

the Exascale Report™



The voice of the emerging exascale community

<http://TheExascaleReport.com>
December 2012

The Uber-Cloud Experiment: An Exascale Report Special Feature

By: John Barr, European Correspondent

In this feature interview with HPC consultant, Wolfgang Gentsch, John Barr, The Exascale Report's European Correspondent, reviews the Uber-Cloud experiment, a project to help researchers explore the end-to-end process for scientists and engineers, from technical challenges to social barriers, as they access remote HPC facilities on which to run their applications.

Introduction

Optimising High Performance Computing applications is all about understanding both the application and the target platform. Application developers, or those that port an application to a new supercomputer architecture, worry about memory bandwidth, data placement, cache behaviour and the floating point performance of the target platform in order to deliver the very best performance. Cloud Computing, on the other hand, is all about virtualisation, which hides details of the target architecture from the application. Cloud offers both technical and business model flexibility, which is good – but perhaps not for HPC applications. How does the Uber-Cloud project square this circle? I spoke with Wolfgang Gentsch to try to answer this – and many more – questions.

Wolfgang Gentsch

Wolfgang Gentsch founded GENIAS Software in 1990 which merged with Chord Systems to become Gridware in 1998 before being sold to

Sun Microsystems in 2000, where Wolfgang led the development of commercial opportunities for

Grid Computing. As an industry expert and independent consultant Wolfgang now chairs the ISC Cloud Conference for HPC & Big Data in the Cloud, as well as being a driving force behind the Uber-Cloud Experiment.



Wolfgang Gentsch
HPC Consultant

Background

The objective of the Uber-Cloud HPC Experiment is to explore the end-to-end process for scientists and engineers as they access remote HPC facilities on which to run their applications. The first round of projects ran from August to October 2012, with the second round kicking off at SC'12 in Salt Lake City, and running to February 2013. The project investigates the technical difficulties encountered, and also explores the social barriers, especially for SMEs, to access remote facilities in order to run their applications.

the Exascale Report™



The voice of the emerging exascale community

Although we have very quickly become used to working in the cloud (Google, mail, storage, etc.), there is still a strong tendency to hug servers, or applications, where the functions supported are crucial to a business. If a company's differentiation is aligned with the use of an application, there is often a feeling that running that application in the cloud will impact the value add it delivers. But is that feeling justified, or is it an unreasonable emotional response? And can the technical requirements of performance hungry HPC applications really be met by cloud resources? The bottom line is this. Is industry missing a trick by not making more use of cloud-based HPC facilities to deliver both additional, flexible, cost-effective capacity and new business models?

The motivation for the project came from a series of conversations between Wolfgang Gentzsch and Burak Yenier who wanted to better understand how real the perceived problems are that constrain running HPC in the cloud. These problems include concern about privacy and security, unpredictable costs, ease of use, software licensing and application performance. The experiment was planned to help address these concerns. The project has no funding and is backed by no commercial or governmental organisation. It is a labour of love, and an opportunity to build a community that may change the way that high performance computing delivers value to businesses.

The Experiment

The goal of the Experiment is to form a community to explore the challenges and benefits of running HPC applications in the cloud, to study the end-to-end process, learn what works (and what doesn't), and to document the findings to help the next group of potential participants.

More than 260 organisations and individuals, from 26 countries, are involved in the Uber-Cloud Experiment, with more than 150 of these directly participating in an experiment. An experiment consists of a team of four entities: SME end user, resource provider, software provider and HPC expert. The makeup of this team is crucial to the process, as SMEs that try to exploit the benefits of the cloud without having the buy-in of the resource or software provider can feel that they are banging their head against a brick wall, and the addition of an independent HPC expert is the final, important, piece of the jigsaw. The goal of the experiment is not only to run a number of applications, but also to document a range of potential solutions to perceived roadblocks to delivering high performance computing as a service.

Resource providers include Universities, mainstream Cloud providers, HPC system vendors and SaaS providers. Software providers support manufacturing simulation, CAD, CAE, electronics, building design, data management and visualization. Expertise is sourced from a mix of small and large companies, including specialist consultancies, hardware and software vendors and academic HPC groups. Users come from academic research, as well as the computer, construction, energy exploration, minerals and cement, and shipbuilding industries.

Why do these companies participate in the Uber-Cloud experiment?

(Continued on next page)

the Exascale Report™



The voice of the emerging exascale community

Why do these companies participate in the Uber-Cloud experiment?

Group	Benefits
Users	The process is fully guided, enabling an inexperienced user to gain experience on the job. A range of resources are available, and the one that best suits a user's needs can be selected. It enables users to do things they couldn't do before. In some areas the initial benefits of a new approach are easy to access – the “low hanging fruit”. This is not always the case for users making their first foray into HPC, and help is required to reach the tasty, but high hanging fruit.
Software Providers	They can explore new business models without the commercial pressures that accompany rolling out a new service. Many ISVs have a large customer base that they don't necessarily expect to expand through this activity, but what they can do is add new types of services that enable their customers to run more jobs, or bigger jobs, or explore more parameters – all of which will require the consumption of additional licenses.
Resource providers	Most of these are already delivering resources using a cloud model, so they don't have to invest a great amount of effort, but if Uber-Cloud is successful they will have developed a new business opportunity, especially if the take-up with SMEs is strong
HPC Experts	Uber-Cloud gives domain and HPC experts the chance to take their skills and experience to the next level. Perhaps they are CAE experts, or cloud experts, but not both. This gives them the chance to share their skill set, while broadening it at the same time.

So it's a win-win-win-win situation.

The initial round of experiments was focussed on manufacturing, but the second phase adds computational biology, a segment with very different characteristics. As the community expands the experiments could run forever, but sponsorship is needed to augment the many hours of unfunded effort invested by Wolfgang and Burak.

Conclusions

Delivering easy, cost-effective, cloud-based access to high performance computing facilities for SMEs who might not otherwise be able to afford a dedicated HPC capability, provides not only cost savings, but great flexibility – and opens up new business models that exploit HPC, helping SMEs to build better quality products and services.

How does Uber-Cloud address the issue we raised at the start – namely that optimised HPC applications often run poorly on virtualised cloud infrastructures? For a start – the point of Uber-Cloud was to explore if problems like these are real, or if they are hype. Their report will answer this question – and many others – for would be HPC in the cloud users. Delivering applications in the cloud as an optimised Software as a Service that hides the complexity of the architecture from the user is one answer. The software provider is then responsible for delivering performance, and the HPC expert joins the dots if anything is incomplete or unclear.

Uber-Cloud is quite remarkable. It has no funding and no institutional backing, yet it has delivered more real value to its 260 participants in less than a year than many well funded research projects (on both sides of the Atlantic) do in three years.

the Exascale Report™



The voice of the emerging exascale community

I asked Wolfgang about the most important lesson that he had learned working on the Uber-Cloud Experiment. His response was: "Before doing the Uber-Cloud Experiment, I thought I knew quite a bit about HPC in the Cloud. Now, after 3 months, at the end of the first round, I am amazed about what I have learned from just watching our 25 teams working hard and hands-on to bring HPC to the Cloud." So the most important issue is that even experienced HPC professionals still have much to learn about HPC in the cloud - the results from Uber-Cloud will answer a lot of questions that are being asked by users, ISVs, resource providers and HPC experts.

###

Talk to the Emerging Exascale Community

Don't miss out on your opportunity to talk to the emerging Exascale community by submitting a sponsored editorial opinion article.

Talk to the community in your own words. Address the issues you feel most strongly about. Raise your concerns or ideas, or present a challenge to policy makers, funding sources, vendors or end users.

For more information:

<http://theexascalereport.com/pdf/advertorial.pdf>

John Barr is the European correspondent for The Exascale Report.



Barr is a widely recognized independent industry analyst, formerly Research Director of HPC at the 451 Group, who brings 30 years experience in the HPC industry to the publication.

Barr covers the European perspective on exascale issues, including coverage of users, vendors, and European Commission funded research programmes.

<<<<< >>>>>

For subscription information, go to:

<http://theexascalereport.com/user/register>

The Exascale Report

The Voice of the Emerging Exascale Community.

© 2012, The Exascale Report™