REPORT

ANSYS IN CLOUD – SCALING UP COMPUTATIONAL CAPACITY

OCTOBER 09, 2014

Disclaimer

Techila Technologies Ltd. disclaims any and all warranties, express, implied or statutory regarding this document or the use of thereof by you to the full extent permitted by law. Without limiting the generality of the foregoing, this document provided by Techila Technologies Ltd. in connection therewith are provided “as-is” and without warranties of any kind, including, without limitation, any warranties of performance or implied warranties of merchantability, fitness for a particular purpose, title and noninfringement. Further, Techila Technologies Ltd. does not make, and has not made, any presentation or warranty that the document is accurate, complete, reliable, current, error-free, or free from harmful information.

Limitation of Liability

In no event shall Techila Technologies Ltd. or any of its respective directors, officers, employees, or agents, be liable to you or any other person or entity, under any theory, including without limitation negligence, for damages of any kind arising from or related to the application of this document or any information, content, or materials in or accessible through this document, including, but not limited to, direct, indirect, actual, incidental, punitive, special or consequential damages, lost income, revenue or profits, lost or damaged data, or other commercial or economic loss, that result from your use of, or inability to use, this document, even if any of those persons or entities have been advised of the possibility of such damages or such damages are foreseeable.

Use of this document and copyright

No part of this document may be used, reproduced, modified, or transmitted in any form or means without the prior written permission of Techila Technologies. This document and the product it describes are considered protected by copyrights and other intellectual property rights according to the applicable laws.

Table of contents

1 Introduction....4
2 Motivation of This Study....5
3 Current Solution....6
3.1 Current Solution Architecture....6
3.2 Challenges of Current Architecture....7
4 Target Solution....8
4.1 Target Solution Architecture....8
4.2 Requirements for Target Solution Implementation....8
5 Cloud Alternatives....10
5.1 Initial Cloud Alternatives....10
5.2 Availability of Microsoft Windows Sperating System....10
5.3 Use of eCADFEM License....10
5.4 Cloud Platforms Included In the Technical Evaluation....11
6 ANSYS Performance Comparison....12
7 Cost Comparison....14
7.1 Instance Price Comparison....14
7.2 Techila Worker Deployment Time Comparison....14
7.3 Total Cost of Ownership Comparison....14
8 Conclusion....17

1 Introduction

Analysts following the HPC Market have reported frequently that software costs and complexity of cluster systems are major customer pain-points in HPC. A recent trend has also been that HPC customers have started to pay more attention on the return-on-investment (ROI) of their HPC systems (Source: IDC HPC Update at ISC'14).

Techila Technologies' discussions with customers have confirmed the reported HPC market trends. The pain caused by third party software licensing models, cost of HPC system ownership, and complex usability have been key topics, especially when talking with engineering simulation software users.

In this report, sponsored by MacGregor Finland, Techila Technologies investigates possibilities to optimize the cost of HPC system ownership with cloud services integrated to the enterprise IT. This report looks at the feasibility of computing clouds as a platform for engineering simulations. We will also study how to manage the software license costs related to ANSYS computing in a business, where the demand for resources can vary drastically from one time to another, and usability will be critical because all users are not experts in cloud or HPC technologies.

Building readiness to respond to this challenge requires either a HPC solution, which is designed to the average peak workload, or a new kind of HPC solution, which can provide the required elasticity by other means.

- Jouni Lehtinen, MacGregor Finland

I am excited to see the potential, which modern cloud platforms can offer to engineering simulations.

- Tuomas Eerola, Techila Technologies

2 Motivation of This Study

Cargotec's MacGregor is a leading brand of solutions and services for marine cargoes and offshore loads. In MacGregor's business, an optimized structure translates into optimized and flexible space for transported cargoes.

"The role of computing is critical in MacGregor's business. Nearly all marine structural components are custom designed for a specific application. Tendering stages are quick, and producing a price quote and an initial design proposal need support from the computing. The hectic nature of business causes fluctuating workloads to computing. Building readiness to respond to this challenge requires either a HPC solution, which is designed to the average peak workload, or a new kind of HPC solution, which can provide the required elasticity by other means," says Project Manager Jouni Lehtinen, MacGregor Finland.

Techila Technologies is a leading provider of High-Performance Computing (HPC) middleware solutions and one of the pioneers of cloud-powered HPC.

"Based on the experience of Techila Technologies, problems related to optimization of a large number of variables, sensitivity analysis, and scenario simulations are good examples of simulation areas, where large-scale distributed computing and cloud technologies can offer valuable benefits. Techila is a middleware solution for integration and orchestration of HPC resources. We are looking forward to understanding, how engineering simulation software users could benefit of cloud-based processing and Techila," says Vice President Tuomas Eerola, Techila Technologies.

MacGregor Dry Cargo engineering department and researchers from Tampere University of Technology studied in 2013, how MacGregor could benefit of the HPC middleware technology of Techila Technologies in their ANSYS computing. The solution developed enabled efficient optimization of a large number of design variables. (Source: "A Perfect Fit", ANSYS Advantage, Volume VII, Issue 2, 2013)

This document is a report from a study, where Techila Technologies investigated possibilities to benefit of cloud computing in MacGregor's ANSYS computing scenarios. In this study, Techila Technologies continued the work and combined it with the results of the cloud benchmark experiment, which Techila Technologies conducted in May 2014. (Source: Techila Technologies Report: Cloud Benchmark – Round 1) In this report we will look at the results from three key angles: IT architecture point-of-view, engineering simulation performance point-of-view, and usability point of view. The results will be compared to a reference non-cloud infrastructure running in a hosted datacenter.

This report uses some technical terminology, which is specific to the Techila technology. If you are not familiar with the Techila solution, please refer to Techila Fundamentals tutorial available on the web site at www.techilatechnologies.com.

3 Current Solution

3.1 Current Solution Architecture

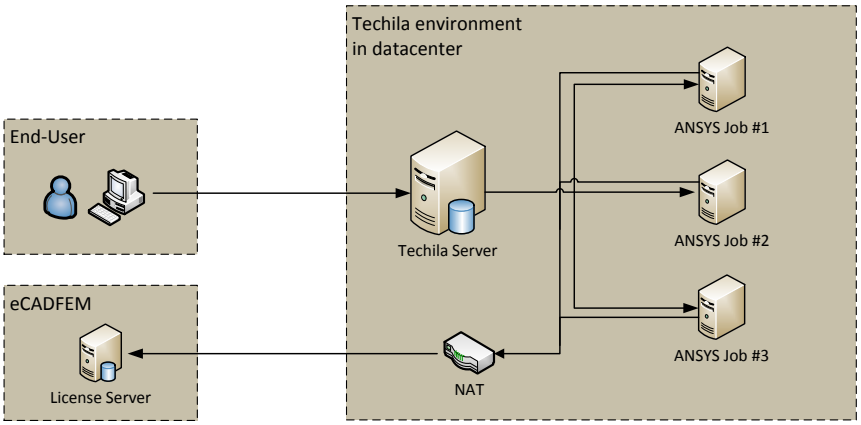

Operating System	Microsoft Windows Server 2008 R2
CPU	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz
Number of Cores	7

Operating System	Microsoft Windows Server 2008 R2
CPU	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz
Number of Cores	7

Operating System	Microsoft Windows Server 2008 R2
CPU	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz
Number of Cores	6

Figure 1. Current Solution

MacGregor Finland has a long history in outsourcing IT infrastructure services. MacGregor Finland's current solution architecture is built around a Microsoft Windows computing platform, which runs in an outsourced computing datacenter, and uses Techila HPC middleware to manage the computing resources. In this study, we will compare our findings from cloud-based environments to a reference non-cloud infrastructure, which is similar to MacGregor Finland's actual datacenter environment.

Common goals of IT infrastructure outsourcing include building a proactively managed IT infrastructure and a more agile and flexible foundation, empowering a flexible operating model to aid in achieving operational excellence. (Sources: 3 Key Advantages of Outsourcing Core Infrastructure Business Functions by Dean Schauer, VDN, Enterprises Seek The Benefits Of Hybrid Cloud, And Work To Overcome The Challenges by Forrester Research)

The purpose of Techila HPC middleware in the computing environment is to simplify the management of a distributed computing infrastructure, and to provide automatic configuration of compute nodes, and an automatic and instant recovery from unexpected situations, such as a node crash or network problem. (Source: The UberCloud HPC Experiment: Compendium of Case Studies 2014).

A traditional licensing model for engineering simulation software is an annual license, which defines the number of simultaneously running software instances. The software license can be installed on the user's own computer or the licenses can be made available from a license server running in the user's corporate network. Even if the IT infrastructure would be designed to support agility and flexibility, this licensing design can cause limitations for the system's abilities to respond to the needs of business.

Motivated by the results of the MacGregor Dry Cargo engineering department and Tampere University of Technology project in 2013, MacGregor started seeking for possibilities to find ANSYS licensing solutions, which can fully benefit of an agile and flexible IT infrastructure foundation. As a result of this process, MacGregor Finland selected eCADFEM's Engineering Software on Demand services (eCADFEM). eCADFEM supports many leading ANSYS products including ANSYS Multiphysics, ANSYS Mechanical and ANSYS CFD. (Source: eCADFEM Preise und Konditionen September 2014)

When coupled with Techila's autonomic computing features, the eCADFEM solution can enable easy scalability of the ANSYS computing environment to respond to the needs of business.

3.2 Challenges of Current Architecture

The role of computing is critical in MacGregor's business. Nearly all marine structural components are custom designed for a specific application. Tendering stages are quick, and producing a price quote and an initial design proposal need support from the computing. The hectic nature of business can cause computing workloads, which vary drastically from one time to another.

IT infrastructure outsourcing can increase the agility and flexibility of the IT foundation compared to internally managed solutions. However, the agility and flexibility comes with a price, which is often linked to the Service Level Agreement (SLA), length of service contracts, and required degree of flexibility.

4 Target Solution

4.1 Target Solution Architecture

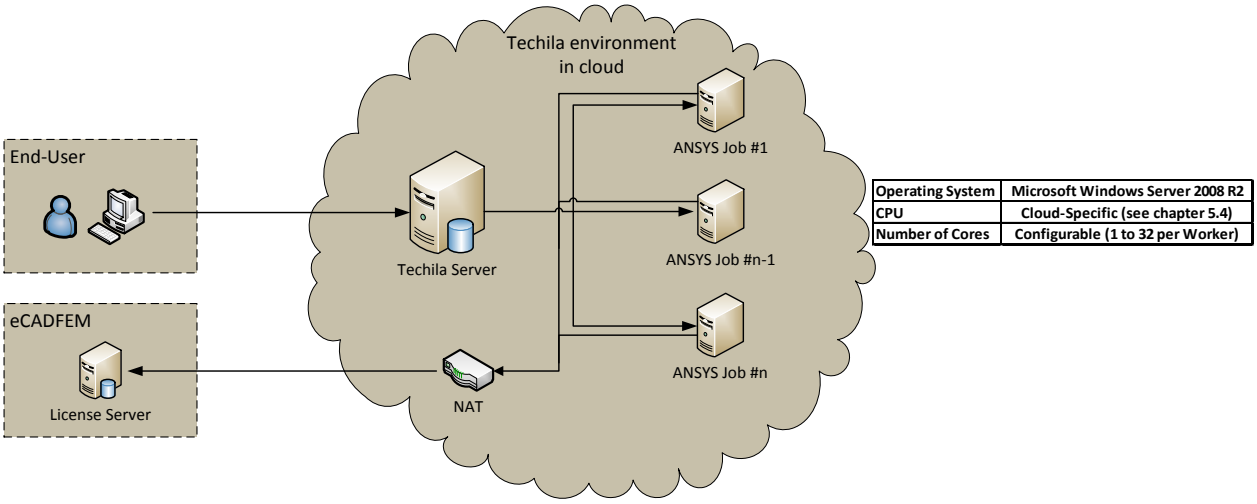

Figure 2. Target Solution

The role of computing is critical in MacGregor's business. In MacGregor's business case the demand for computing resources can vary drastically from one time to another. Meeting the highly fluctuating demands of business efficiently and at a reasonable cost, made MacGregor interested in studying possible benefits of computing on public cloud platforms.

Questions asked in the beginning of this study were: Would it be possible to integrate cloud-based computing infrastructures to the existing Microsoft Windows –based ANSYS solution? Could the leading public cloud platforms enable more efficient scaling of ANSYS computing resources than the current datacenter solution? Would it be possible to automate the capacity management processes for ANSYS throughput and link them to the business? Could the clouds offer finer granularity billing in ANSYS HPC than a hosted datacenter solution?

Techila Technologies has worked several years in co-operation with the leading cloud computing platform providers Amazon, Google and Microsoft. The Techila HPC middleware is productized for these cloud platforms, which made this group natural focus group for this study. This study was also seen as an interesting continuum for the cloud benchmark experiment, which Techila Technologies conducted in May 2014. During this first round of the cloud benchmark experiment, Techila Technologies observed that not all clouds are built similar. This study will show if there are differences, which might become visible in engineering simulations, particularly in ANSYS computing.

4.2 Requirements for Target Solution Implementation

Because MacGregor had a solution, which was proven and working in their current enterprise IT infrastructure, it was decided that changes resulting of the integration of cloud-based processing should be minimized. Based on that, the following requirements were set:

- Previous installations were made on Microsoft Windows systems, and MacGregor has good experience from Microsoft Windows, therefore it was agreed that the cloud platforms included in this study should support Microsoft Windows operating system
- MacGregor was satisfied with the on-demand licensing model for ANSYS tools, and it was agreed that it should be possible to benefit of this licensing model in the cloud, too

When defining the scope of this study, Techila Technologies' team raised the questions: What if the organization has an existing ANSYS license manager in their on-premises network? Could they benefit of these existing licenses with cloud-based processing?

Some cloud platforms support setting up so-called Virtual Private Clouds (VPC), where cloud resources are seamlessly integrated to the enterprise IT. The VPC isolates the cloud computing resources from the rest of the cloud and the Internet.

If an organization has an existing ANSYS license manager in their on-premises network, there are no technical limitations, why the cloud-based computing infrastructure could not benefit of these licenses. It was decided that a more detailed analysis of a VPC-based system would not be included in the scope of this study, and the focus of this study should be put on use of eCADFEM.

5 Cloud Alternatives

5.1 Initial Cloud Alternatives

The initial list of cloud platforms included in the scope of this study contained the three most famous competitors: Amazon, Microsoft and Google. The Techila HPC middleware, which is used also to manage MacGregor's on-premises system, offers a fully productized support for all of these clouds. (Source: Amazon, Google Cloud Price War Benefits CIOs by Steven Norton and Clint Boulton, CIO Journal)

In this report we will use following abbreviated names for the cloud services:

Abbreviated name	Full name
AWS	Amazon Elastic Compute Cloud EC2
Azure	Microsoft Azure
GCE	Google Compute Engine

5.2 Availability of Microsoft Windows Sperating System

Cloud platform	Supports Microsoft Windows operating system
Azure	Yes
AWS	Yes
GCE	Yes[1]
[1] **In developer preview mode**	

Table 1. Availability of Microsoft Windows operating system

5.3 Use of eCADFEM License

In this study, we will focus on use of eCADFEM in cloud-based computing environments.

Use of eCADFEM requires that the virtual machines in the cloud have an ANSYS installation, and are able to establish a network connection to the eCADFEM license server in the Internet. The virtual machines will use this connection to verify the availability of an applicable ANSYS license for performing the requested ANSYS operations.

The eCADFEM service is linked to the IP addresses. Only license request made from valid IP addresses will be successful. IP addresses are valid, if they are included in the list maintained by eCADFEM. The customer must send the IP addresses to eCADFEM prior to usage. The simplest way to manage this requirement is if the cloud platform supports reserving static IP addresses for the virtual machines. This can be made even easier by using NAT (Network Address Translation) which makes all the virtual machines appear to be connecting from a single IP address. All the cloud platforms included in the scope of this study offer support for NAT and some of them use it automatically for all the virtual machines.

The service's requirement for the IP addresses will set requirements to the cloud computing platform and defines also the way how integration of cloud-based processing to the enterprise IT needs to be done. Requirements for the integration may be different if using a VPC solution.

Public cloud platform	Use a static IP address
Azure	Yes
AWS	Yes
CGE	Yes

Table 2. Static IP address in the cloud.

5.4 Cloud Platforms Included In the Technical Evaluation

As a result of the cloud platforms' readiness to support the technical requirements, the scope of the technical performance comparison was limited to following public cloud platforms. GCE was left out because its support for Microsoft Windows operating system has not yet ready for general availability level.

- Azure
- AWS

Based on the results of the cloud benchmark experiment, the following virtual machine types were selected for further analysis. In cloud-based environments, virtual machines are commonly called "instances". Two different instance types were selected from Azure because of the significant differences in their CPU architectures and price levels. (Source: Techila Technologies Report: Cloud Benchmark – Round 1)

Instance Type	CPU Architecture
AWS (c3.8xlarge)	Intel Xeon E5-2680 v2 @2.80GHz
Azure (A9)	Intel Xeon E5-2670 @2.60GHz
Azure (A4)	AMD Opteron 4171HE

Table 3. CPU Architectures of Cloud instance types

6 ANSYS Performance Comparison

This Chapter contains the results of the ANSYS performance comparison for the following public cloud platforms:

- Azure
- AWS

The tests performed to collect data for comparing execution speeds in different infrastructures were implemented by launching an ANSYS Mechanical optimization process, which was delivered by MacGregor. The ANSYS optimization process was executed identical in the different infrastructure environments.

Moving the process from one infrastructure to another without changes was enabled by the Techila HPC middleware. The Techila HPC middleware provides the application layer with an abstraction, which removes lock-in to a specific infrastructure.

In this document, we will call each execution of a test run a Project. In Techila terminology a Project is a computational process, which is started by an End-User or application. A Project can consist of one or several Jobs. A Job is a task, which can be executed independently, and contributes to the goals of the Project.

In the ANSYS performance tests described here the Techila HPC middleware executed the ANSYS optimization process as a Project, where each Job consisted of one ANSYS solver run.

The performances of different computing infrastructures were measured by analyzing the wall clock times spent in each Job.

Because of the internal design of the ANSYS solvers, the ANSYS Jobs used in the tests of this study were able to use multithreading. In practice, this means that each Job is able to divide the computational load onto more than one CPU core. By using multithreading the wall clock time of each Job can be reduced by reserving more CPU cores for the Job on a Worker computer. “Worker” is a term used for a compute node in a Techila environment.

The use of multithreading does not decrease the total CPU time used for the Project, but it can decrease the wall clock times of the Jobs. This also decreases the wall clock time used for the license reservations and therefore it can decrease the ANSYS license costs.

For example, multithreading could be taken into use by assigning 8 ANSYS Jobs on a Worker computer, which has 16 CPU cores. This would allow each Job to use up to 2 CPU cores worth of computing capacity on the Worker.

Environment	Wall clock time per ANSYS Job / (Number of ANSYS Jobs equal to Worker computer’s CPU core count)	Wall clock time per ANSYS Job / (Number of ANSYS Jobs ½ of Worker computer’s CPU core count)
Hosted Datacenter (reference) [1]	50,78	31,00
AWS (c3.8xlarge)	65,68	36,59
Azure (A9)	55,03	34,18
Azure (A4)	146,69	88,66

[1] See **Figure 1** for instance details

Table 4. Execution times of ANSYS Jobs in different infrastructures

When comparing the performance data measured from the public cloud infrastructures with the performance data measured from the reference non-cloud infrastructure running in a hosted datacenter, we can see that ANSYS Jobs require more wall clock time in the included public cloud infrastructures than in the reference infrastructure.

AWS c3.8xlarge cloud instances and Azure A9 cloud instances are designed for computational workloads. When comparing the performance of these instance types and the reference infrastructure, we can see that if the cloud instance type is designed specifically for computational workloads, the difference in wall clock time per Job is relatively small.

Azure A4 is an instance type, which has been on the market for a longer period of time, and comes with an affordable pricing. However, Azure A4 is not specifically designed for computational workloads. When comparing the performance of Azure A4 with the reference infrastructure, we can see that there is a significant difference in the wall clock times per ANSYS Job. As reported in Techila Technologies' Report: Cloud Benchmark – Round 1,, a similar difference was seen in model calibration computations executed in MATLAB. This was discussed to be potentially caused by a different CPU Architecture (AMD Opteron.)

7 Cost Comparison

7.1 Instance Price Comparison

When considering the cost of cloud-based infrastructure as a part of the enterprise IT, it is important to notice the cost structure of cloud-based services.

Generally, we can say that cloud costs start when the resources are requested and end when they are released. The actual point of time when the invoicing starts and ends depends on the cloud platform in question. From an engineering simulation software user's point of view this means that he will pay for the actual usable computing time, but also for the time required to start and prepare the compute nodes to support the engineering simulations.

The different billing granularities of Azure (1 minute) and AWS (1 hour) will affect the total amount payable, too. If the business requires an infrastructure, which enables catering HPC for short discrete Projects, a finer billing granularity can be beneficial. On the other hand, if the business requires a solution, which includes a cloud infrastructure that runs for long periods, the role of billing granularity will reduce.

In contrast to the public cloud platforms, the SLA of hosted datacenter services include often monthly pricing and they can define minimum contract periods for different services.

Techila Technologies analyzed provisioning of cloud-based services for computing in Techila Technologies' Report: Cloud Benchmark – Round 1, published in May 2014. For examples about cloud pricing specifications and observations done in the course of the benchmark experiment, please refer to the Cloud Benchmark Report.

7.2 Techila Worker Deployment Time Comparison

In the tests of this study, the Worker nodes in AWS and Azure were started and stopped using automated IT system management tools, which are included in the Techila solution.

Once started, they received a custom-made Windows image, which included the Techila Worker software and the ANSYS software. When the solution is transferred into production, the entire deployment process can be automated.

When the ANSYS process consists of independently executable ANSYS Jobs, such as an optimization process, sensitivity analysis, or a parametric sweep simulation, more Worker nodes translates to more computing throughput. This throughput can be used to serve one or several simultaneous users.

Below table contains average deployment times for Workers deployed in AWS and Azure.

In contrast to AWS and Azure, starting more capacity in a datacenter service the SLA can require issuing a request to the datacenter service provider. In many cases this process includes human intervention and the response times are often measured in days rather than minutes.

Environment	Worker Deployment Average Time
AWS (c3.8xlarge instances)	5 min
Azure (A4 instances)	9 min

Table 5. Worker deployment times

7.3 Total Cost of Ownership Comparison

The cost of a cloud infrastructure looks often very affordable compared to an internally managed infrastructure or a hosted infrastructure. This benefit gets highlighted especially when the infrastructure

needs to support business, which causes highly fluctuating demands for computing. AWS and Azure are designed to support scalability and to benefit of economies of scale.

As noticed in Chapter 7.1, when designing a cloud solution, we need to acknowledge the time required for starting cloud-based services. This is important both for responding to the requirements of business, and for estimating our Total Cost of Ownership (TCO).

When designing a computing solution, it is good to remember that the cost of infrastructure does not give the full picture of the TCO. When looking at engineering simulation software, software licensing plays a significant role in the TCO, too. Because of this, it is important to understand the requirements of business and design the infrastructure and licensing solutions accordingly.

Because of this, there will not be one answer to the TCO, but several answers, which depends on the usage patterns of the business. When we know our computing usage patterns, we can design the cloud-based infrastructure and software licensing in a way, which provides us the best Return On Investment (ROI).

Below table and figure include Techila Technologies' estimated TCO of ANSYS computing in different clouds. The costs presented are only indicative estimates. Techila Technologies recommends every customer to estimate his/ her own TCO based on his/ her own business requirements.

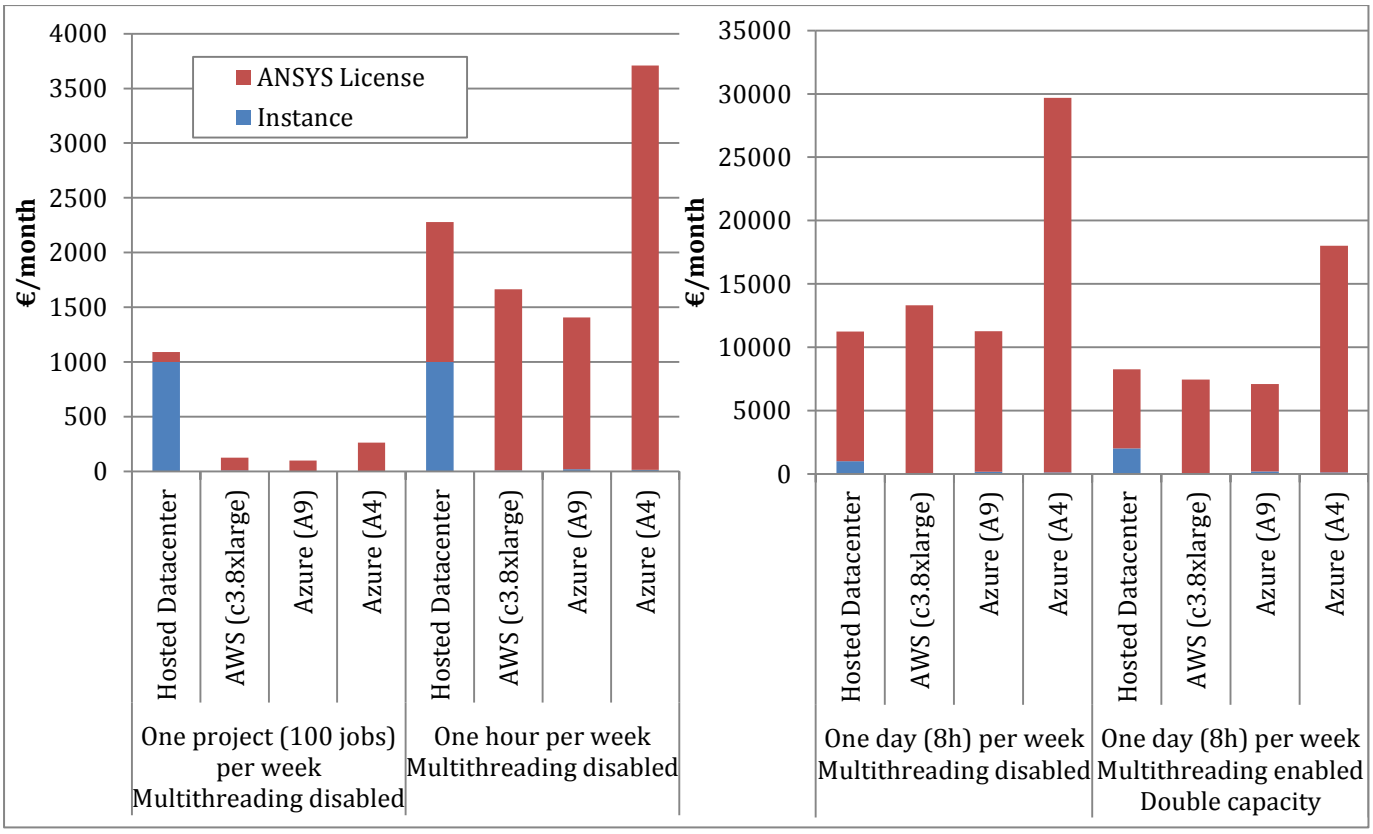

Figure 3. TCO of ANSYS computing. Notice different scales. Please see Table 6 for exact numerical values.

	Environment	Instance Cost €/month	ANSYS License Cost €/month	Total Cost €/month
One project (100 jobs) per week	Hosted Datacenter	1000,00	90,00	1090
	AWS (c3.8xlarge)	9,00	116,76	126
	Azure (A9)	1,39	97,83	99
	Azure (A4)	0,94	260,78	262
One hour per week (on Reference)	Hosted Datacenter	1000,00	1280,00	2280
	AWS (c3.8xlarge)	9,00	1655,58	1665
	Azure (A9)	19,78	1387,13	1407
	Azure (A4)	13,29	3697,58	3711
One day (8h) per week (on Reference)	Hosted Datacenter	1000,00	10240,00	11240
	AWS (c3.8xlarge)	63,00	13244,65	13308
	Azure (A9)	158,22	11097,03	11255
	Azure (A4)	106,31	29580,65	29687
One day (8h) per week (on Reference) Using multithreading with double capacity	Hosted Datacenter	2000,00	6251,28	8251
	AWS (c3.8xlarge)	72,00	7378,53	7451
	Azure (A9)	196,55	6892,54	7089
	Azure (A4)	128,50	17878,66	18007

Table 6. TCO of ANSYS computing. These numerical values are visualised in Figure 3.

Based on the tests performed in this study, we can say that considering high-performance cloud instances can be useful in many scenarios. In this study, we were using eCADFEM for ANSYS licenses. The current eCADFEM licenses use a wall clock time-based billing model. Even if Azure A4 instances are affordable, the ANSYS Jobs take so much longer to complete on Azure A4 than on the higher performance instances, that the growing ANSYS license costs will eat the benefits of affordable infrastructure.

In the course of this study there was also discussion about building a solution, where there would be a pool of traditional licenses to serve the needs of daily computing and possibility to activate eCADFEM for high-demand peak periods. Because the traditional licenses are not linked to wall clock time, in some scenarios this could also enable benefiting of lower cost cloud instances.

8 Conclusion

This study showed that there are scenarios where ANSYS users can improve the TCO of their HPC systems by integrating cloud services to their enterprise IT. These users include businesses, which have highly variable computing workloads in parallelizable computing problems, such as optimization of large number of variables, sensitivity analysis, and parametric sweep simulations.

“This study will help MacGregor Finland in defining the next-generation architecture for engineering simulation platforms. I believe that cloud computing technologies will play an increasing role in the future of simulation driven engineering. Usability and integration solutions, such as the Techila HPC middleware have quickly made cloud a viable option in business HPC. With Techila, the benefits of cloud can be integrated to applications and processes in a way which does not cause complexity to the IT or users,” says Project Manager Jouni Lehtinen, MacGregor Finland.

When we compared the pricing and availability of AWS and Azure services and an outsourced IT infrastructure, we noticed that the leading public cloud providers are able to offer computing resources at attractive rates and SLAs, which are designed to support even higher level of agility and flexibility than traditional IT outsourcing.

We noticed that the ideal solution depends on the requirements of the business. We noticed that if the business demand for resources varies drastically from one time to another, the SLAs of traditional IT outsourcing are not always able to support these requirements. In these scenarios, the benefits of AWS and Azure in ANSYS computing get amplified, especially when combined with a usage-based ANSYS licensing solution, such as the eCADFEM service.

On the other hand, we noticed that despite of the public cloud providers’ economies of scale, hosted datacenter services can become an interesting solution, if the ANSYS computing demand of business is relatively constant and if the primary pain-points in on-premises solutions reside around floor space and system administration. In other words, based on what we learned in this study, we believe that a successful IT architecture design is business driven.

When performing tests related to this study, the project team noticed that the usability of a cloud-powered ANSYS computing solution can benefit significantly of a solution which orchestrates the enablement of cloud-based processing. An orchestration solution can simplify the IT management and end-user experience which shows as better productivity.

In this study, we used the Techila HPC middleware, which is an autonomic computing solution. When the tests were started, the cloud nodes were automatically configured to support the ANSYS computing, and if there were any issues during the simulations, Techila made the process “self-heal”. The experiences from this study show that autonomic computing features, which orchestrate the system without human intervention can enhance the ROI of the solution a lot. Simplified usability can make computing available also to those users who are not experts in cloud or HPC technologies. The wider the benefits are, the bigger the ROI potential for the business will be.

“Techila Technologies thanks MacGregor Finland for their co-operation in this project. With the help of MacGregor we were able to benchmark the leading cloud platforms in a real ANSYS business scenario from the industry. I am excited to see the potential, which modern cloud platforms can offer to engineering simulations. Based on the findings of this study, we can say that cloud computing has rapidly become a true alternative to on-premises infrastructures and hosted datacenter solutions in engineering simulations. We hope that the ANSYS community will find the observations presented in this report interesting and useful,” says Vice President Tuomas Eerola, Techila Technologies.

In the course of this project, Techila Technologies’ team noticed that a potentially interesting topic, which was left out from the scope of this study is use of a so-called hybrid infrastructures in ANSYS computing. A hybrid infrastructure would consist of resources which are operated in an on-premises datacenter or in a hosted datacenter, and could be extended using the Techila HPC middleware with AWS or Azure resources on demand. This included also discussion about benefiting of an existing ANSYS license server

in the on-premises network outside high-demand periods. Based on the experience, which Techila Technologies has from other similar implementations for example in MATLAB scenarios, the technical implementation of this solution is realistic. If your business could benefit of this solution and you are interested in continuing this study with an analysis in a hybrid infrastructure, please contact Techila Technologies at info@techilatechnologies.com.