between investments in HPC and the resulting financial ROI and innovation, https://hpcuserforum.com/roi.html, 2013. [4] 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/ [5] The Statistics Portal about SaaS: http://www.statista.com/statistics/273642/worldwide-software-as-a-service-revenue-forecast/ [6] Steve Conway et al. IDC HPC End-User Special Study of Cloud Use in Technical Computing. June 2013. [7]Jon Peddie. Workstations continue solid growth in Q3 2013. http://jonpeddie.com/press-releases/details/workstations-continue-solid-growth-pace-in-q313-jon-peddie-research/ [8] Chris Andersen, 2006. The Long Tail. http://www.longtail.com/about.html [9] The UberCloud Marketplace, https://www.TheUberCloud.com [10] The UberCloud Experiment, https://www.theubercloud.com/hpc-experiment/ [11] Mike Kavis, 2015. 3 Arguments Against Enterprise SaaS That Fall Flat. http://www.cloudtp.com/2015/01/12/3-arguments-enterprise-saas-fall-flat/ [12] Alert Logic, 21013. https://www.alertlogic.com/resources/security-intelligence/download-cloud-security-report-spring2013/ [13] Wolfgang Gentzsch, 2014. How cost efficient is HPC in the Cloud. https://www.theubercloud.com/cost/ [14] Cost of a Microsoft Azure HPC Cloud G5 instance in June 2015. http://azure.microsoft.com/en-us/marketplace/?term=openfoam [15] The UberCloud Compendium of case studies, 2013, 2014, and 2015: https://www.theubercloud.com/ubercloud-compendium-2013/ https://www.theubercloud.com/ubercloud-compendium-2014/ https://www.theubercloud.com/ubercloud-compendium-2015 /

About the authors

Wolfgang Gentzsch is the President of The UberCloud. He is an industry executive HPC and Cloud consultant and the Co-Chairman of the Intl. ISC Cloud & Big Data Conferences.

Previously, he was an Advisor to the EU projects DEISA and EUDAT, directed the German D-Grid Initiative, was a Director of the Open Grid Forum, Managing Director of the North Carolina Supercomputer Center (MCNC), and a member of the US President's Council on Science and Technology PCAST. In the 90s he founded several HPC companies, e.g.Gridware (with its resource and workload manager Grid Engine) acquired by Sun Microsystems, and he became Sun's Senior Director of Grid Computing.

@wolfgent
http://www.linkedin.com/in/wolfganggentzsch

Burak Yenier is the CEO of The UberCloud. He is an expert in the development and management of large-scale, high availability systems, and in many aspects of the cloud delivery model including information security and capacity planning. He has been working on Software as a Service since early 2000 and held management positions in software development and operations in several companies. In his most recent role as the Vice President of Operations Burak managed the multi-site datacenter and digital payment operations of a financial Software as a Service technology company located in Silicon Valley where he built cloud infrastructure and operations from scratch and for scale.

@BurakYenier
http://www.linkedin.com/in/buraky

Lab Review

How do we test ?

HPC Labs

HPC Labs is the technical unit of the HPC Media group and totally independent of the manufacturers. HPC Labs' mission is to develop methodologies and materials testing and software metrics in the high performance IT world. Capitalizing on best practices in the field, these tools are based on several decades of joint experience of the laboratorys' management.

HPCBench Solutions

Specifically designed for HPC Review, the HPCBench Solutions assess not only performance but also other equally important aspects in use, such as energy efficiency, sound volume, etc. To differentiate synthetic protocols like Linpack, these protocols allow direct comparison of solutions pertaining to the same segment, resulting in a single index taking into account the specific hardware or software tested. For example, an SSD will be tested with the HPCBench Solutions> Storage, while a GPU accelerator will be tested with the HPCBench Solutions> accels. Rigorous and exhaustive, these protocols allow you to choose what will be for you, objectively, the best solution.

HPC Bench Global index 9 108

A single synthetic index to help you compare our test results

HPC REVIEW AWARD

A technical recognition Award

OCZ Saber 1000 960Gb

The Saber 1000 is an enterprise-class SSDs designed for data-intensive applications handling large volumes of data as part of a HyperScale, Web and cloud hosting deployment duties. OCZ developed the 1000 Saber to eliminate input / output bottlenecks in enterprise and datacenter infrastructures.

Aan enterprise-class SSDs designed for data-intensive applications handling large volumes of data as part of a HyperScale, Web and cloud hosting deployment duties.

CONSISTENT PERFORMANCE

As the number of ever-growing dynamic application workloads increase significantly within the enterprise, the performance of sto-

The Saber 1000 incorporates PFM + (Power Failure Management +) which in the case of a sudden power failure, powers the SSD for a period sufficient to preserve the integrity of data. PFM+ is based on a 16-byte log file containing the last instructions to perform in order to avoid corruption associated metadata. Once power is restored, the SSD resumes its duty.

rage resources become more important. SSDs are part of accelerators capable of providing known and predictable performance for standard applications. OCZ has designed the SSD to address this issue and thus be able to reduce the data read queue and thus help improve widely distributed application responsiveness.

A CONSISTENT ARCHITECTURE

Technically, the 1000 Saber uses 19nm technology NAND flash from Toshiba, its parent company, coupled with its in-house Barefoot 3 controller. OCZ has taken great care to incorporate flash memory and controller in order to preserve the architectural coherence and deliver a sustained data rate. The result is a fairly balanced performance of 550 MB / s read and 448 MB / s write. OCZ has specifically designed the Saber 1000 for read-intensive applications. Thus this drive data rate delivers results that show how it favors read operations: 91000 IOPS (read) against 16,000 IOPS (write). A 65% (read) by 35% (write) workload averages 33000 IOPS. A score that meets the requirements of enterprise applications and hyperscale workloads like databases and web hosting.

POWER CUT PROTECTION

The Saber 1000 incorporates PFM + (Power Failure Management +) which in the case of a sudden power failure, powers the SSD for a period sufficient to preserve the integrity of data. PFM+ is based on a 16-byte log file containing the last instructions to perform in order to avoid corruption associated metadata. Once power is restored, the SSD resumes its duty.

A WIDE RANGE OF USES

The Saber 1000 Series, available in capacities of 120, 240, 480 and 960 GB, is suitable for read-intensive applications, such as caching, reading and indexing databases, video on demand, VDI environments, multimedia streaming, web server front end, cloud infrastructure, online archiving, CRM and ECM applications, where read performance is crucial in these environments. In terms of consumption, the SSD consumes only 3.7 watts in operation, a criterion to consider for reducing energy requirements and adding savings due to reduced cooling requirements.

A MANAGED SSD

OCZ developed StoragePeak 1000, a centralized management and administration tool for monitoring SSDs deployed in the enterprise or datacenter infrastructure from a single web interface. StoragePeak 1000 connects to Windows and Linux Pcs and servers and integrates an alert mechanism. The warranty period of the Saber 1000 is 5 years. **RAMON LAFLEUR**

Datacore SANsymphony v10

The enterprise storage virtualization offers many benefits compared to proprietary hardware solutions. DataCore is a key player, who created and developed the concept in 1998. The DataCore SANsymphony-V application is a great example and its latest version 10 offers new features and features that make storage virtualization a persuasive discipline.

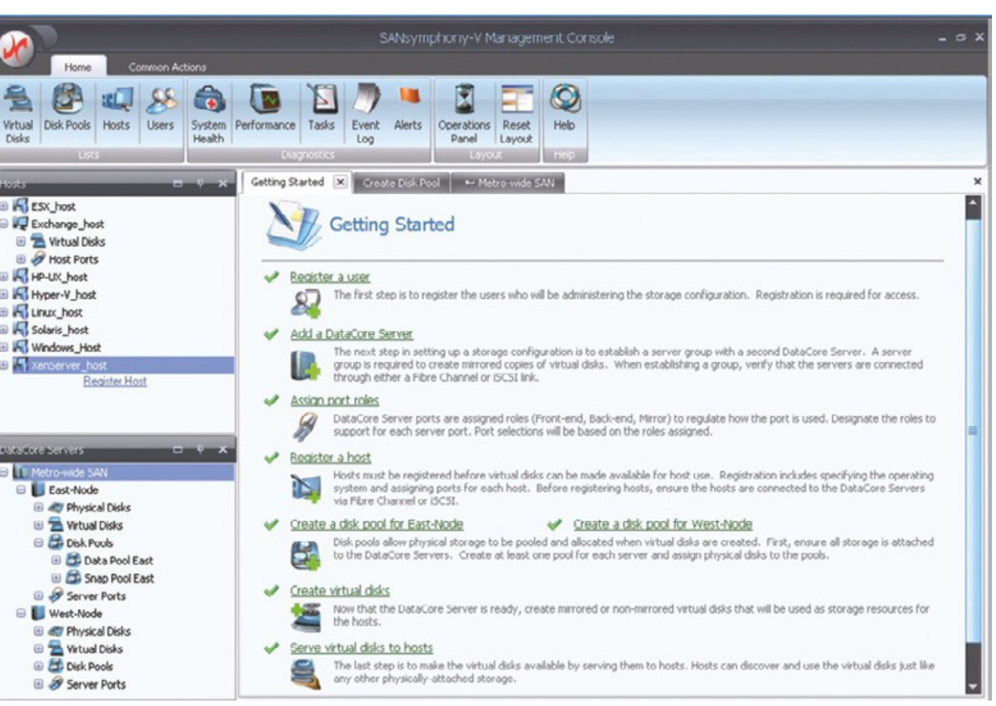

The new features focus heavily on the highest performance possible by the use of Flash. Normally reserved for high-end enterprise disk arrays, SANsymphony-V allows disk activity monitoring and live data migration to improve storage infrastructure performance while allowing the use of heterogeneous vendor devices. The maximum size of the storage cluster supported by SANsymphony extends to 32 nodes and 32 Petabytes. Synchronous mirroring replication also allows up to 100 km distant N + 1 redundant networks.

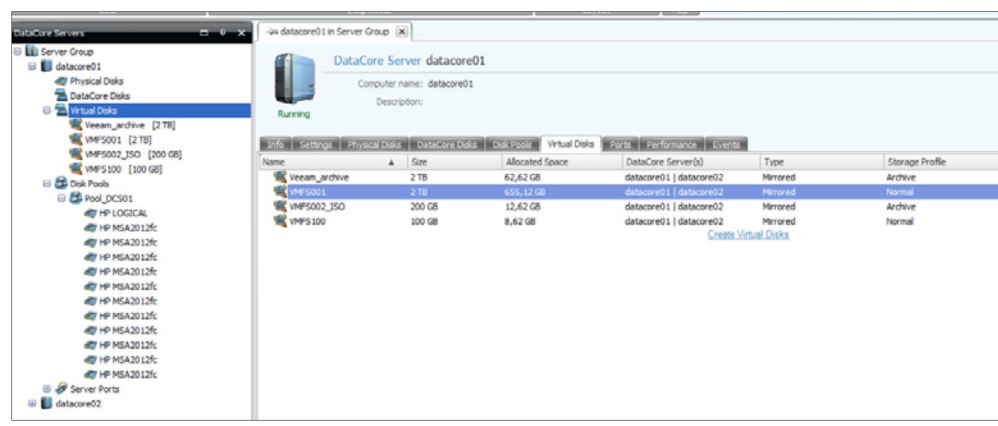

Upon installation, an intelligent deployment wizard simplifies the transformation of local disk and flash storage on the application servers in an elastic storage pool, using the concept of Virtual SAN. The storage space can also be shared on a server cluster, eliminating the need for an external physical SAN. SANsymphony-V is nevertheless able to unify physical and virtual SAN. The software works on the server side as a virtual SAN and is able to administer the related physical SAN.

AN AGNOSTIC APPLICATION FOR HETEROGENEOUS STORAGE

With its independent approach towards the installed storage equipment, SANsymphony-V can use any device or storage resource compatible with the Windows server on which the software is installed. This includes local hard disks, SSDs, SAS storage arrays, SAN and FCoE. This release also supports high-speed 40/56 Gigabit Ethernet interfaces and 16 Gbps Fibre Channel.

Once the software is installed on Windows Server 2012 R2, you can from the management console create a virtual SAN covering local storage and network systems. Before al-

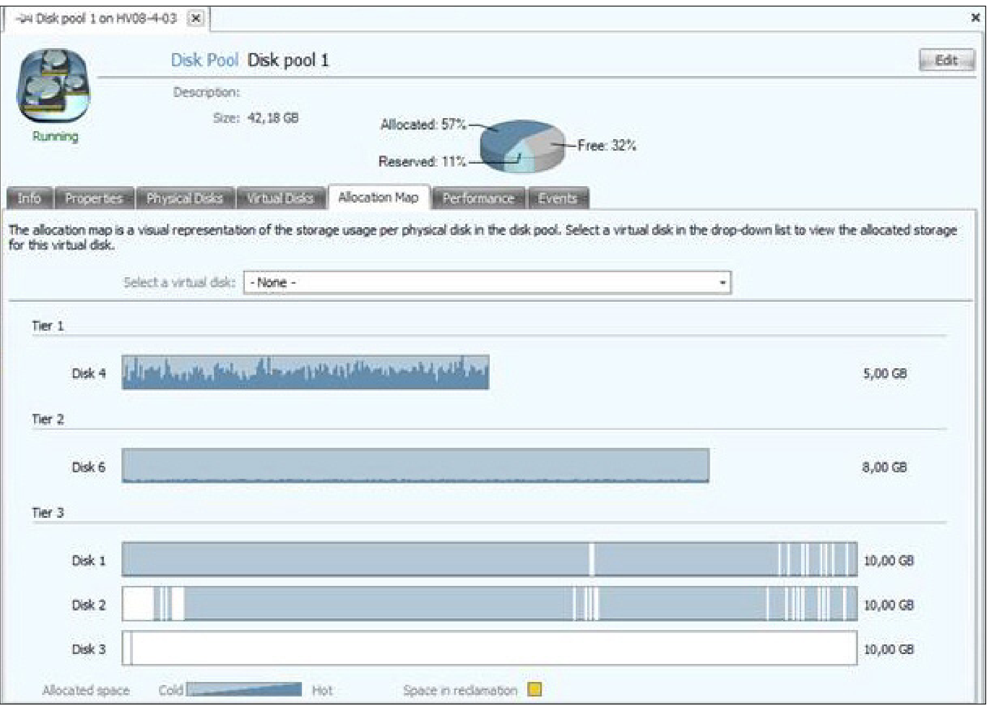

locating the storage resources, the roles of the ports are set to determine if they are reserved for storage or for dedicated access to the physical disks of the local mirrored node (or several nodes). This feature works in real time. If replication were to fail, the virtual disks are available to provide redundancy throughout for the connected hosts and clients.

VIEW AND OPTIMIZE INFRASTRUCTURE PERFORMANCE

Once the hosts connected to the storage node, they are saved using the console. Physical disks and storage devices are grouped into storage pools, and each pool can contain a heterogeneous mixture of drives and interfaces. When the pool is created the administrator should determine the number of required levels (maximum 15) for each pool. The first level is reserved for the fastest resources, typically SAS drives or SSDs. This arrangement has two benefits : inform the software about the hierarchy of the elements of the storage infrastructure, and activate the auto-tiering feature of SANsymphony-V, which migrates data on the fly from the slow resources to the faster ones in order to optimize the infrastructure's global performance. The administrator's dashboard shows the use of the storage resources in graphical form. A visual heat map also indicates the performance of each resource and pool of the infrastructure in realtime, eliminating the tedious work of needing to measure - or guessing - when it comes to having an idea specifies the composition of the storage infrastructure at a given time.

Storage profiles are rule-based to provide automatic data migration. A critical profile instructs the software to maintain the data in the most efficient resources regardless of the access frequency. A Normal profile migrate data blocks through the different levels of the infrastructure according to their use.

THIN PROVISIONING ENABLED BY DEFAULT

Thin provisioning is applied by default to save actually consumed disk space. SANsymphony-V dynamically allocates disk blocks as soon as they are required. The management of an unlimited number of snapshots for backups allows to return to any point in time in the past. Virtual disks in mirror are easy to create: select a pool on a server, enter a pool from another server and let the software do the rest. In addition, activation of the CDP function (Continuous Data Protection) creates a history log file for each virtual disk which can be reviewed at any time from activation.

SANsymphony features a particularly innovative cache function, since it uses the server memory - up to 1TB - to optimize writes to the disks. This function buffers write requests, improving performance of the same infrastructure, especially if it consists mainly of mechanical disks. Read requests and random writes are grouped to reduce latency, immediately releasing the targeted resources.

In conclusion, SANsymphony v10 makes storage defined extremely simple to understand and to deploy for companies looking for a way to manage and provision their heterogeneous storage. With this version of DataCore has dramatically improved the scalability and availability of data prioritization and the display of storage resources according to their performance makes optimizing the storage infrastructure of business readily accessible.

RAMON LAFLEUR

Understanding Hadoop-as-a-Service Offerings

ABOUT THE AUTHOR *Raymie Stata is CEO and founder of Altiscale, Inc. Raymie previously served Chief Technical Officer at Yahoo! where he played an instrumental role in algorithmic search, display advertising and cloud computing. He also helped set Yahoo's Open Source strategy and initiated its participation in the Apache Hadoop project.*

Hadoop has clearly become the leading platform for big data analytics today. But in spite of its immense promise, not all organizations are ready or capable of implementing and maintaining a successful Hadoop environment. As a result, the need for Hadoop coupled with the lack of expertise in managing large, parallel systems has resulted in a multitude Hadoop-as-a-Service (HaaS) providers. HaaS providers present an outstanding opportunity for overwhelmed data center admins that need to incorporate Hadoop but don't have the in-house resources or expertise to do so.

But what kind of HaaS provider do you need? The differences between each service offering are dramatic. HaaS providers offer a range of features and support, from basic access to Hadoop software and virtual machines, from preconfigured software in a "run it yourself" (RIY) environment to full service support options that include job monitoring and tuning support.

Any evaluation of HaaS should should take into account how well each of the services enables you to meet your business objectives while minimizing Hadoop and infrastructure management issues. Here are five criteria that help distinguish the variety of HaaS options.

HAAS SHOULD SATISFY NEEDS OF BOTH DATA SCIENTISTS AND DATA CENTER ADMINISTRATORS

Data scientists spend significant amounts of time manipulating data, integrating data sets and applying statistical analyses. These types of users typically desire a functionally rich and powerful environment. Ideally, data scientists should have the ability to run Hadoop YARN jobs through Hive, Pig, R, Mahout and other data science tools. Compute operations should be immediately available when the data scientist logs into the service to begin work. Delays in starting clusters and reloading data are inefficient and unnecessary. "Always on" Hadoop services avoid what can be frustrating delays that occur when data scientist must deploy a cluster and load data from non-HDFS data stores before starting work.

For systems administrators less is more. Their job typically entails a set of related management tasks. Management consoles should be streamlined to allow them to perform these tasks quickly and with a minimal number of steps. If the administrator must configure a set of parameters then they should be exposed while avoiding parameters that are managed by the HaaS provider. Similarly, low-level monitoring details should be left to the HaaS provider. The administration interface should simply report on the overall health and SLA-compliance of the service.

HAAS SHOULD STORE "DATA AT REST" IN HDFS

HDFS is the native format for storing data in Hadoop. When data is persisted in other formats it must be loaded into HDFS. Storing data persistently in HDFS avoids the delays and the cost of translating data from another format to HDFS.

After initial data loads, users should not have to manage data in storage systems that are not native to Hadoop or be required to move data into and out of HDFS as they do their work. HDFS is industry tested to provide cost effective, reliable storage at scale. It is optimized to work efficiently with MapReduce and Yarn-based applications, is well suited to interactive use by analysts and data scientists, and is compatible with Hadoop's growing ecosystem of third-party applications. HaaS solutions should offer "always on" HDFS so users can easily leverage these advantages.

HAAS SHOULD PROVIDE ELASTICITY

Elasticity should be a central consideration when evaluating HaaS providers. Another consideration when evaluating HaaS providers is the ease with which the service manages elastic demand. In particular, one should consider how transparently the service handles changing demands for compute and storage resources. For example, Hadoop jobs can generate interim results that may be temporarily stored. Does the HaaS transparently expand and contract storage without system administrator intervention? If not, Hadoop administrators may need to be on call to adjust storage parameters or risk delaying jobs.

Also consider how well the HaaS manages workloads. Environments that support both production jobs and ad hoc analysis by data scientists will experience a wide range of mixed workloads. How easily does the service adjust to these varying workloads? Can it effectively manage YARN capacity and related CPU capacity?

Ideally, the elastic expansion and contraction of resources should create minimal numbers of configuration and administration tasks.

HAAS SHOULD SUPPORT NON-STOP OPERATIONS

In production environments with fixed workloads, system administrators can tune operating systems and applications to optimize the processing of those workloads. They can achieve non-stop operations by crafting the best set of configuration parameters and monitoring key operation metrics to ensure jobs run as expected. Hadoop environments are rarely so predictable.

Big data environments are large, complex, distributed, and parallel systems. Such sys-

After initial data loads, users should not have to manage data in storage systems that are not native to Hadoop or be required to move data into and out of HDFS as they do their work. HDFS is industry tested to provide cost effective, reliable storage at scale.

System administrators rarely have too little to do or have deep expertise in every area of systems management. A HaaS that provides self-configuration allows system administrators to focus their time and efforts on tasks that cannot be easily automated.

tems present more challenging operating conditions than one finds in non-parallel applications including:

The need to restart failed subprocesses of a large job to avoid restarting the entire job

Jobs that starve for resources and finish late (or not at all), even when resources are available.

Deadlock, which occurs when one process must wait for a resource held by another process while the second process simultaneously waits for a resource held by the first process.

Non-stop Hadoop operations address these and other problems unique to the Hadoop environment. In-house and RIY environments are especially prone to problems maintaining non-stop operations because it requires deep Hadoop expertise and tooling.

HAAS SHOULD BE SELF-CONFIGURING

One of the advantages of using a HaaS is that it minimizes the need for Hadoop expertise. A HaaS should configure itself for optimal numbers and types of nodes. Data scientists well versed in statistics and machine learning may have deep knowledge about when to apply a particular statistical test or use a specific machine learning algorithm but may have no foundation for deciding on the configuration of a Hadoop cluster needed to run their workflows.

System administrators rarely have too little to do or have deep expertise in every area of systems management. A HaaS that provides self-configuration allows system administrators to focus their time and efforts on tasks that cannot be easily automated.

HaaS solutions should dynamically configure the optimal number and type of nodes and automatically determine tuning parameters based on the type of workload and storage required. These optimized environments dramatically reduce human error, reduce administration time, and provide results faster than customer-tuned environments.

LOOK BEFORE YOUR LEAP

Hadoop as a Service is a promising option to building and maintaining Hadoop clusters on premises. There is a wide range of features and support offered by existing HaaS providers. Carefully evaluate HaaS providers before committing to one so you have the best chance of selecting a HaaS that meets both your Hadoop management and data science support needs.

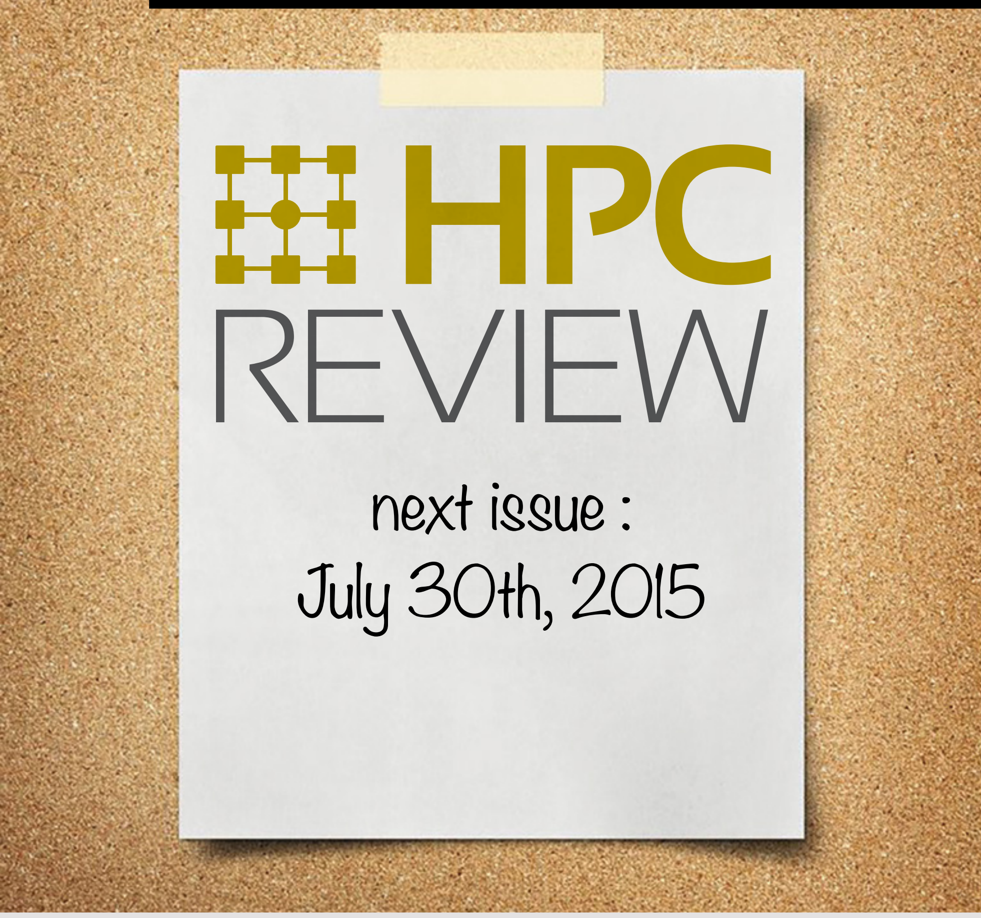

LAB REVIEW

Fujitsu Eternus CD 10 000
OCZ Z-Drive 4500
HP 3PAR 20800

VIEWPOINT

CERN infrastructure
Meet a technology giant

EN COUVERTURE

Flash Storage: the ultimate frontier

Breaking the $1 per gigabyte barrier

HPC Media
The HPC and Big Data Global Media

HPC Media
11, rue du Mont Valérien
F-92210
St-Cloud, France

CONTACTS
editorial@hpcreview.com
subscriptions@hpcreview.com
sales@hpcreview.com

PUBLISHER
Frédéric Milliot

EDITOR IN CHIEF
Joscelyn Flores

EXECUTIVE EDITOR
Ramon Lafleur

CONTRIBUTORS TO THIS ISSUE
Amaury de Cizancourt
Steve Conway
Roberto Donigni
Ramon Lafleur
Stacie Loving
Renuke Mendis
Marc Paris

ART DIRECTOR
Bertrand Grousset

SUBMISSIONS
We welcome submissions. Articles must be original and are subject to editing for style and clarity