# Cloud Computing Roadblocks for Engineers and Scientists and their Service Providers

**Protocol of a Discussion on LinkedIn**

Wolfgang Gentzsch, The UberCloud, February 08, 2015

This report is a collection of discussions on 11 LinkedIn Groups which took place in November and December last year. The aim was to explore the current level of acceptance in our community about engineering and scientific application simulations in the cloud. This report is mainly meant for cloud professionals, for example as resource, software, or expertise providers. For those interested just in the essential results we recommend the short summary article here: http://www.theubercloud.com/workstations-servers-clouds-comparing-apples-apples/

## Introduction

To increase innovation and competitiveness, today, we are recognizing in governments and in our community an increasing interest in creating new high-quality jobs especially in manufacturing, and here especially in digital manufacturing which focuses on the design and development of new products and services for the benefit of our whole society. To accomplish this successfully our engineers and scientists are using powerful computers and sophisticated simulation software as their major tools; computers like workstations and high performance computers, in-house and in public clouds; and simulation software from ANSYS, Autodesk, CD-adapco, ESI, Matlab, and Simulia, for example. Especially the hardware tools – workstations, servers, and clouds – come with great benefits, but also with real challenges to acquire, access, use, and maintain them.

Over the past two years we at the UberCloud have consistently tried to contribute to a better understanding of the benefits and challenges of these different tools, by helping our community of engineers and scientists to perform real-life experiments (more than 160 so far), by publishing about 100 articles and monthly newsletters, and by organizing webinars.

In November and December 2014 we kicked off a LinkedIn initiative to find out the current level of acceptance from engineers and scientists to move their application simulations into the cloud. We posted the following simple question:

**What if ALL Cloud Computing roadblocks were removed, today? Would you move part of your CAE simulations to the Cloud?**

If we were able to remove ALL the roadblocks (security, licensing, data transfer, cloud access and control, cloud expertise, etc.), would you then incorporate on-demand cloud bursting into your daily CAE projects in cases where your desktop workstation can't handle the overload?

We posted this question in the following 27 LinkedIn Groups, and received 71 comments from 48 experts posted in 11 groups (market in bold):

> **HPC Cloud, Abaqus,** ANSYS**, Autodesk,** Bio-IT World**,** Bioinformatics Computing**,** Bioinformatics Geeks, CAD/CAM/CAE**,** CAE Professionals**,** CD-adapco**, Computational Biology, CFD Computational Fluid Dynamics,** Computer Aided Engineering**,** DM Digital Manufacturing**,** Digital Manufacturing**, Finite Element Analysis,** Finite Element Analysis (FEA)**,** High Performance & Super Computing, HPC Executives**,** HPC Services for SMEs**, Innovative Use of HPC,** Matlab Users**,** ME Mechanical Engineers Network**, Molecular Modeling in Life Sciences, New Trends in CAE, Simulation Driven Engineering,** and **Simulia.**

We are presenting the comments and our corresponding responses as they were made, unchanged, except correcting a few typos. And because these discussions happened in parallel in 11 LinkedIn Groups, many of the topics and themes overlap, and the challenges and concerns raised in the different groups are similar and redundant. The major concerns and challenges mentioned in these 71 contributions are (and how often they are mentioned):

23 x Uncertain costs of cloud computing
22 x Slow data transfer and bandwidth
16 x Security
14 x Cloud licensing
14 x Cloud access
4 x Losing control
4 x Intellectual property and corporate assets

We also summarized and published comments and responses in a short article discussing and comparing 14 major topics of concern about in-house servers as well as technical computing clouds, and how much we believe the cloud roadblocks are (can be) reduced or even removed by recent technology, today and in the near future. These topics are: **system or service procurement, budget/cost, system operations, system and services maintenance, business flexibility and agility, reliability of services, average resource utilization, security, new technologies, data transfer and bandwidth, full control over your assets, software licensing, access to computing and storage resources, and wait time to computational results.** In the contributions below, we have highlighted these key cloud challenges and other important pro/con arguments in **bold**.

One technology stands our here: Linux containers, especially containers from a company called Docker, and based on Docker the UberCloud Containers especially developed for containerizing engineering and scientific applications. Please see the summary article here:
http://www.theubercloud.com/workstations-servers-clouds-comparing-apples-apples/

| Feature / Functionality | In-house HPC Server* | Remote HPC Cloud** |
|---|---|---|
| **Procurement** | can take months | depends on your company |
| **Budget** | CAPEX, inflexible | OPEX, flexible |
| **Operations, maintenance** | complex | none |
| **Flexibility, agility** | low | high |
| **Reliability** | single point of failure | redundant |
| **Average utilization** | around 20% | up to almost 100% |
| **Security** | medium | high |
| **Technology** | ages quickly | frequently updated |
| **Data transfer speed** | high | depends on your last mile |
| **Full control over your assets** | high | medium to high** |
| **Software licensing** | expensive (e.g. annual) | pay per use |
| **Access** | mastering many scripts | seamless |
| **Wait time** | depends on job queue | none |

*On average. **Depends on cloud provider.

**Acknowledgement**

For their valuable contributions to this discussion we'd like to thank our expert colleagues (in alphabetical order): Ardalan Amiri, Sadegh Ben Babaii, Peter Bartholomew, Lawrence Bertoldi, Ziad Boutanios, David Byrne, John Cary, Vladimir Chupakhin, Chris DeGroot, Alvaro Fernandez, John Giles, Merle Giles, Richard Gosling, Ben Grigg, Martin Hegedus, James A. Holst, Michael Jandron, Sharan Kalwani, William Kilpatrick, Pradyumna Kulkarni, Zhengjiang (ZJ) Li, Greg Locock, Pedro Marcal, Marek Michalewicz, Hammad Mubeen, Samuel Murgie, Dennis Nagy, Pierre Malouf, Pedro Marcal, Tom Miller, Oleg Moskvin, Theodore Omtzigt, Mohan P, Chris Pegg, Mihai Prună, Leo Reiter, Harry Schaeffer, James Shaw, Francisc Simon, Ebbe Smith, Kartik Srinivas, Robert Strong, Adam Towse, Herman Van Niekerk, Joe Walsh, Jeff Waters, and Herman Westra. To respect their privacy we anonymized their comments in the following (referred to as XYZ).

# Discussions

**Comment 1:** Good question, Wolfgang -- let's see who answers.

**Response 1:** I posted this question into 27 LinkedIn Groups, CAE and HPC oriented and got already many excellent comments. Now on my way from Munich to the supercomputing conference SC'14 in New Orleans, there's a lot of good time to start with my responses. And I decided to collect all comments and my responses in a report available for download on our website soon after.

**Comment 2:** The common perception is that human beings are rational animals. So ideally if all roadblocks and potential objections are overcome then rationality demands that we move part if not all of our CAE simulation to the cloud. But, humans are not rational animals. We are emotional beings. It is just that some realize this while others do not accept this fact. **Our emotions and inherent biases will always cloud our judgments** and our behavior i.e. decisions.

Now, having said this - I would venture to answer your question rhetorically - Can a person who is not comfortable with the idea of the "cloud" really embrace it as an acceptable way of working?

Consider the **demographic of engineers** in general - as of 2008 the median age of scientists and engineers in the US was 41 years and **the proportion over 50 was 27%.** This is the demographic that (for the most part) is **making the buying decisions** for the younger engineers. These are very smart people but they did not grow up with the notion of data flowing freely into and outside their realms of control and this introduces the emotional bias.

I believe we see evidence of this from the companies still **engaging in build and break instead of simulate** in their product development process! And this, after simulation having been around and gathering momentum over the last 30+ years, aided by ubiquitous computing technology.

So to conclude - I will say that advances in technology and business models are definitely needed but the shift will only occur **when decision makers are emotionally comfortable with the way the cloud works**. In my opinion this means it will take another 5-10 years to achieve the demand momentum. While technology is making huge progress it is still ways away from being truly seamless and easily available and I believe your contention that technical and business model roadblocks will be removed will be proven soon enough.

So while users and buyers make the journey to adopting the cloud, work needs to continue on **evaluating and exploring business models and technology stacks** to make simulation on the cloud a reality.

Wolfgang, I realize that I may have gone off on a tangent to the question you posed but I believe that this (**human irrationality**) is an important aspect that we engineers tend to overlook and get frustrated when smart engineers refuse to accept this obviously elegant solution to their own challenges.

**Response 2:** What a great observation. And that's (among others) what I meant in my answer to XYZ above, what you nicely call the **'emotional bias'**. We know how much we rely on this precious attitude in our lives, so I won't accept this to be a negative. Removing a roadblock therefore can only mean that we have to work hard on removing the real obstacles such that the emotional bias disappears over time.

Concerning your demographics about the median age of an engineer, good news is that it is a median. In smaller, younger companies engineers tend to be younger, more open to new developments, more innovative (on average). Therefore, I believe that **change will first come from smaller companies**. And then it will influence the majority of engineers in the mid-term.

Also, we often tend to generalize things too much; when we think of Cloud we tend to see ONE Cloud, ONE application, ONE user, ONE requirement, ONE scenario. But in reality, this space is n-dimensional with n being a big number. On the one end there are very natural and easy **scenarios which are ready for the cloud today**; and on the other end there are very challenging and **complex scenarios which will never be well-suited for the cloud**. You know what? We humans tend to pick the last extreme when we face a new situation, a disruptive technology, to justify a decision based on emotional bias.

How can we accelerate the process of cloud acceptance? I think just **by developing real-life (cloud) examples**. That's why we at the UberCloud first started with what we call 'experiments', i.e. small projects where we help engineers to **explore the end-to-end process of moving and running their application in the cloud**, with all the benefits and challenges; and publishing an objective case study at the end. That way, we did **162 engineering cloud experiments so far, and published 42 case studies,** hoping that more and more case studies will convince more and more engineers. If you haven't seen the case studies, you can download them here: https://www.theubercloud.com/ubercloud-compendium-2014/

**Comment 3:** Wolfgang, Let me insert here a few of the slides I've presented a couple of days ago:

1. In what space does the CLOUD live?
REAL Computer Resources
REAL Datacenters
REAL Energy Requirements
REAL Data Storage
REAL Data Transmission
REAL Data Processing

2. CLOUD is
Business Model
Ownership Model
Charging Model
Provisioning Model

3. But, it is all possible due to … Alan Turing. Software <–> Hardware . . . . . (Something) <– as –> (Something).

Now, in order to remove Cloud roadblock, we need to work on removing real roadblocks related to real resources.

We at A*CRC in Singapore are collaborating with the National Computing Infrastructure at ANU in Canberra and are on our way to remove two of the roadblocks:

1. Provisioning of large instances of HPC Cloud *with* InfiniBand interconnect (this is excellent NCI work). [This is about intra-cluster **high bandwidth and low latency**]

2. Very fast access and data transfer over large distances using InfiniBand (including trans-continental InfiniBand), allowing easy and efficient access and data transfer between virtual HPC resources and a remote user (A*CRC work + Obsidian Strategic Technology). [This is about **long-distance big data transfer**]

Then, this might be one step closer to your vision!

**Response 3:** One by one to your comments: I fully agree with you: the Cloud too lives in a REAL space; it's made of real components, and in this sense it's like your real desktop computer or your server, a piece of hardware and software connected via network, just a few more miles away. That's why I think the terminology 'Cloud' is a somewhat misleading analogy. It is even unfortunate, because with 'Clouds' we often intuitively connect things with it which have a negative impact on our lives (rain obviously when we want to relax at the beach...).

To 2. First and foremost, **Cloud is a Service Delivery Model**, and as such it comes hand in hand with the other models you are listing.

The Cloud is not (primarily) technology as some people tend to think, but it uses a lot of existing technology to function properly (like VMs, Linux Containers, etc.), and to help remove some of the roadblocks, like the ones you mention: using InfiniBand interconnect for building an HPC Cloud, and using transcontinental Infiniband to improve GB data transfers. However, as we know, not all roadblocks can be removed by technology though, especially not those which reside in our heads and minds…

**Comment 4:** The **last mile** is a huge hurdle. In the US, cable and DSL are still the two most common ways to get on the Internet. If you leave the cities and go to a smaller town, it is by no means unusual for the nearest DSL junction box to be 10 miles away, dropping your effective **bandwidth** by a factor of 2 or more. And it is also possible that you won't have cable connection at all.

So, for all the great 21st century infrastructure upstream of the user, the user himself has an infrastructure that is more reminiscent of the late 20th century.

The user's experience of the cloud is therefore that it is slower and less reliable than the on-premise solution you are trying to get him to replace. This means selling the cloud to these clients implies selling the local telco infrastructure provider on the wisdom of expanding their service to your prospect. I've experienced this first-hand.

The point is - the biggest *roadblock* IMO is not the technology of the cloud, but the **technology to bring the cloud to the end user**. This is lagging in our country, and it informs the attitude of many of the prospects I've dealt with, not all of whose installations are in major metropolitan areas. Let's include this as one of the roadblocks that got magically removed in this thought experiment.

**Response 4:** Indeed, we can't remove the last mile problem in a sudden. That's why we try to soften it, by-pass it. Engineers in our experiments have used several ways to get their GBs of resulting data from the Cloud back onto their workstation: obviously **data compression**, I have seen up to a factor of 8. Then **remote visualization**, and that's a big one: often an engineer just needs to have a look at an intermediate result, checking accuracy, convergence, quality of a result. Then changing a few parameters and re-run the job; several times, until the final solution looks good enough. That's the result you want to have back and archive. So you either send the entire dataset back to your place which can take a few hours, **or you fedex it**. Sure, not an elegant solution, but reliable, secure, and fast enough. At the end of the day, what's the difference? You'll always have your final dataset next morning. At least it's a reasonable work-around, for now.

BTW, there is new technology on the market, which accelerates **interactive remote visualization** with GPUs, e.g. NICE DCV. And there is **data streaming** from the gaming developers around the corner . . .

**Comment 5:** Probably not. If I needed a huge computer to run my analysis, I probably haven't thought cleverly enough about how I posed the problem. I also don't believe the roadblocks you stated could be removed to the satisfaction of most analysts, particularly the issue of security. Allowing licensing on demand is a better solution.

**Response 5:** I agree that many problems we face today still can be softened or even removed when rethinking the approach or the underlying algorithm. A fresh look from a different angle can often help.

But, and I guess you agree, many more problems are simply big by their nature. Like an engineer who has to do a digital crash simulation with a complex car model, 250 million elements, using a commercial code like Pamcrash. This problem won't fit anymore into the memory of his workstation, and if so it would run a whole week. But the engineer might have to do 150 runs to find the best geometries and materials. **In the Cloud,** you can do each run on e.g. 4-times more cores than on your workstation, reducing one run down to one day. And you could do all the **150 runs in parallel, still one day.** There is a nice case study from Mark Lobo in Desktop Engineering these days, about virtual testing of severe service control valve, who experienced exactly that, http://www.deskeng.com/de/testing-simulation-cloud/.

**Security**: Every cloud provider knows that security is an issue, especially for his customers. Without providing sufficient security a cloud provider would not even exist, no customer would come. Therefore, every cloud provider has at least one security expert, usually very expensive people an SME (small and medium enterprise) can't afford. Therefore, **cloud infrastructure tends to be at least as secure as the SME's,** if not better. So if an SME is working on some secrete stuff internally, at least NSA knows about it… and certainly some foreign countries… A great and un-biased article about security in the cloud is here: http://www.businessspectator.com.au/article/2015/1/12/technology/clearing-skies-cloud-getting-safer/.

And **licensing on demand** can solve some problems. But when you have an average workload load of X licenses, but sometimes needs X+10, you would still need the hardware to run your X+10 jobs then. But what SME would easily buy additional hardware just for a few peak demand times? Especially for these cases Clouds are really ideal. I have written a little cost article for this situation which you can download from http://www.theubercloud.com/cost/ which shows that you can save a lot of money with this '**hybrid'** solution (in-house **and** cloud).

**Comment 6:** This is interesting. If this included the ability to use thousands of CPUs then yes, it might rival HPC clusters. I think it's silly to say we don't need the cpu power especially when we are just at the cusp of highly coupled multi-physics simulations (e.g. FSI, LES CFD) and even something as simple as resolving high frequency waves (ultrasonics, photonics). **Massively parallel algorithms are of huge interest** and cloud computing might be an interesting way to make it happen. Of course **bandwidth** is an issue today and **security** would obviously be a nogo for most.

**Response 6:** The Cloud is told to provide **infinite amount of resources** (at least theoretically). So, Yes. Still the algorithms have to be **highly parallel** to be able to scale. Luckily, some do.

And yes, **bandwidth** IS an issue. Engineers in our UberCloud experiments have used several ways to get their GBs of resulting data back from the Cloud into their workstation: obviously **data compression**, I have seen up to a factor of 8. Then **remote visualization**: often an engineer just needs to have a look at an intermediate result, checking accuracy, convergence, quality of a result. Then changing a few parameters and re-run the job; several times, until the final solution looks good enough. And that's the result you want to have back and archive. So you either send the entire dataset back to your place which can take a few hours, or you **FedEx** it. Sure, not an elegant solution, but reliable, secure, and fast enough. At the end of the day, what's the difference? You'll always have your final dataset next morning. At least it's a reasonable **work-around**, for now.

BTW, there is new technology on the market, which accelerates **interactive remote visualization** with GPUs, e.g. NICE DCV. And there is **data streaming** from the gaming community around the corner . . . Concerning your remark on security, please see my comment to XYZ, above.

**Comment 7:** **Security** is very important when we are dealing with ITAR controlled projects or other confidential work.

So my answer would be "No" without security. Like mentioned before, on demand computing power with all the security is preferable. There is of course a lot of demand for faster computing power. But many small companies can't justify procuring a cluster with 1000 cpus as the usage may not always be close to 100%. So renting computing power as needed would be a better option, with all the security.

**Response 7:** Concerning your remark on security, absolutely! And please see my security comment to XYZ, above.

To support your case about justifying a 1000 CPU cluster, I did an example **cost analysis** for a $70K cluster which according to IDC costs $1 million of Total Cost of Ownership (TCO) over three years, and I compared this to buying a smaller cluster for your average workload and bursting into the cloud for peak demands. You can download the little paper from https://www.TheUberCloud.com/cost/ which shows that you can save a lot of money with this 'hybrid' solution. BTW, say after one year, you still will get the latest and greatest hardware from the cloud provider, while your in-house server is getting older and older…

**Comment 8:** I think this is very interesting indeed. Given the possibility to rent computing power on bleeding-edge servers and on-demand licenses - with no overhead for server management, I think this would benefit medium-to-smaller companies - as they would be able to increase capacity without having the overhead of maintenance or (hopefully) annual **cost**.

On **security**, I do believe this is very solvable. There are good existing solutions out there that provide **high-speed encrypted connections**. I do recognize that such a service will be a more lucrative target for hackers to obtain IP and general know-how, but having a company that works dedicated with this kind of expertise and continuous monitoring of traffic might be more secure in the end (especially for medium-to-smaller companies with limited IT resources). It will however be hard to convince the average-Joe.

**Response 8:** Thanks, very right. Especially **small and medium size companies** will benefit most, because they cannot afford to buy their own systems. And they don't have to, anymore. But also the big companies will benefit; not directly, because of their complex and bureaucratic compliance regulations; but indirectly: these big guys often have a huge network of small and medium size **suppliers** which with using cloud resources are now able to provide parts with **better quality, in shorter time, to reasonable cost**. Improved parts will improve the whole, easily.

**Comment 9:** Interesting to me. The question is at what magnitude and in which context we are considering it for CAE supercomputing.

a) How do you protect me the small guy in CAE against big corporations taking over all your computing power (cpus) to free up their facilities from time consuming tasks during the day? My jobs will be at the end of an endless queue of big jobs from the big guys.

b) What about **licenses**, should I have my licenses or your CAE pricing includes the cost of the software licenses?

c) Your offer of much more cpus than I can afford onsite, plus prospect of as many licenses that I only pay while they are actually used (not per day or month, etc.) I am tempted to run jobs greater than I could handle. The painful upload and download CAE mega file (**data transfer, network bandwidth)** would be my super great road block. CAE is very demanding and interactive work and painful up and down transfers make it very inefficient despite the quick and economic service you could offer.

d) My CAE platform of choice will be offered by you or should I train myself to your choice of CAE platform and therefore become vulnerable to price hikes from you because I cannot **migrate to another cloud** supercomputing service with more economic pricing as I have to invest and learn all over again.

e) Last but not the least, you become dependent to a single big corporation customer using all your cpus as they can afford to pay for it (and all points above doesn't really matter to them anyway) but every year you should reduce your price to keep your single customer or stop operating.
I could write a paper on these points and the list is by no mean complete. I don't see them as hurdles but as **economic realities**.

**Response 9:** To a) First we know the big guys will come late to the (cloud) game, because of their very complex compliance rules. So for now all the cloud resources are yours! Later, when the big guys join, there will be more business, and therefore more cloud providers with more resources (demand - supply). And, **you have the choice**, if one provider doesn't serve you well, move on to the next provider. Now you might argue how to move your stuff from one to another provider with all these different providers not having one standard (for now). But there is new technology, **containers,** which lets you move your stuff around seamlessly; **no vendor lock-in anymore**. See the article in Scientific Computing World http://www.scientific-computing.com/features/feature.php?feature_id=322.

To b) **Licensing** in the Cloud: there are several **license models** today: the first two ones you mention: **BYOL, Bring Your Own License**: say your CAE software is already sitting in the cloud, in a container. You access the cloud (via a link and password), and you are asked for a license key for temporary use. You go to the ISV's website, buy 25 tokens for X CPU core hours of software use, provide it to the cloud container, submit your initial data set, and you are ready to go.

Or you buy an application **package** (see e.g. OpenFOAM https://www.TheUberCloud.com/store) for $199 for 24 hours on 16 parallel cores) with **EVERYTHING included** (bundled hardware, software, license (if needed), container, tech support). There are other models, such a **subscription** per week, month, or year. Or the ISV has its own cloud, like Altair or Autodesk, and everything then comes from them, including the software license.

To c) Data transfer: or **bandwidth**; please see my comment above to XYZ about this important topic.

To d) Great question about **vendor lock-in**. Working directly with one single cloud vendor has indeed the danger of depending on this and only this one provider. We have clearly seen this two years ago in our early UberCloud experiments: many teams needed weeks to get access to a certain cloud provider, get their software installed, get their initial dataset in the cloud, run the job(s), and get the results back. And a vendor-independent CAE Cloud Marketplace with dozens of providers doesn't solve this problem. Every Cloud provider is different, so lock-in can happen. But with the new **container** technology (based on **Docker**) which I mentioned in a) the software is already waiting for you in a container (identical application containers can sit in different clouds), and as soon as your initial dataset arrives the job can run; same process with every provider, fully standardized. Not happy with your provider? Go to the marketplace, select another provider which offers the same software, and move on.

To e) This point is important, but is independent of Cloud. One way out of this is, never depend on just one big customer; easier said than done, I know. The other way out has indeed to do again with Cloud: with a constant quality of your product or service which you deliver in a certain time, your big guy can slowly squeeze all blood out of your veins. BUT, you can counter with keeping your price at least constant (if not increase it) by at the same time improving quality and shortening time. The Cloud in a sudden provides you with two additional dimensions: with more and finer simulations you can **improve the quality**; and with doing more simulations in parallel you can **shorten time to delivery**.

**Comment 10:** A lot of good questions and points here! Let me share with you what I think - as a possible future customer for one of these services.

To a) If it's like that, you don't get what you pay for. For such as service, experts will monitor server loading to not let scenarios like that happen. Most likely you'll end up to pay for a certain quota.

To b) good question indeed. I personally would love on demand licenses, but I think a license scheme similar to what we see today is most likely.

To c) Another good point! For download - it is technically possible to **stream** the huge odb-files back to you as they are written - but for sure, having a **high-speed internet connection** would be essential. Maybe this service is also something you do for those 'enormous' jobs you can't do on-site....?

To d) What would you accept? Personally I'm only interested in paying for raw power and the expertise of maintenance - I'll choose my solver.

To e) Well, don't know. I guess you'll end up paying for a certain quota of power - and this kind of services is possible to do without depending on one big customer. You'll expand your server/hpc park as you get customers?

What would you accept as a user, what is your utopia?

**Response 10:** I didn't see your comments to the previous contribution when I answered. So I agree with your usage scenario in a). Workload cluster management software like Grid Engine can easily handle what we call **fair share scheduling**. Maybe the priorities (which you can impose) are not right. But anyway, Cloud can also fix that.

On the **licensing**, we see many different scenarios: **subscription** based (Autodesk), power on demand **tokens** (CD-adapco), **pay-per-use** (eCADFEM), and **no license** required with free software like OpenFOAM.

**Data transfer**, yes, streaming will become a great solution, quite a few data streaming companies in the gaming industry are working on this.

Choosing your own solver is not always an option. Many engineers have used one and the same solver for years, they are now great experts just with this one software which often is not easy to handle. And for starting to learn another solver there is often no time in-between all the different projects and day jobs. So **having the same solver on different Cloud platforms** can be a great alternative.

**Comment 11:** I'm a consultant and travel around a lot. I do not have reliable access to the internet. Often times I don't even have cell service. Even out of my home base office the **connect speed is too slow to use with CFD file sizes** and in bad weather it goes down a lot. I don't see how you can solve that problem. Even if you could solve that problem I would not want all of my business to be dependent on one weak link. I have more than one computer now and could transfer the license to another machine if one went down.

Cloud based computing is what we used to do 30-yrs ago and then we got powerful computers for our desk and life became much better. Why would I give that up?

**Response 11:** Today we would call the 'cloud based computing' 30 years ago 'private cloud', the mainframe in your organization's data center. Still today big organizations have these 'private clouds', and they are reasonable and economic as long as you have a high average utilization (at least 70%). But industry average today is between 10% and 20%, a lot of waste of expensive (un-used) computing time.

An un-reliable Internet access is indeed a big roadblock. Often however the Internet access can be improved if necessary, and usually this comes with a price. But this is a cost/benefit calculation. With an excellent Internet access Cloud is no longer out of reach, and even with a lot of travelling you can then transfer the larger files to your office, and do the checking of your cloud jobs and intermediate results (via remote visualization) on the road.

**Comment 12:** Much the same as XYZ above, I travel quite a bit but I also live and **work in a rural area where internet speeds are erratic** at best and can drop from 2mb to around 200kb on rare occasion, it's not really an efficient way to work for my small business. My computer is sufficiently fast for the simulation work I carry out and as correctly stated above why have the weak link, I've spent thousands on the software and licensing as everyone else I don't imagine many who invest in things like simulation mechanical and CFD are going to have poor computers to run the software on. In short no.

But then again I like to own my software also and receive my cd/usb stick each year so the idea of logging in to use my investment does not appeal much either, this maybe an antiquated way of looking at things but never the less it's my preference, though I do understand it's the way the world is moving though for me in a rural area the superfast internet highway is some time in the future.

**Response 12:** Same answer about an un-reliable Internet as for XYZ, above. Not much can be done with an erratic Internet speed. So for now, you are not in an ideal situation to access and use cloud computing.

And why should you use the cloud that your own computer is sufficiently fast? **Cloud is for those who are in need of even better simulation quality (100 million elements instead of 10 million; adding multi-physics; LES; doing 150 runs in parallel to find the best material, etc.), or**

**faster time to delivery of results (with doing all 150 runs in parallel), or saving costs (using cloud instead of buying your own supercomputer).**

**Comment 13:** Daily? No. If the need was that great, then **owning the hardware** is cheaper than cloud computing. Not to mention, it is a one-time capital expense, rather than a recurring expense, and so there are **financial benefits** to purchasing it. I am honestly trying to incorporate it into my workflow (for the occasional CFD analysis which exceeds my resources) and I have just not gotten it to work as smoothly as I had hoped. The **cloud scene is very dynamic** right now, so I am thinking about waiting a year or two until it has reached a steady state. Keep in mind, I am an independent consultant as well, so my opinion is based on similar assumptions as XYZ and ZYX above.

**Response 13:** **Owning the hardware** is cheaper than doing the same work in the cloud **ONLY IF your hardware is utilize more than 80%,** on average. **Otherwise, Cloud is cheaper**; and much cheaper if you consider the industry average which is between 10% and 20% utilization. I did the math, it is here: https://www.TheUberCloud.com/cost/. And the best price/performance can be achieved with a **hybrid cloud usage** model: buying a server which can cover the (e.g.) 20%, and using the cloud for your peak demands (beyond the 20%).

**Purchasing** a bigger computer usually needs **long procurement** cycles, has to be **justified**, **approved** from upper management, and competes with other projects in-house, while **Cloud** usage is operational expenditure which is much **more flexible**, agile.

BTW, with your situation similar to XYZ and ZYX (above) it might make good sense to wait for another year, until technology improved especially your boundary conditions (limitations).

**Comment 14:** I agree with those above, my business involves mobilising structures onto construction ships for the offshore industry, as well as designing them in the first place. I use my FE in the office but also often onsite (a ship) where often wifi doesn't work, besides I like to be independent of my clients system.

Having dependable, stable software that will switch on and work quickly each time is far more important

**Response 14:** Absolutely agree. With changing environment conditions the access to remote resources is not fun. Only if you have a real need for more compute power, be it for many more simulations all at a time, foe many more elements, etc. which cannot easily be handled by your existing computer, then you will have to think about alternatives. Those might not work on a ship, but they will work from your office. And dependable, stable software is a prerequisite for a functioning cloud, today.

**Comment 15:** I don't want to sound like a dinosaur. I tried using 360 CFD **a couple of years ago** when it first came out for about 3 months. The potential advantage I saw was the ability to run **multiple load cases in parallel** possibly saving me wall clock time. First problem is the UI was not setup to make this very easy. That problem I could see you solving.

The next problem I had was the **huge amount of CFD data** that needed to be transferred at slow **internet speeds**. This was a big problem with restarts where the big data files need to go both ways, up and down. You may solve this problem putting everything on the cloud and using a web based UI, however, the virtual PC programs I've tried seem very slow and clunky.

The third problem I had was that **jobs would fail and I had no idea why**. Sometimes I would not know they failed for hours or more. The exact same jobs would run fine on my system. I suppose this problem could be solved as well.

I found that when jobs did run without issue they were not any faster than running them on my system, plus I had to deal with the upload and download times. Fast computers are cheap to buy these days.

The last problem I had was **my own internet reliability**. I'm not always online. I don't think you can solve that one. It's out of my control as well.

In the end, over those 3 months I worked on three large projects running them in parallel on my in house machine and I was not able to get any of them completed on the cloud for one or more of the reasons above.

I can see the cloud has potential as an available option but I would not want to give up the ability to run local and **have local control of my core business**.

Now for some dinosaur talk. Thirty years ago I could create my NASTRAN models in my home office but the 'cloud' computer I needed to run on was 1000 miles away at Boeing. All I had at home was a 300bps dial up modem. I had to travel 20 miles to a data center where there was a high speed (9600bps) data transfer modem to submit my jobs. When I got a copy of Algor FEA to run on my PC, life changed and the data center went out of business. The way things are going, I may be spending my time at Starbucks to do my work.

**Response 15:** Great dino story :-) You are a very early adopter of cloud computing, perhaps a bit too early? All problems you report about CFD 360 are a few years back, but CFD 360 has changed a lot over the years (not just its name). We recently conducted an UberCloud Experiment with **Mark Lobo**'s application on virtual testing of a severe service control valve with CFD 360 (version 2014), and his conclusion is not that disappointing; in contrary (published recently in Desktop Engineering, http://www.deskeng.com/de/testing-simulation-cloud/). With all your past experience about **CFD 360** it would be a great experience to now perform a new experiment and compare both experiences.

**Comment 16:** I agree with a lot of what has been said previously.

I just purchased a workstation which allows me to run all of my simulations in my home office. From that point of view it is cheaper in regards to the cost of running in the cloud and extra costs to upload/download the model.

I am also concerned with running the simulation in the cloud. If there is a problem I simply stop the run, fix the model and run it again. It is quicker this way. I am not sure it is that simple in the cloud.

I would consider running in the cloud if I have many jobs to run with a tight deadline.

**Response 16:** As an engineer, your own **workstation is your daily workhorse**, there is no alternative. First, you have to have one for your designs, analysis, tests, and so on. And then, it's even quite powerful for many of your simulation work. Nobody ever would ask an engineer to give up his major working tool. Analogy: I've never give up my car at home IN FAVOR of a taxi; a taxi is simply too similar, and my car is much more flexible. But, I drive my car to the airport to get on the plane. And I don't own that plane, but I rent a seat from Munich to Houston, and another one from Houston to New Orleans to arrive just in time for the Supercomputing Conference. You get my point: **using your workstation for most of your daily work, and using the cloud for everything which you can't easily do with your workstation.**

But already with a second car this is different. Only if you fully utilize your second car as well then it's worth it. But if that second car sits in the garage most of the time you might think of using a taxi instead, from time to time. This **'hybrid' model gives you all the flexibility** you need plus it's most economical.

**Comment 17:** I agree with the above. The only time I found cloud simulations useful was when I had to create performance curves for multiple variations of a product that took 4 years to develop. It took thousands of simulations to come up with a product that performed as required. I tried to do some of the analysis work on CFD 360 when it came out but quickly realized that I was not going to make the deadlines or stay in budget because of all the issues discussed above.

**Response 17:** This is one of the **perfect use cases for the Cloud: parameter studies**. They can nicely be done in parallel; you do 1000 simulations in the time of one. Maybe you ran into the same issue as XYZ above: you were simply too early. Vendors often tend to provide a new service too early so then it's the customers who have to do all the testing. This can be very damaging, for both sides, unfortunately.

I can only repeat what I mentioned to XYZ (above) about our UberCloud experiment #142 with engineer Mark Lobo, about virtual testing of severe service control valve, on CFD 360, the case study has just been published in Desktop Engineering, http://www.deskeng.com/de/testing-simulation-cloud/.

BTW, if anybody is interested to **perform your own CFD 360 experiment** we have developed an experiment platform and provide all the support you need, and it's free. The only request is to write a case study at the end, see: https://www.TheUberCloud.com/hpc-experiment/.

**Comment 18:** As a fresh master student of Mechanical Eng. having dealt with many CAE projects, I will definitely go for that kind of computing method in order to save my CPUs' life time and **get the job done as fast as possible**. I believe Cloud Computing is one of those nice sharing ideas in which one can find human community spirit. Moreover, nowadays co-simulations need more solving potential; I mean in the time that even some simple app like Google Chromium uses your graphic card to accelerate itself (!!) it's not surprising that CAE software are going beyond. After all we can buy an extremely powerful PC and put it in our home or office, but I see future engineers pre-processing their CAE scenarios by their smartphones and let the - may I say - kind Cloud solve it for them. That gonna be beautiful.

**Response 19:** What a great vision you are sharing with us! You are representing the younger generation of (what we call) digital natives. You grew up with zeros and ones and find it natural to use all these new digital tools. More senior people like me tend to be still suspicious about the 'negatives' of such tools and deeply investigate and try to balance benefits versus barriers. I am sure **most of us will change minds over time** when others have adopted it **and we get used to it**. Still there will be areas and applications which will need a long time to get on the cloud, especially when your company's assets are at stake. And then even with your own data, hackers and even your own employees might try to steal your data. Our data is not even safe on our own premise today which is a scandal. And believe it or not, in such situations your data might be safer in the cloud because of the professional infrastructure and expertise of most of the cloud providers today.

**Comment 20:** Yes u R absolutely right Mr. Wolfgang Gentzsch. Well . . . it has always been like this through the history of civilization . . . Fathers work . . . Children harvest ;) However, we as "the younger generation of digital natives" expect the moral postmodern technological era mainly from you dignified professional Germans' old children ;) So thanks to you.

**Response 20:** Although this is not our topic here, I want to contradict humbly to your kind statement: Fathers work, children harvest. IMHO, it is vice versa: we (fathers) are living at the expense of our unique planet, and our children have to pay for it...

Back to the topic: I believe indeed in your 'postmodern technological era' **with a fundamental contribution from the cloud services paradigm.** Just one example from even today: 15 years ago if you wanted to start an Internet based company you had to have enough VC money to be able to buy a big server (easily a million, if not more). Today, you start your services company, without expensive server and software needed upfront. You develop your services with free software (like WordPress, Google Drive,...) on a small server in the cloud. When the first customer comes (and often pays upfront) you offer the service instantly, through your cloud infrastructure. And this is just one example, and it's only the tip of an iceberg, and it is only the beginning. This new paradigm means there are great new jobs for those creative minds out there who know how to handle this new technology. The implications are simply overwhelming.

**Comment 21:** In an ideal world, yes! There is so much to be saved in operating costs and hassle of maintaining in-house HPC systems.

**Response 21:** I admit I was suggesting an ideal world, without any cloud roadblock anymore. We will still face **a few roadblocks** for some time, those which we can resolve by technology (e.g. **bandwidth and data transfer, easy access**) and those in people's minds (e.g. **security, losing control**). We can resolve most of the other roadblocks today, with Linux container technology which tackles (softens, or even removes) **security, portability, compliance, data transfer, standardization, software licensing, resource availability, cost transparency, and need of cloud expertise:**
http://www.scientific-computing.com/news/news_story.php?news_id=2552.

**Comment 22:** Different issues have different roadblocks. If all were removed, then the only thing keeping companies from cloud bursting would be inertia or FUD - both of which can be powerful motivators.

Respectfully, I think a more useful conversation is to identify which are the major roadblocks for each industry, and see to what extent current or near-term technology can address them.

Let me get the ball rolling with two examples: oil and gas geoprocessing, and animation rendering processing, just addressing the **security** dimension:

Oil and Gas has huge amounts of data whose **confidentiality** is extremely valuable. Thus, security and scaling are critical.

Cloud-based animation rendering, while certainly of scale comparable to Oil and Gas, has different confidentiality / security requirements I argue- if you're rendering a major movie, stealing a frame or two from the cloud render farm is not as bad as stealing the whole movie, something much harder to pull off due to the number of GB involved.

So for rendering, security certainly cannot be porous - but the key is to prevent wholesale theft of large datasets - e.g. whole movies. People will still pay to see a movie even if a frame showing the new Batmobile gets posted on the Web.

For e.g. O&G, any portion of the dataset could potentially reveal the GPS coordinates of a play. Thus, having a hacker get his hands on just O(100) bytes of data could represent a huge hit to the cloud client. For this type of cloud application, security has to be hermetic against any data escaping and hacking attempts must be detected and stopped instantly.

I will amend my comments though - it might be extremely valuable to find out that Animation Studio X is working on an adaptation of a Hans Christian Andersen tale regarding ice and snow. :-) You could rush production of a copycat. And that is also just a few bytes of information. But in general, for some industries the whole is far more valuable than the intermediate products, whereas for others even the intermediate products are extremely valuable.

**Response 22:** Thanks for getting the ball rolling, by selecting mainly security out of the most common roadblock today (which are: **security, portability, compliance, data transfer, standardization, software licensing, resource availability, cost transparency, and need of cloud expertise**). Many of these roadblocks can be tackled (softened, or even removed) with Linux container technology mentioned in my answer to XYZ, above. Concerning security, however, as you rightly point out, there is a whole spectrum of applications which require highest possible security, all the way to those which don't require any. The strategy seems easy: **don't use the cloud for applications where security is paramount. For any other case it might not be such a problem.** Because this is a multi-dimensional problem (it really depends on the application, the data, the infrastructure, the people, etc.) this is often a case by case decision making.

**Comment 23:** It really depends on the **cost**.

**Response 23:** Indeed, if you **compare in-house server cost versus cloud usage** for the same amount of workload (say per year) then there is a cross-over point in the diagram with two curves, one for the server and one for the cloud: the higher the utilization of a server the better its economics (cost per FLOP), i.e. cost per FLOP goes down the more the system is utilized, and vice versa. Cost per FLOP in the Cloud is almost constant, more FLOPs, higher cost (one core hour today is 5 – 10 cents). The **cross-over is somewhere by 70% to 80% of server utilization**. I have made a little calculation which you can see here: https://www.TheUberCloud.com/cost/.

**Comment 24:** I'd prefer that my data was only kept on my own machine. Maybe the software I'm running could address the cloud CPU cores somehow more directly in parallel, so that nobody snooping in from the outside could get any idea of what I am trying to model.

**Response 24:** You are pointing towards the issue of **security**. Every cloud provider knows that security is an issue, especially for his customers. **Without providing a high level of security a cloud provider could not exist, no customer would come.** Therefore, every cloud provider has at least one security expert, usually quite expensive people an SME (small and medium enterprise) can't afford. Therefore, **cloud infrastructure tends to be at least as secure as the SME's, if not better**. Please see also XYZ's comment on security below.

**Comment 25:** Yes, definitely. Let's talk about the roadblocks and how to remove them.

For me the biggest roadblock for CFD in the cloud is **licensing**. I'd like to be able to use commercial CFD software like ANSYS Fluent for my consulting work, but the usage I would get out of my license is too low to consider this option. As a result I only use **OpenFOAM** for my consulting because (a) **it is free** and (b) it is **straightforward to deploy on cloud**-based HPC systems. I could definitely increase my business if I could license Fluent by the hour, but I could not make a return on my investment in a yearly license. Sometimes I do have to turn work down because the client requests the work be done using Fluent and I can't support that.

Based on my research all of the companies offering **ANSYS products on cloud systems require you to bring your own license**. This only solves half of the equation. Yes, you've solved the problem of investing in HPC hardware, but the licensing problem is still there.

**Response 25:** Licenses for consultants is certainly different from licenses for companies and their R&D departments which often focus on one specific project, product, service, which requires a very special software, say Fluent. And then they have to buy a certain number of licenses. When they need more (peak demands, deadlines) they can cover this peak with additional cloud cycles. And because they are ANSYS customers already they can manage to get temporary cloud licenses. For consultants this is tricky, either you work for one or only a few customers who ask you for specific Fluent expertise such that your Fluent license is fully utilized and you get your money back from your few customers, or you do exactly what you do, trying to minimize license cost, by e.g. using free software. And using free software in the cloud comes also with a great benefit of no license cost.

**I currently see ANSYS changing a lot.** And I can already see the day where you don't have to be an ANSYS user to use Fluent in the Cloud, pay per use.

**Comment 26:** **Security** is a major concern when it comes to the cloud, but has been mostly addressed by providers. What we call "cloud" is really a number of connected servers which are (can be) maintained at the same security levels as customer in-house servers and even better. **Cloud connections through SSH and VPN are common practice and extremely** secure with even more strict measures available (dedicated random number generators, digital signature

checking, IP whitelisting, etc.). In fact security is so good that even banks and insurance companies operate on the cloud now and several companies are offering secure and efficient big data-sharing services and software, and they are doing very well (NetApp, DDN, etc.).

**Licensing is a bigger hurdle** than security in my opinion and one that will take several years to address, mostly due to provider inertia and conservatism. As XYZ pointed out **ANSYS Fluent**, one of the two major CFD codes in the world today, is still not available under a compute-on-demand license. The other major code, **STARCCM+**, is available under a **power-on-demand** license which goes by blocks of hundreds of core-hours valid for a year (~10K$). It is not ideal but allows much more flexibility than ANSYS and is a step in the right direction. **Open source** software such as OpenFOAM is great for the cloud since it is GPL and one can use it without having to pay for the GPL. However, GPL is also cause for a major concern for people like us who develop their own OpenFOAM solvers or modify existing ones with proprietary models or source code. Such code can, and much of it does in fact, fall under the definition of GPL code. But we don't want to release such code for competitive reasons and are generally wary of uploading it to the cloud since there does not necessarily exist any protection for our proprietary code from third-parties under the GPL. We're not sure our code can be treated as confidential without infringement to the GPL so we prefer to hang on to it and use it strictly internally. All to say that even GPL is not perfect and **that licensing on the cloud is still a work in progress**. I for one am looking forward for what's to come...

**Response 26:** Valid comment on security, nothing to add. O.k., **Licensing**, but this landscape is currently changing dramatically. **Many ISV's are currently working on an additional more flexible licensing** scheme. Especially ANSYS has an interesting reseller community which is quite independent from ANSYS. For example, CADFEM in Central Europe offers a dynamic license scheme with eCADFEM. Ozen Engineering in California is preparing a temporary license offer for the Cloud. And many others.

Your OpenFOAM GPL issue is an interesting one. Is there anybody out there who can comment on this concern?

**Comment 27:** Wolfgang, it's good to know things are moving wrt to commercial licensing. This could also be an **interesting opportunity for smaller software providers** to increase their market share. Picture a small CAE software provider who makes their product available on a commercial cloud platform. It comes installed and ready to run with all the trimmings, all you have to do is purchase the appropriate license. Basically you got an advertising platform inviting potential users to test your software free of charge in the exact same way they would be using it. The "Cloud" has this kind of potential: :$ :$ :$.

The GPL issue is a tricky one. The most viable solution in my mind would be some sort of a limited-GPL caveat offered for cloud environments under the classical GPL, but then you would need to introduce an appropriate legal definition of a cloud environment as an extension of the

end-user's "home" computing platform as far as running their GPL application on the cloud. Getting a headache already... Might be worth starting another discussion on it.

**Comment 28:** Haven't tried this yet, but a quick search says that you can create shared compiled libraries to use in your custom solvers. http://www.openfoam.org/features/creating-solvers.php

Would obfuscation work for this? Keep the public functions names visible but obfuscate everything else in your compiled library? Basically you would have a public API layer which calls other functions inside the precompiled shared libraries. This way even if someone decompiles them, they won't be able to make much sense of it.

**Comment 29:** Based on the assumptions stated there would be no reason not to use cloud. It may also allow access to greater CPUs and align with actual license need.

**Response 29:** Thanks, good point. Often, while you bought **an HPC system for in-house use, this system gets older quickly, while in the cloud you can always get the latest technology.**

**Comment 30:** Wolfgang, exactly a good point. Some vendors use a pay as you solve approach. This means you are more efficient with your license costs with a 'peak and trough' workload.

**Response 30:** Thanks, interesting. I tried to find this model on the Internet, but was not really successful. Would you give me an example please? Is there a difference to 'pay as you go'?

**Comment 31:** An example would be the simulation tools offered by **Autodesk** - fine for simple LE analysis. **Simulia** rolling out the 3D experience will offer a similar approach, I believe. It's worth looking on their website for more info.

**Response 31:** Thanks. This reminds me to point you to Mark Lobo's UberCloud experiment #142 about virtual valve testing with **Autodesk Simulation CFD**, see the article in Desktop Engineering here: http://www.deskeng.com/de/testing-simulation-cloud/. During his experiment he had indeed kind of a peak and trough workload.

**Comment 32:** I do sporadic freelance work - this could be very useful for me, if it meant that I could access licenses only when I need them.

**Response 32:** Especially **for freelancers like you clouds are very beneficial:** it doesn't make sense at all to buy a server just for sporadic work, because of the server's high TCO (total cost of ownership). **According to IDC a $70K server comes with a TCO over three years of $1 million**; not including the application licensing cost. Together, hardware, software etc. sums easily up to $1.2 million over three years, $400K per year, or $2000 per work day. Hey, a one-day ANSYS usage in the cloud (s/w and h/w) is cheaper today. So **the ideal combo is to do your daily stuff on your workstation where you certainly have your engineering application license, and then use the same application on a larger system in the cloud**, either for the same configuration (then your jobs run faster) or for a much larger one (with e.g. finer mesh, or more complex physics etc. (then you get higher quality results; what a great spectrum of new opportunities in a sudden.

**Comment 33:** Does lack of tightly-coupled I/O count as a roadblock?

**Response 33:** Yes it does. But then, as we all know, you are simply not using the right architecture for your application. At UberCloud, for example, we are checking the user's requirements and, in such a case, would offer only tightly-coupled HPC 'Cloud' systems, like yours, or XYZ's (see below).

**Comment 34:** We do have an HPC cloud in operation, but it is not a capability supercomputer. Instead, we are focusing on low-latency edge computes that can be used to deploy sensor fusion and other deep computes at the edge at latencies that are matched to the I/O rates. Certainly high performance computing, and certainly a cloud deployment model, but not a centralized capability supercomputer.

We looked at the commercial use cases of HPC and concluded that **hybrid clouds are the better architecture to deliver value to the market place**. 90% of our HPC software is organized as an elastic micro-services architecture, with the deep computes a mixture of scale-up through hardware accelerators, and scale-out delivered through public cloud resources (HP Helion because it offered much faster VMs than AWS). The ability to move workloads around different infrastructures, the hallmark of a hybrid cloud architecture, was for us the key value creator, and we continue to pursue perfecting that solution.

The final non-technology attribute that governed our decision to pursue a hybrid cloud architecture for HPC is the observation that when value-added data sets grow past 10TB or more, they effectively become beholden to the infrastructure that is used to manage these data sets. Since thousands of organizations are now generating and refining these value-added data

sets, it is not hard to postulate that to mash-up these data sets and leverage deep computes, **you need some sort of hybrid/federated cloud capability**.

**Comment 35:** Lack of tightly coupled I/O does count as a roadblock, but not all clouds suffer from this indignity ;-)

**Comment 36:** How we can remove the IP roadblock?

**Response 36:** Clearly **Intellectual Property** is among a company's most valuable (and most vulnerable) assets. A (nasty) counter question could be: How do you protect your IP today? From **outside threats**, from **in-house threats**? If your company has implemented strong rules and regulations, then this might possibly lead to a similar solution in the cloud. A **secure end-to-end solution might do it, with encryption on your side, cloud connections through SSH and VPN which are common practice and extremely secure with even more strict measures available (dedicated random number generators, digital signature checking, IP whitelisting, etc**. as pointed out by XYZ in the LinkedIn CFD Group, **and a single-tenant secure server environment in the cloud**. But, it's still very much **dependent on the company's compliance regulations**.

**Comment 37:** Yes, I'll transfer my work projects to on-demand cloud.

**Comment 38:** Objection! Leading the witness! With those givens, I think most people would have to say yes.

With my blinders on, I'd say the one factor to assess for yourself is just how "**burstable**" your usage patterns are... the less dynamic, the easier it is to buy some blazing local hardware spec'd for that load without breaking the bank.

With my blinders off, I'd say the question is moot with all those givens. The cloud is more like a platform option at that point. Should I move my goods by ship or jet? Depends on what I'm hauling.

**Response 38:** Thanks. But wrong: I posted this question onto some 30 LinkedIn Groups, and got about 50 great answers so far, but even with this obvious question (expecting a Yes or a No), most colleagues started with bringing up and discussing some of their 'favorite' roadblocks. This shows one thing to me: although many of the roadblocks have recently been resolved or at least

softened, most of these roadblocks are still in people's heads. Take security; is my office or home environment still more secure than the cloud? Licensing: there are many ISVs today who finally start offering more flexible cloud licenses. No standards (resp. portability, resp. vendor lock-in): recent container technology will solve this one pretty quickly. And so on.

Letting your usage pattern and your wallet speak is a good thing, not emotional, either I need additional capacity or not, either I can pay for it or not, that's all measurable and decides for ship or for jet.

**Comment 39**: Ah, I'm wrong in the respect that people couldn't follow the directions of the question?

**Response: 39:** My question was pretty clear (and I expected a short and clear answer, like: yes, because I have this and that need; or no, I don't see a benefit...), but most of our colleagues seem to be still quite concerned about the current roadblocks that they still insist to discuss them, instead of such a 'Gedankenspiel'. And I have to say that I love these roadblock discussions very much, because they give me a chance to reconsider our UberCloud approach and check it again and again, critically.

**Comment 40:** I'm afraid I agree with XYZ simply that I think the question is really just baiting people. If we take the question literally, it has no meaning. It is like saying "If there was no differences between widget A and widget B, except that widget A was cheaper would you still consider buying widget B? Why or why not?" It makes no sense. If there were no difference except price, then why wouldn't somebody choose widget A?

What you are pointing out is that people essentially ignore the question and say "But there is a difference between widget A and widget B!" and then proceed to argue about why one is better than the other. It strikes me as more a psychological experiment more than a sincere question.

**Response 40:** Well said. If all cloud roadblocks have been tackled, there is no reason not to use cloud. The only 'no' could come from colleagues who simply have not a big problem to solve which goes beyond the capacity of their workstation, so they don't need additional compute power. That most of the discussions still circle around the common roadblock proofs that many colleagues simply can't imagine that one day all roadblocks will disappear. In that sense these discussions / contributions are reminding us to work hard to remove these roadblocks, one by one.

**Comment 41:** Rather than paying for my computing in dribs and drabs, I would install an Intel **Xeon Phi co-processor** (240 cores at 1/3 power of a PC) and effectively have a small super computer with more power than I can use for 4 years. At which time I will buy another desktop plus co-processor. The challenge is to develop software to sop up the computer overhang.

**Response 41:** As long as you can solve your problems with your current workstation, fine, why should you use additional resources (be it server or cloud). If your **problems are much bigger** however (e.g. 500 million elements), or they should **run much (10 times) faster** because of time-to-market is an issue, or you want to do a **parameter study** with 100 cases in parallel for e.g. finding the best material for your product, then yes you should **use much faster, bigger, cheaper resources**, which go far beyond what any workstation will be able to do within the next 5 years.

**Co-processors:** Unfortunately, all the results which e.g. ANSYS (and other commercial software) has achieved with Fluent on Nvidia 1000-core $3000 GPGPU so far show improvement factors of just 2.5, not very encouraging for using co-processors. **Our complex engineering applications will not be ready for these sophisticated and memory-bound computer architectures within the next couple of years.**

**Comment 42:** I vote "Yes" and lets just see what the marketplace says after this tool is developed. If someone loses their shirt over this, so be it. It's not my money. BTW, I do agree, it is a poor question. If ALL obstructions are taken way, then sure water flows.

**Response 42:** Water flows, Martin, yes almost. Except I have enough water already from the well in my back-yard. Meaning, I don't need additional resources, my applications are small enough to fit into my workstation nicely. The reason for my question was two-fold: first I wanted to know first-hand what the actual opinion about clouds is in our community, and second I wanted to **create awareness for CAE in the Cloud**. And perhaps there are additional concerns which we haven't seen so far, but which we should tackle as well to ease the additional use of Clouds for CAE.

**Comment 43:** Yes, I would relocate all the CAE I could to the cloud in a heartbeat. **The Cloud provides real time access** for both CAE Engineers, and Customers' engineers anywhere in the world. And it creates a natural platform to "sell" into customer engineering functions.

**Response 43:** We are inches away from the vision you describe. The technology to do this is here. Sharing of cloud resources, applications, development tools, collaboration space like BaseCamp, and more, all this is possible today. One major enabling technology to accelerate this trend is Linux container technology like Docker. The next two big hurdles to overcome are: will

your company policies allow you to use public clouds; and (how fast) will engineers MENTALLY adopt and join this cloudy trend.

**Comment 44:** If all the obstacles are removed ... then (as per XYZ)... The cloud is more like a platform option at that point.

@Wolfgang, Thank you for your gentle response to what must have seemed to you as a "Stick in the mud attitude!" Perhaps I should try to elaborate on some of my concerns as well as reasons.

[1] First off let's dispose of your Fluent example of failure to take advantage of the GPU. Fluent is a Legacy Code. To my mind going back at least to the 80s. No amount of conversions will allow it to run in a parallel environment. The same would happen on the cloud.

[2] My world also includes computational Photography, additive manufacturing (3d printing), optical scanning. micro-fluidics and the internet of things. All of which requires me to have plenty of local compute power to aid in their processes.

[3] For the really modern parallel software, there is plenty of computing power left for accommodating these in Virtual Machines. Also communications are free for support in the form of Facetime, Hangout, Skype.

**Response 44:** Thanks for elaborating on some important details. Concerning **Fluent it runs today (if needed) on many hundred cores in parallel** already successfully (and, to my surprise there was just 2 weeks ago an announcement from ANSYS that Fluent was running on 36,000 cores on the NCSA supercomputer). The real problem is the GPUs, or more accurate, the **complex memory hierarchy of GPUs** which is unable to sufficiently feed the 1000+ data hungry parallel GPU cores; and although the solvers in many legacy codes run with speedups between 5 to 10, the overall code then often runs with a speedup factor of around 2 (because there is a lot of serial (non-parallelizable) code around these solvers. On the other hand, there are indeed applications (algorithms) which run with 100+ speedup on GPUs but unfortunately only very few ones in CAE.

Concerning your world of computation it really is a multifaceted one. But as I mentioned in one of my comments before, **not every application is well-suited for a parallel system,** including the cloud, and it's very natural then that you try to run them on your desktop workstation (if your workstation – as mentioned before – is big enough and fast enough).

This new area of cloud computing will develop in the same service-oriented way as we see it within many of our daily life applications: take our mobility: to cover our average needs we buy and use our own car; in certain situations I can't use my own car (business trips; others are already using it; etc.) so I rent a car for a day or a week; or I am in a hurry so I call a taxi and accept a higher price; or I have a real peak demand (e.g. going from Munich in Germany to SFO) so I take a plane. All depends on your needs and on your budget.

**Comment 45:** @Wolfgang, continued. Sorry I missed your comment on what would I do with a 500 million element problem? First of all I would conclude that whoever thought of it did not understand the fundamental engineering problem. Secondly I would conclude that the analyst did not understand the FEM process and is using the wrong type of element. Finally I would feel nostalgic about the old days where great engineering feats were achieved with just a slide rule and/or a room full of people cranking on calculators.

**Response 45:** **Hundreds of millions of cells are quite common today** when complex geometries, physics, and other requirements come together, see e.g. the 3-years old example of 210 million grid cells for the simulation of energy-efficient truck, https://www.olcf.ornl.gov/2011/03/02/bmi-uses-jaguar-to-overhaul-long-haul-trucks/ or the 300 million cells for race cars and over 500 million cells for large architectural projects, http://www.wirthresearch.com/about/development-in-the-digital-domain.html. Many of these examples demonstrate a need for higher resolution, quality, accuracy, therefore **the need for High Performance Computing**, far beyond the performance of workstations, for example.

**Comment 46:** @Wolfgang, to continue the discussion on FLUENT and Legacy Codes. All codes come up against Amdal's law sooner or later (http://en.m.wikipedia.org/wiki/Amdahl%27s_law). My point is that Legacy codes were not written with parallelization in mind, especially with local GPUs. That is why I chose the full virtual core Intel Xeon Co-processors. Let me propose a test, let us call it the Amdahl Test. 'The limit of parallelization on a system is reached when a doubling of the co-processors produces less than 25% improvement in overall computing times.' This is then a function of algorithm, computer implementation and localization of memory by the type of GPUs. If you tell me the point at which a program in a system fails the Amdahl test I will tell you the point at which it is no longer cost effective to throw computer hardware at the problem. In the past I have written software to determine the parallelizability of codes. But Amdahl's test appears to be a good simple test. I will get back to you on the second part of your discussion.

**Response 46:** I am in full agreement with you here. I like especially your 'summary' statement that the limit of parallelization is "a function of algorithm, computer implementation and localization of memory"; absolutely. And while parallelization on Intel Xeons for many of our engineering codes often works nicely up to a few hundreds of processor cores, GPGPUs add another level of parallelization and complexity, which often requires an additional effort on our side to push parallelization further.

**Comment 47:** @Wolfgang, thanks, for those of us confined to speed ups of 8, on our PCs, I would gladly take a speed up of 100 (divided by three to adjust for the speed of the Xeon co-

processors).
You are right to point out that this only takes care of the parallel or vector problem but not necessarily the parallel imbedded vector problem. The additional use of GPUs will help resolve that. Fortunately the addition of more GPUs is relatively inexpensive. I cannot speak for other languages but the Microsoft AMP feature makes the use of imbedded GPUs fairly straight forward at the coding level of C++.

**Response 47:** I guess the speedup factor you mention comes from an 8-core Xeon E-5 workstation? On these nice machines well-tuned application codes tend to deliver a speed up of 5 – 7, which is a great economic value. Now adding one GPGPU, e.g. Tesla K40 for a whopping $3,000, comes along with an additional complexity as we discussed before. You are talking about C++ which implies that you have access to the source code and you have learned to program and optimize such kind of huge application codes, which most engineers however don't. And most of the engineering codes are not accessible for you, so you rely on the ISV. And as we already discussed, commercial engineering software like ANSYS for now just gives you a GPU speedup factor of 2 – 3, that's it. So, my conclusion is that if you require capacity and capability which goes far beyond your (pimped) workstation, now you have a great additional opportunity: the CAE Cloud ☺

**Comment 48:** @Wolfgang, here I will try to comment on the existence of problems requiring 500 million elements. Let's simplify the problem and call it 500 million nodes. This resolves roughly into 800 x800 X 800 mesh in 3d. The error in a simplex element stems from two sources. The first is its truncation order of the polynomial interpolation function 0(h**n), and the second is the acute angle of edges at its nodes. In the early introduction of the isoparametric hexa, Zienkiewicz and his co-workers showed the efficacy of their elements culminating in their summary that a hexa was worth a thousand tetras. They did not delve too deeply into the cause of that. The tetra elements and even the serendipity 20 node Hexas suffer from truncation error due to premature truncation of their interpolating polynomials. But not to the extent that they require 800 meshes in 1 dimension. So the answer could be the distortion of the tetras, our second source of error. My recent studies of using 27 node hexas for thin shells and Tom Hughes use of isogeometric Nurbs based elements show a further advantage of using Hexas over Tetras. There is mainly one reason for the widespread use of tetras in FEM and that is the existence of automatic Delaney triangulation mesh generators. However these generators have led to increasingly larger mesh sizes because of the lack of ability to determine the accuracy of their results.

In the past I have advocated the automatic subdivision of each tetra into four hexas so that we could better control the mesh accuracy. Recently, with further discussions and demonstrations from Bob Rainsberger, the developer of True Grid, I have come to understand that semi-automatic generation of hexas is quite practical. So when we have used up all our size capacity with tetras, we can fall back on the use of hexas and/or their refinements isogeometric elements. I think I have displayed my biases, I thank you for your examples, they are indeed a

tribute to our profession in which we are able to use any and every tool at our disposal to solve the outstanding problems of the day.

**Response 48:** Thanks for this instructive argument. Indeed you are sketching a second route to improving design and product quality and time to result (or market): improving our numerical algorithms at the heart of many engineering codes. A great example is indeed the more recent area of isogeometric analysis with using NURBS and T-Splines as a basis for constructing element shape functions which is being championed by Tom Hughes at the University of Texas at Austin (who I know well). It's a great way to seek better accuracy to compute stresses, velocities, pressures and buckling loads.

The problem with all new methods is that it might take 10 years until they get into commercial codes. And then we are again at the beginning of our discussion: Imagine, with such a method, you would speedup your computations by a factor of 10 (to get to the same accuracy) resulting say in a wall time of one day for simulating a complex product, then the same computation might take only one hour on a 100-node server, and if you can't afford to buy such a server for $500K, then you get this computation done in the cloud within one hour for a few hundred $$. Again, my recommendation is to make use of (and benefit from) the whole spectrum of using the workstation, and buying a server, and renting compute power in the cloud.

**Comment 49:** @Wolfgang, please let me finish my thinking on the Tetra vs Hexa argument. Another way to look at the Zienkiewicz rule is to take the cube root of 1000 and say that the equivalent mesh in each direction can be reduced by a factor of 10 in each dimension. Then our 800 X 800 X 800 node mesh of tetras reduces to a 80 X 80 X 80 mesh of hexas with an error of 0(h**2) assuming an 8 node linear hexa. Or if you adopt a 27 node hexa with error 0(h**3) and equating errors to get the same error, we get an equivalent mesh of 18 X 18 X 18 for the same error. I would be prepared to solve the problem with either type of hexa on my desktop in minutes. In terms of parallelization, the 27 node element transfers its complexity to the generation of the element stiffness and the parallelization problem is spread out between the stiffness generation and the equation solver. I do both of these in my MPACT program, so I prefer the use of the 27 node hexa. Of course, the ultimate complexity of the geometry will also dictate the size of the mesh and the type of hexa selected. In the spirit of finite elements, I have chosen to stay at the quadratic 27 node hexa. Tom Hughes on the other hand has chosen to push beyond to the cubic 64 node hexa and beyond because he feels confident he can handle the CAD equivalent geometry problem with isogeometric elements.

Finally I agree with you that we should preserve our flexibility for the future by being prepared to adopt whatever computer system type and size necessary to solve our problems.

**Response 49:** This confirms your and my argument that real progress in numerical simulation can only be made when following these two directions: working hard on improving our numerical algorithms, AND using the right computing tools (workstations, servers, clouds) and architectures appropriate for running these numerical simulations.

Unfortunately, it's my own experience that it might take 5 - 10 years until new (and complex) methods like the one you describe become available in popular commercial engineering codes and thus will be widely used by our engineering community on a daily basis.

**Comment 50:** 1) Yes, **I would move to the cloud if I had absolutely no choice (i.e. if my local cluster could not handle the load at all)** and 2) No, if given a choice I would not move any of my analysis to the cloud.

Reason: I do not know how the **security** is handled back of the scene. Concerns about cloud security, definitely yes!

**Response 50:** Your point 1) sounds like you feel forced into the only one possible choice. But you still have a choice, just stick to those simulation jobs which can run on your workstation and don't do complex and high-quality simulations.... Yes, I know, you wouldn't be satisfied with the low quality of results coming from your workstation for such complex geometries / physics.

O.k., then, move to the cloud, and take your security concern seriously. You can very well know about **how security is handled** in the cloud. Just **ask your cloud provider** for a very detailed description of his security measures he has taken. And if you don't get a satisfactory answer, move on, to the next one. We can help you with your selection; we know a lot of **clouds which are more secure** than many company computers and data centers out there.

BTW, I covered the security topic (among others) in an article about: Workstations – Servers – Clouds – Comparing apples to apples – A little decision-making support for the undecided, here: http://www.theubercloud.com/workstations-servers-clouds-comparing-apples-apples/.

**Comment 51:** @Wolfgang: Your question is prejudiced isn't it? It supposes that the CAE applications on the Cloud are those that currently exist and can be run on the Cloud or on Pedro's machine. Then we get into your question of removing all "roadblocks" to using the Cloud. My position is that the Cloud can offer capabilities not found on Pedro's machine. These include **access to open source programs** that are **available on Simscale.com and Rescale.com**; and, new paradigms that take advantage of the cloud ecosystem.

The benefit of Simscale and Rescale sites is that an array of **simulation capability is available in my browser**. Simscale and Rescale have done all the hard work: they have downloaded, compiled and linked the codes - no simple task for most of us who probably have little to no experience with the Cloud environment, Linux and the programming tool chain; and, they have developed an **interactive framework in which engineers can perform simulation tasks.**

But, the real attraction of the cloud, at least for me, is that it supports a revolutionary paradigm for performing simulations as described on the MyCosim.com website. This paradigm

represents complex assembly in terms of independent components. One can envision, as this architecture matures, that the 500 million element model can be represented by 1000 components each of which can be simultaneously processed using component mode synthesis to reduce the size of the final solution component as well as the time to do so. Full disclosure - I am the developer of COSIM.

**Response 51:** I wouldn't say 'prejudiced': after running **163 engineering cloud experiments** (status Nov 2014) in CAE, Finance, Life Sciences, and Big Data, with almost half of them not successfully completed, we believe we know all kinds of **potential roadblocks**, and we have developed a solution to most of them (our UberCloud **software containers**). Now, for comparison with the reality, my question intends to explore how ready our engineering community really is for the cloud. And therefore I tried to be a bit provocative by pretending all roadblocks were removed. I am collecting all the answers and intend to publish a summary article. I already extracted the major roadblocks discussed in about 10 LinkedIn Groups and included my answers here: http://www.theubercloud.com/workstations-servers-clouds-comparing-apples-apples/.

And yes, I know your work on COSIM which seems to be already in beta. Looks like a great platform for building substructures which can dramatically simplify engineers' work and reduce computing time.

**Comment 52:** In regards to CFD, the **bottleneck is the memory bus,** and not the CPU. Also, the legacy CFD codes, in general, were programmed from the bottom up to be vectorizable and to some degree parallelizable. (In general they are easy to parallelize for multi core CPUs. GPUs are more along the lines of vectorizing than parallelizing) I cannot speak explicitly for Fluent, but I suspect the heart of it is similar; however, maybe, recent additions/changes have broken that. Coders are a little sloppy thanks to Moore's law. However, in regards to CFD codes, coding (other than performance optimization), is not that challenging. UIs are much more complex. And, even to this day, there never seems to be enough computer capacity for CFD. Turbulence models are just that, a model.

In regards to using Cloud computing (i.e. time sharing 2.0), that has been done for a while in one form or another. Changes have basically been evolutionary. And slow at that with the occasional discrete jump, such as the introduction of Linux. Regardless, in the end what needs to be done from a business point of view is trade off on such things as cost, turnaround time, minimization of analysis error (i.e. oops, wrong units or wrong geometry), and the clients possible need to use a program of record. If the resources exist in house, great. Otherwise off premises resources need to be found if expanding the in-house resources is not justified. But sometimes if in-house resources do not exist then the client is willing to reduce/change the requirements. And, sometimes, pre-existing in-house engineering expertise goes a long way in reducing the amount of resources required. In other words, creativity can be adjusted to match the available in-house resources.

**Response 52:** I fully support your points regarding CFD, especially those using explicit discretization schemes. ANSYS Fluent for example recently ran on 36,000 cores of the NCSA supercomputer. Concerning GPGPUs unfortunately the combination of parallel CPUs with vector-parallel GPUs makes full optimization and scaling a challenge (GPGPUs are many-times more data-hungry than CPUs), especially for implicit codes.

To your 2nd paragraph: Cloud is not necessarily time sharing, i.e. many users use a cloud server at the same time (it could be, yes). But **many cloud providers offer bare-metal single-user servers especially for high-performance engineering and scientific codes**; sharing resources (multi-tenancy) would affect performance dramatically. And in this case, sharing is not many users at the same time, but one user after another. Your suggestion to expand in-house computing resources (95% of engineers are just using workstations) by off-premise resources on demand seems to be the way to go, and the **best economic model is indeed a balanced mix of resources (the hybrid model), including workstation, in-house server, and cloud**, for a detailed cost analysis please see my cost paper here: https://www.theubercloud.com/cost/.

**Comment 53:** BTW, in regards to workstations vs. servers vs. cloud, I would like to point out that having a cluster is no big deal. Email, VPN, and backups (all of which are required regardless of the existence of the cloud) are more of a pain than the compute server. And the cluster is VERY reliable. And in regards to utilization and CFD, historically, 80% and sometimes more is approached. If we have excessive bandwidth, we use it for non-critical runs. In general, remote computing is required for additional resources or having access to programs not in house. And OPEX spending of any significance needs to be approved. Unless the discussion is about $100 here or there. And I don't understand "And buying a new (another) system during depreciation phase is impossible, even if the system is getting old quickly." The pizza boxes on the rack system are cycled. They come and go. And security is what it is. The security of the cloud is only as good as the security of the company accessing the cloud. The list, at least to me, seems biased.

**Response 53:** Concerning the cluster, IDC research already in 2007 showed that **a $70K server generates TCO (Total Cost of Ownership) of $1 million over three years, nothing an SME can easily lift**. CoC (the Council of Competitiveness) found that 57% of engineers doing simulations are regularly dissatisfied with the limitations of their workstation, but they still don't buy a server which could solve that problem. Why? Because **compute servers are a completely different league, in acquisition, support, and maintenance. If it were no big deal everybody would add a server to their workstations, but they don't.**

Your 80% utilization is indeed quite common in large companies with many engineers in many groups and projects running their simulation jobs. **In small companies however, when the one or two major server users are on vacation, biz trip, ill, etc, then the server would not easily be 80% loaded on average**. The more varying load the better is to burst into cloud resources.

Buying a NEW system (as I discussed) is different from cycling the pizza boxes the latter often being part of a support & maintenance contract. The problem arises if your demands for more complex and high-quality simulations exceed the capacity and capability of all pizza boxes of your rack.

Finally, the list about in-house versus remote High Performance Computing in http://www.theubercloud.com/workstations-servers-clouds-comparing-apples-apples/ stems from the 163 cloud experiments which 163 engineers have performed 2012 – 2014 by moving their real applications and data to the cloud, runing their simulations in the cloud, getting the data back to their workstation, and finally writing a case study with discussing benefits AND challenges. So far 42 of these case studies have been published in two compendiums: https://www.theubercloud.com/ubercloud-compendium-2013/ and https://www.theubercloud.com/ubercloud-compendium-2014/. Certainly this list might seem biased, because I wrote this article with these 163 experiments in mind, and while already knowing that many of the roadblocks with cloud computing are being / will be resolved within the next one or two years.

**Comment 54:** BTW, from my experience almost all engineers doing simulation work are quite capable of running Linux and installing open source programs (assuming they are allowed to do so). To install programs, in general, one just uses the package manager or ./configure, ./make, and ./make install. Assuming the system admin does not do it. Honestly, a kid can do it. And they do. And, even if one uses the cloud, one needs computers at their desk, an intranet, printers, etc. so I assume someone is doing system admin. Sorry, I'm doing a little head scratching here. Honestly, I think some are intentionally underestimating what people are capable of.

**Response 54:** I agree, engineers are the most capable species I know, very creative, very hands-on, very bright (and I mean it). And they are certainly capable of mastering every challenge on a larger HPC server. But would they want to? Would their CEO want them to do that? Are they IT system experts or/and high-paid design engineer to help developing the next-gen product which beats competition? I believe that engineers are the most capable species on the planet, but also the most precious ones and we shouldn't waste this 'resource'.

**Comment 55:** If you truly mean "all" including cost, then clearly people would use it. They wouldn't even notice a difference ;-)

**Response 55:** For some clouds you are right, they are very easy to deal with, almost the same way as your desktop workstation (like a virtual desktop). **Some other clouds are still challenging to access and use**, especially for an engineer whose daily job is not IT, but engineering. Good news, we see a lot of progress recently on the cloud side, and a lot of interest on the engineers' side.

**Comment 56:** When there are no roadblocks, obviously yes. However I'll never believe the **security** part. So no.

**Response 56:** Concerning security: **Every cloud provider knows that security is an issue**, especially for his customers. Without providing sufficient security a cloud provider would not even exist, no customer would come. Therefore, **every cloud provider has at least one security expert**, usually very expensive people an SME (small and medium enterprise) can't afford. Therefore, **cloud infrastructure tends to be at least as secure as the SME's, if not better.**

Security has been mostly addressed by providers. To quote XYZ from the LinkedIn CFD Group, "what we call 'cloud' is really a number of connected servers which are (can be) maintained at the same security levels as customer in-house servers and even better. Cloud connections through SSH and VPN are common practice and extremely secure with even more strict measures available (dedicated random number generators, digital signature checking, IP whitelisting, etc.). In fact security is so good that even banks and insurance companies operate on the cloud now and several companies are offering secure and efficient big data-sharing services and software, and they are doing very well (NetApp, DDN, etc.)." And, by far most of the security breaches in recent years were reported for in-house data centers, not cloud: http://www.businessspectator.com.au/article/2015/1/12/technology/clearing-skies-cloud-getting-safer.

**Comment 57:** Security has been the number one concern from the users that I have discussed this with. I think that security concerns will dissipate over time as more and more things in our everyday lives move to the cloud. IMHO a bigger issue is that it is more work for the engineer. There are slow file transfer and there are other issues. Confusing cost models hurt as well. In essence, **we are a long way from having no roadblocks**...

**Response 57:** Yes, I agree, we are a long way from having no roadblock. Or more correct: **some companies are a long way from having no roadblocks. Other companies however are using the cloud today, with no roadblocks, today**. It really depends on the company (compliance issues), application (no suitable system architecture in the cloud), amount of data to be transferred (MBs versus GBs), security requirements (IP versus pre-competitive R&D). And it depends on the benefits you can get from the cloud (better, faster, cheaper).

Concerning security, please see my comment to XYZ, above.

**Comment 58:** Yes, this discussion starts with a false premise. We *know* that even with smart and big companies, personal information is hacked (cf poor Jennifer Lawrence and other

celebs). We *know* that Google harvests everything that goes on line, and that they scan all gmail. We *know* that the NSA and others have hacking programs that can get probably anything. And many of us know from personal experience that even just shopping at major retail outlets, like Home Depot, can result in the loss of personal and credit card information.

Moreover, as noted above, transfers of multi-terabyte files are not easy, and the ability to generate data is rapidly outstripping the ability to transfer it.

Never say never, but agreeing that it will be a long time, and during that time there will have to be *zero* incidents of information hacking, such as Home Depot, Target, and many others, which occurred in just the last year.

**Response 58:** I agree with your first paragraph, but you agree that it has nothing to do with the cloud. And not many engineers have multi-terabyte result files; **many have a few Gigabytes, and interactively they mostly just use remote visualization (with a few MBytes), and transfer the final result file only once at the end of the project (often using Dropbox or Box.com).**

**There will always be hacking, as there will be steeling, etc. But will my data always be safer in my office, than in the cloud? I doubt. Home Depot, Target, and the many others your heard of all have their own data centers, not clouds. By far most of the security breaches in recent years were reported for in-house data centers, not cloud:**
http://www.businessspectator.com.au/article/2015/1/12/technology/clearing-skies-cloud-getting-safer.

**Comment 59:** Perhaps another question is, does running an internal network/supercomputer cluster really turn out to be more expensive than using the cloud? If not, and in totum I suspect that for a large organisation that is true, why move to the cloud? Fashion?

**Response 59:** This is math and all about utilization: **in-house supercomputer clusters are reasonable and economic as long as you have a high average utilization (at least 70% - 80%). Otherwise, Cloud is cheaper**; and much cheaper if you consider the industry average which is between 10% and 20% utilization. I did the math here: https://www.TheUberCloud.com/cost/. BTW, **the best price/performance can be achieved with a hybrid usage model:** buying a server which can cover the (e.g.) 20%, and using the cloud for your peak demands (above the 20%).

**Also, purchasing a supercomputer cluster is capital expenditure, takes long procurement cycles, has to be justified towards and approved by upper management, and competes with other projects in-house, while Cloud usage is operational expenditure which is much more flexible, agile.**

**Comment 60:** I can think of many problems with cloud computing, like poor internet reliability, data security, cost, slow transfer speeds, sluggish online performance. With cloud computing you introduce one or more third parties into your business who may not have your same priorities in mind. You lose control of your project. Local control is always better than outsourcing but may come at some cost. Not all of these can or will be fixed. Trading the negatives for the benefits, what is the upside to cloud computing?

- Cost (maybe)

- Compute speed and resource capacity (fast computers are cheap these days)

- Automatic software updates (not always a good thing....bugs).

What other benefits may there be to trade against the negatives?

**Response 60:** Not everybody will use additional cloud resources immediately, today. There has to be a need (or even a pain), like for example your **competitor** is coming out in a sudden with a **better product** (functionality, quality, etc.). Further, not everybody has the roadblocks you mention; some might have a great **Internet bandwidth** (like A*STAR in Singapore using long-distance InfiniBand). Some don't care about **data security** because they are working on a pre-competitive R&D project; some just have a few GBytes result file and do the intermediate interactive checking of accuracy or convergence with high-speed **high-res remote visualization** (with tools like NICE DCV). And concerning introducing third parties into your business you might depend on, these third parties treat you highest priority if you pay them for, as in real life. And **losing control in the cloud**, yes, today this is true. But with new technologies like software **containers** which bring a lot of good features such as job monitoring to the table and more, you can **nicely control what happens with your stuff in the cloud**, as you do it on your desktop.

Benefits? See my cost paper at https://www.TheUberCloud.com/cost/.

**Comment 61:** Just to reinforce the point about network bandwidth vs local computing cost, I have a high end mobile workstation which I can take to customer sites to run and visualize simulations requiring tens of GigaBytes of SSD space, but at the same time the network connections are often so slow that it can take 10 minutes to download 1 MB of email. Perhaps this is an exceptional case, but out in much of the real world I don't see network speeds approaching what I would need to replicate my current $5,000 laptop for many years, and by then I'll have a 64-core laptop with 10 TB SSD.

To echo the comment above, apart from specialized applications requiring massively parallel computing or those which only need to be run occasionally on a SaaS basis to save money, what are the real benefits of the cloud? I'll agree it may be useful for backup of critical data precisely due to its lack of a physical location, but then **security** issues may trump performance. Anyway, call me unconvinced that I'll be using the cloud for running simulations in the next 5 to 10 years.

**Response 61:** In your situation with such a limited bandwidth I can only think of kind of a work-around for now, IF you have a need for higher performance: for the time of a compute-intensive project, **using box.com for moving your big data files back and forth to the Cloud**, and only when you are finally done and satisfied with the final simulation result, then let the final big result file send back to you via **FedEx. This is very reliable and secure, and you'll have all your data in your workstation the next day.**

Second, currently, you don't seem to have a real need for more compute power, so Cloud is not a useful alternative for you anyway. BUT, according to analysis from the US Council of Competitive, CoC, from the over 90% of all engineers who for large simulations just (have to) use their workstation, **57% of engineers are regularly unhappy about the limitations of their workstations**. They say that their **workstation is either too slow** (taking days for one simulation), or they are **not able to run a bigger configuration** (like 100 million nodes or elements), **or they can't apply the full physic** (instead doing some first or second order modelling approximation). These engineers up to now didn't have an alternative, except buy a bigger system. However, **a bigger compute server for e.g. $70,000 results in a Total Cost of Ownership of $1,000,000 over three years of life time**; no real alternative for many. Now, with Cloud, they can buy Fluent, STAR, OpenFOAM, and other CAE software fully bundled with hardware in the Cloud, for a couple hours or days, available at your fingertips, for a few hundred bucks. This IS a real alternative now.

**Comment 62:** The issue of **software stability** is critical. We often freeze the software from updates at the start of a 3 year project. If there is a mandatory update in that time then we have to repeat some of the previous work to check that the update has not affected the results.

**Response 62:** This is a great example for how useful the new **container technology** can be for this use case: use the current release of your software within a container in the cloud (and the same container then also runs in-house on your workstation or cluster), and when a new release of this software comes out just wait a few weeks until this release has been containerized, tested, and evaluated, and then switch seamlessly to the new (containerized) release. Because these software containers are part of **a library of software releases** you can always go back to the previous release.

**Comment 63:** It is a while ago but Microsoft Research produced a crystal ball gazing document that attempted to predict trends in Science computing. http://research.microsoft.com/en-us/um/cambridge/projects/towards2020science/. It is not the easiest of reads but there were some relevant observation that would still seem to apply.

"I/O bandwidth has not kept pace with storage capacity. In the last decade [1996-2006], while capacity has grown more than 100-fold, **storage bandwidth has improved only about 10-fold**."

"We believe that attempts to solve the issues of scientific data management by building large, centralised, archival repositories are both dangerous and unworkable."

A message that I believe to be true is that the data repository and high-performance computing should be co-located with only subsets of processed data being transferred for visualisation. That could mean that one would not only need to look to the cloud for peak computing capacity but also build in the costs of using it as part of a long-term data repository.

**Response 63:** Thanks for the link to this famous Microsoft Research document which is still mostly valid. Indeed, a data source too far away from the compute resource and interconnected by (relatively) **slow Internet is a real problem**. For some **'work-arounds'** please see a few answers above, e.g. my comments to XYZ and others (box.com, remote viz, NICE DCV, FedEx, etc.).

**Comment 64:** These days, who even truly believes in the corporate security claims? Sad, but I have to believe a fair amount of our IP isn't just ours. So the bar is set kind of low.

**Response 64:** I agree, we are seeing a very unfortunate trend concerning the **secureness of our assets, caused by a lot of spying**, but also of **the carelessness of companies** underestimating this threat. In that sense, indeed, the security in Clouds is often better than in our own 'datacenter', see an interesting article here: http://www.businessspectator.com.au/article/2015/1/12/technology/clearing-skies-cloud-getting-safer.

**Comment 65:** Why not, just coincides with Internet spirit. While CAE simulation needs isn't and maybe won't be the same popular as Apple App Store, that everyone needs every day. And it is at most enterprise-class needs, most possible, before the big brothers in CAE fields turn into this model to release their products, small companies with great ideas have few chances to shake the market. Or maybe more interesting, will CAE applications go into home of ordinary people? That's the time small companies play the same role as the big guys.

**Response 65:** The trend we are seeing is that the **smaller companies are currently showing more interest in HPC Cloud than the big guys**. Big guys often have their own big HPC systems which SME's can't afford. With Cloud computing now, the SMEs have three computing options at their fingertips: **workstation, HPC Cluster, Cloud**. In this regard they are no longer behind the big guys which **strengthens their ability to innovate and compete**.

**Comment 66:** Wolfgang: As president of TheUberCloud I don't expect you to understand how the design/ simulation process works. I've worked for the software companies selling these applications as well as in the trenches actually using the stuff. There is a big disconnect in what the marketing guys think about how things work and what really goes on in the field. I work out of my home office and have a lot of internet reliability and bandwidth issues. The internet provider is not my friend and only provides one level of service, bad. I also work with a lot of big companies and know very well how they operate and the issues they have. The IT guys down the hall are the last people they want to deal with. The bottom line is most engineers just want to get their jobs done and control their own destiny. They don't want to deal with more IT guys bringing more people into the critical path. They don't want to add more process steps from multiple software companies and service providers who all seem to point fingers at the other guy when things don't work out. They want a single well integrated solution that works the same every time. I see the industry going the way of the CLOUD because it is easier for the software companies, but, it is not what the users want. **I've not seen one comment on this blog or the others where anybody thinks this is a good idea. If they do I would like to hear their success stories.**

**Response 66:** Thanks, certainly good points, indeed, but for me a bit one-sided, too pessimistic. I intentionally asked the question about ALL roadblocks being removed. And we are moving there. Sure, you are talking about bandwidth, but **I can't solve your last mile problem** (I would love to!). I guess you have seen my answers about the current (interim) solution: real-time high-res remote viz, intermediate storage on Box.com, and submitting the final result file via FedEx. Yes, engineers are doing this (**see the UberCloud Compendiums with 42 real engineering application case studies**).

You are talking about this topic as if it were a disruptive technology. It is not, because you will use your existing environment as if nothing happened. And only if you get into time or system limitation, in a sudden, you will be extremely happy that you can tap into resources seamlessly (no roadblocks as I suggested), just pushing a button on your screen, and paying for what you use.

The pressure to include Cloud into your computing tool box will come from several directions: from the engineers themselves because more and more they will understand the benefit of additional (big) resources, on demand. From the enterprise side of the house because they are doing Cloud already for years and will ask the question: "And what about you?" at the latest when you want a new compute server. From everybody else (including you and me) because whatever we do with our smart appliances we are doing it in the Cloud. I can see that this will be a matter of time, say 1 – 2 years.

**Comment 67:** Thank you for your reply. I've always embraced new technology and ideas when it improves the process. For the past 20-yrs I've given countless seminars on how to merge design and simulation. Back then, that was considered crazy talk, now it's almost main stream. I hope this cloud idea does go the way you say and it all works like it should, however,

you made one of my points in that you have introduced three new players into my business; the internet provider, Box.com and FedEx. This may be OK for the 10% of the time when I wish I had more capacity but how does this help 90% of the time. I don't think this is going in the right direction but I will look at your case studies and see if I'm enlightened.

I've worked for 3 FEA software companies and know many others in the business, so yes, I am a bit pessimistic. They all go by the 80/20 rule. Their goal is to satisfy 80% of the need. That last 20% does not have a good ROI. I can see that when 80% are up on the cloud, other solutions will not be supported. That's how business works. So yes I do think this is disruptive technology, and it's not new. 30-yrs ago we had to run all of our jobs on the big CRAY computer in the sky. We had to send our data over slow land line connections, wake up the SysOp (with a bell) at the other end of the line to mount our tape (when he had time) and then wait.... Make sure we never go back to that!

**Response 67:** ☺ I think with all the new opportunities coming along with Internet based services we will introduce even more players, always looking for the best suited most efficient service (solution, product, expertise). We as engineers are very used to this method, because we are using special tools when it comes to solve special problems, even at home, when I look e.g. at the tools collection in my little private workshop. So I think there is nothing wrong with it.

For example Box.com, or similarly Dropbox. Just one hour ago I uploaded my presentation slides for CICT'14 into my public file on Dropbox and gave a colleague at Sullivan University (who is organizing CICT Conference) access to this file: 56 MB because I added voice; impossible for sending it via email. My take: I use additional tools when I need to, and vice versa.

To your 80:20: I think the contrary will become true. Imagine you could package every data set, every software release, every modified customized something, including the other 20%, each into a Linux container (like Docker), then encrypt it, compress it, and store it into an image library, according to: built once – run anywhere – any time; then how much easier is it to maintain such a library for your customers. Yes, certainly, we have to build the infrastructure for this, but we are only 1 – 2 years away. For example, the Docker containers came out 1.5 years ago, and now there are already over 12.000 Docker projects in our community (as of Dec 2014); I find this speed and mass amazing; and it makes me optimistic.

Last but not least: I love your Cray example. But in 1980 (at the DLR in Gottingen) we first had to port 24 aerodynamics codes to this engineering marvel, until finally the reward was a speedup factor of between 5 and 25 compared to the IBM 360, which was disruptive enough to take this burden.

**Comment 68:** As I read this path I notice the little icon that tells me things are downloading to my page is still moving. That combined with the **security** problems insures that I would never trust critical data to the cloud. Could you imagine if Tesla or Ford stored there next gen plans for

development on the cloud and it was compromised? I think we have a way to go yet! This is an old technology with a new name. Don't be fooled by marketing!

**Response 68:** Yes, Tesla or Ford would never store their next gen plans in the cloud, nor would Boeing, nor Pratt & Whitney, . . . UPS – I just read the an "Ex-Pratt Worker Allegedly Tried To Ship F-35 Files To Iran". And another one: "Chinese Businessman Charged With Hacking Boeing". The big guys certainly have expensive security 'walls', and it still happens, and most crimes are done by their own employees. But what about the smaller companies? Boeing's suppliers, defense contractors. Do they have the money and expertise to defend themselves effectively? I doubt. Take Clouds: a cloud provider is very aware of the security concerns of his customers. He will do everything (e.g. hiring high-paid security experts) to avoid security breach and damage to his reputation (which is equal to death). And, please read this very important article about security breaches mostly done in in-house data centers, and not at all in clouds: http://www.businessspectator.com.au/article/2015/1/12/technology/clearing-skies-cloud-getting-safer.

I don't think cloud is marketing hype. Over 70% of American companies – according to IDC – are already using cloud service in their daily business. And over 90% of the American people like you and me, on Amazon, Apple, Dropbox, Google, Microsoft, and many 100s more.

And still, you are right. There are many applications and data which can easily be run in the cloud, because they are not really confidential. But others are, and they should stay within your firewall, AND surrounded by expensive security technology (which unfortunately today they are not).

**Comment 69:** I just noticed a BBC news report from a couple of days ago that Microsoft's Azure cloud had gone down, taking out critical software, services and websites from many companies across the world. Customers weren't able to inform Microsoft (or complain) because their email also went down. Right now I'm glad I've got good old Office 2010 with Outlook sitting on two laptops next to each other and another at home, with all my simulation software and data files on a RAID NAS and mirrored on another large backup drive :-)

**Response 69:** We all know that this happens with all technology in its early days, and especially if human beings are involved and make mistakes as in this case, so no surprise. For now, why not keeping your data on your RAID NAS, and sending only the few MB data files over to Abaqus sitting in a secure container in the cloud, ready to use, at your fingertips? And then sending the intermediate results via high-res remote viz (e.g. NICE DCV) back to your desk? This works today, flawlessly ☺

**Comment 70:** XYZ's comment reminds me of my personal use of the "cloud". I use TurboTax for my tax returns. One year I decided to use their online version. One day during a snow storm,

our cable TV and Internet went down. I thought "this is a perfect time to finish my tax return". But of course, with no Internet I could not do that. I went back to the desktop application the following year.

The "cloud" must always be available ("always" is pretty hard to accomplish) for it to gain widespread acceptance. ...or at a minimum, there must be a fairly useful offline mode for the application.

**Response 70:** In contrast to your tax return which has to be at IRS before the deadline (you might consider starting early ☺ ) there is usually a bit more project flexibility in sending your data file to the application in the cloud. And, **if that cloud goes down (very rare) you can send it to another cloud, and to another cloud, packaged securely and in a container, fully portable**. In contrast to your online TurboTax, your application software is already sitting on your workstation on your desk, and you can do a lot of preparatory work with it, including preparation for a big run in the cloud. And when time comes, you can send this fully prepared 10 MB file with your geometry and physics data to the cloud and perform 25 simulations in parallel, getting high-res remote viz pictures back and decide for the best configuration, in one twenty-fifth of the time for one cloud run, and perhaps in one hundredth of the time on your workstation.

Unfortunately, TurboTax is not parallelizable ☺

**Comment 71:** From my perspective this question is similar to the question some years ago, would you use mainframe or pc's? By the way using pc's leading you to build a mainframe with pc's. Ohhh, wait a minute, isn't that the cloud? So using the cloud I think is only a question of time, because everybody will use the cloud in the future if he/she wants or not. Twenty years ago the today's common computer was a big mainframe or even more it was a grid. I'm not sure if grid, cloud and mainframe will not be words for the same thing in the future :-)

As a conclusion the question if you want to use it or not should not be very important. The only question, from my perspective, should be: When, do we will use our own cloud, grid, mainframe as we do now with the mainframe's that were available some years ago?

But let's say that will happen in some years. The next question from my point of view can be for example: why should anybody move a CAE simulation to a cloud which will be available in some time on smaller hardware with less cost, more power and less energy consumption to a prototype?

OK, let's suppose again it is because you really need it now. The question about using a technology or not is not limited to the technology itself. **Trust, is from my perspective the reason why cloud does not work as well as selling computers**. Even if it is in fact the same thing. But only the fact that if you buy a computer gives you the illusion of "totally control" over the computer you do not need the new vector "trust". Trust is the main point with cloud form

my perspective. Most people which do make decisions are not computer scientists so they do not understand, that if you need an anti-virus, firewall etc. on your own pc you will never be in the position of "totally control" over the computer even if you hold it in your own hands.

In fact not understanding that holding something in your hands which is yours but you cannot "totally control" is a different paradigm in comparison with the rest of our daily lifes. I'm totally curious how long it takes for the people to understand and make the shift from the illusion to reality.

**Response 71: Using the Cloud is already happening today**; admittedly by a very small number of early adopters who have a real need for it or feel a real pain. I am more curious about when the early majority will jump on the Cloud. I feel this will be very soon. Then the Cloud will become a viable business also for the providers, who are currently mostly still investing in the future. Even that desktop technology will become faster and better all the time, that same technology will be in the Cloud as well, and while your desktop sticks on your desk for three years, the **Cloud technology refresh happens in much shorter cycles**. And your product design requirements will also increase and get more complex. What I am trying to say is, that **technology will advance, as will your requirements**; nothing new though, we know this for 40 years.

**Trust: This will happen over time, and it's again not Cloud specific**; it's always with new technologies; early adopters try it, help to improve it, start loving it, talk about it, and then the early majority is joining. One major contribution to this cycle we were able to provide are these 42 CAE Cloud use cases which we published in two compendiums (just search on Google) and in magazines like Desktop Engineering and HPCwire. **Reading about other engineers who succeeded in running their applications in the Cloud is slowly but continuously building trust**; definitely a very important element of the adoption.

Another contribution to **building trust are our new application software containers** who have additional **features like job monitoring, status reporting, email messaging, visualization**, and more **which gives you control** over your jobs running in the cloud, similar to the jobs on your desktop workstation.

## Reference

Workstations - Servers – Clouds: Comparing apples to apples. A decision-making support for the undecided. Wolfgang Gentzsch, The UberCloud, December 17, 2014: http://www.theubercloud.com/workstations-servers-clouds-comparing-apples-apples/

The UberCloud Compendium of case studies 2013 https://www.theubercloud.com/ubercloud-compendium-2013/, and 2014 https://www.theubercloud.com/ubercloud-compendium-2014/

Mark A. Lobo, P.E.; Jon den Hartog, P.E.; Derrek Cooper and Wolfgang Gentzsch, 2014: Testing simulations with cloud computing: http://www.deskeng.com/de/testing-simulation-cloud/

Michael Swinson: Clearing skies: Is the 'Cloud' getting safer, 2015: http://www.businessspectator.com.au/article/2015/1/12/technology/clearing-skies-cloud-getting-safer/

Wolfgang Gentzsch: How Cost Efficient is HPC in the Cloud? A Cost Model for In-House Versus In-Cloud High Performance Computing, 2014: http://www.theubercloud.com/cost/

Wolfgang Gentzsch: Take HPC to the cloud – in a container, in Scientific Computing World, 2014: http://www.scientific-computing.com/features/feature.php?feature_id=322

The UberCloud Marketplace: https://www.TheUberCloud.com/store

The UberCloud Experiment: https://www.TheUberCloud.com/hpc-experiment/

The OpenFOAM Foundation: Creating OpenFOAM Solvers: http://www.openfoam.org/features/creating-solvers.php

BMI Uses Jaguar to Overhaul Long-Haul Trucks, 2011: https://www.olcf.ornl.gov/2011/03/02/bmi-uses-jaguar-to-overhaul-long-haul-trucks/

Digital flow simulation: http://www.wirthresearch.com/about/development-in-the-digital-domain.html

Amdahl's Law: http://en.m.wikipedia.org/wiki/Amdahl%27s_law

Science 2020, Microsoft 2006: http://research.microsoft.com/en-us/um/cambridge/projects/towards2020science/