engineering.com

Cloud Computing for Engineering Simulation

intel

TotalCAE

This report is submitted by TotalCAE.

TABLE OF CONTENTS

INTRODUCTION

Engineering simulation has become essential to product development allowing companies to reduce or eliminate costly and time-consuming physical prototypes, deliver optimized product performance, and improve product robustness.

Ever-increasing competition in the marketplace is creating additional pressures to reduce product development time while increasing product quality. Engineers are responding with more simulations than ever before, which need to be done faster and with higher fidelity. Reducing simulation turnaround time is critical to successful, on-time product launches.

The speed of simulation depends on the speed of the workstation or server running the simulation software. To address the computation time challenge, many firms have tried to buy bigger workstations. Yet even the most powerful deskside workstations do not offer enough compute capability to achieve the required turnaround time on large projects. To gain additional compute capacity engineering firms are adopting cloud-based High-Performance Computing (HPC) which offers virtually unlimited compute capacity.

TotalCAE is a leading solution provider of managed HPC on the cloud for engineers and has helped hundreds of clients adopt HPC to solve their most complex engineering challenges while reducing their time to market. TotalCAE has partnered with AWS Cloud as its preferred HPC cloud infrastructure provider because of their depth of services, innovation in HPC, reliable global infrastructure, and variety of Intel-based compute instances suitable for simulation and other Computer-Aided Engineering (CAE) workloads.

This white paper explains how TotalCAE makes it easy for organizations to utilize cloud computing on AWS to accelerate product innovation, solve complex engineering problems, and reduce the time waiting for CAE simulation results.

WHY USE CLOUD COMPUTING FOR CAE?

Cloud computing offers engineers access to virtually unlimited compute capacity for scalable HPC. This in turn can be leveraged to reduce turnaround time from weeks to hours. Engineers can access a wide variety of Intel-based HPC compute machines (instances) on-demand with minimal investment or commitment.

The gain in computing power from cloud computing allows engineers to run larger, more complex simulations with greater fidelity. A recent study by Ansys found that engineers were reducing the fidelity and detail of their models more than 50% of the time to meet required turnaround times. With access to unlimited computing power, engineers can not only get their existing model results back faster but also get their results with higher fidelity.

This combination of scalability and flexibility enables clients to do more simulations faster without being constrained by the limits of fixed on-premises resources.

The benefits don't end with scalability and flexibility. There is no guessing about capacity; engineers always have the computing power required without needing to go through a lengthy procurement cycle. The compute resources are globally available, which makes it easier to collaborate or to monitor simulation jobs anytime, anywhere. Finally, small businesses have access to the same HPC resources as the largest companies to level the competitive playing field.

CLOUD SIMULATION PLATFORM PROCUREMENT

When an engineering team decides to move to cloud HPC, the first question is, “How do we get there?” As with many business advancements, the issue boils down to if the company should use internal IT resources to build a simulation environment or use a turnkey Software-as-a-Service (SaaS) provider like TotalCAE.

TotalCAE offers a managed simulation platform running on AWS, the leader in cloud infrastructure and platform services. AWS offers data centers around the globe, each securely hosting reliable cloud computing services including Amazon Elastic Compute Cloud (EC2), which provides secure, resizable compute capacity in the cloud. Amazon EC2 instances, powered by Intel® Xeon® Scalable processors, have the largest breadth, global reach, and availability of compute instances across AWS geographies. TotalCAE leverages the global scale of AWS and its HPC related services technology to deliver a turnkey simulation platform to clients around the globe.

BUILDING AN HPC SIMULATION ENVIRONMENT FOR ENGINEERS ON THE CLOUD

Building a custom HPC Cloud simulation environment is a complex undertaking. Public cloud providers such as AWS offer a broad range of components to develop an HPC platform, but these services must be assembled into a usable system. There are very few IT teams with the experience to develop a reliable turnkey SaaS HPC environment for CAE from scratch. It is an alphabet soup of things to learn such as:

1. Programming experience (Python, YAML, Terraform, JSON, API’s, Git, SQL).
2. HPC experience (compute, networking, storage, orchestration, automation, data management, data transfer, security and remote visualization).
3. Expertise in Linux systems administration.
4. Knowledge of AWS core concepts and services.
5. CAE application experience (applications like Ansys; how to run distributed; MPI; HPC support; RSM; Workbench integration; using the native HPC fabric, application aware monitoring).

For most clients, the time building and managing their own environment in the cloud to run their engineering applications is an opportunity cost against the real work of design and development. One client that approached TotalCAE had already spent over eight months and multiple consultants without success. As a result, this client was highly skeptical of using the cloud for CAE workloads. After engaging with TotalCAE the customer was up and running production workloads in one day. TotalCAE believes engineering teams should build the products that differentiate the company, and buy all else.

“Your engineers can use a production-ready cloud platform like TotalCAE for simulation, or spend time building the same tools,” says Rod Mach, CEO and founder of TotalCAE.

Engineering teams can use TotalCAE’s turnkey managed cloud platform to take advantage of HPC for CAE without either the engineering or IT departments dealing with cloud deployment and management complexity.

HOW TOTALCAE ON AWS WORKS

To get started, the client chooses either a dedicated AWS account hosted by TotalCAE, or an AWS account owned by the client. Regardless of which account hosts the environment, TotalCAE fully manages the deployment and ongoing management; little or no client involvement is needed.

Once TotalCAE sets up the service, users log in to a client-specific URL or via the command line to access their engineering simulation programs of their choice. TotalCAE will install software to access the clients existing software licenses, which are shared between on-premises and the cloud, so clients do not need to purchase separate applications licenses for cloud use. TotalCAE can also host your license servers in the cloud if the client does not have on-premises license server infrastructure. TotalCAE supports over 100 engineering applications, enabling clients to run almost any CAE application able to benefit from cloud computing resources. If the client has a custom or unique collection of libraries, applications, or home-grown applications, TotalCAE works with the client to integrate these tools into their environment.

WHY USE TOTALCAE ON AWS

TotalCAE on AWS makes it easy to run a CAE simulation in just a few clicks. The process entails the following three steps:

1. Upload the model.
2. Select the number of CPUs.
3. Click on “Submit” to solve.

It’s that simple. Projects can be run interactively through the use of remote visualization or in batch mode. Users can choose to do pre- and post-processing on the cloud or locally on their workstation.

TotalCAE on AWS offers seven specific advantages to their customers: ease of use, cost management, license flexibility, hundreds of integrated applications, application-specific monitoring, data management tools, and one telephone number for support.

EASE OF USE

> *“TotalCAE is simple to use, straight-forward and enables us to focus on engineering”*
>
> *—Guang Dong*
> *CAE Crash Lead, Lucid Motors*

Lucid Motors uses TotalCAE to optimize quality, cost, and development of their new long-range luxury EV, the Lucid Air. Guang noted how the TotalCAE portal was easy to use for job submission, enabling them to focus on engineering.

Here is an example screenshot of the TotalCAE portal with several popular applications Engineers can submit jobs for analysis on AWS in just a few minutes without being burdened by configuring and setting up their own cloud environment and associated engineering applications.

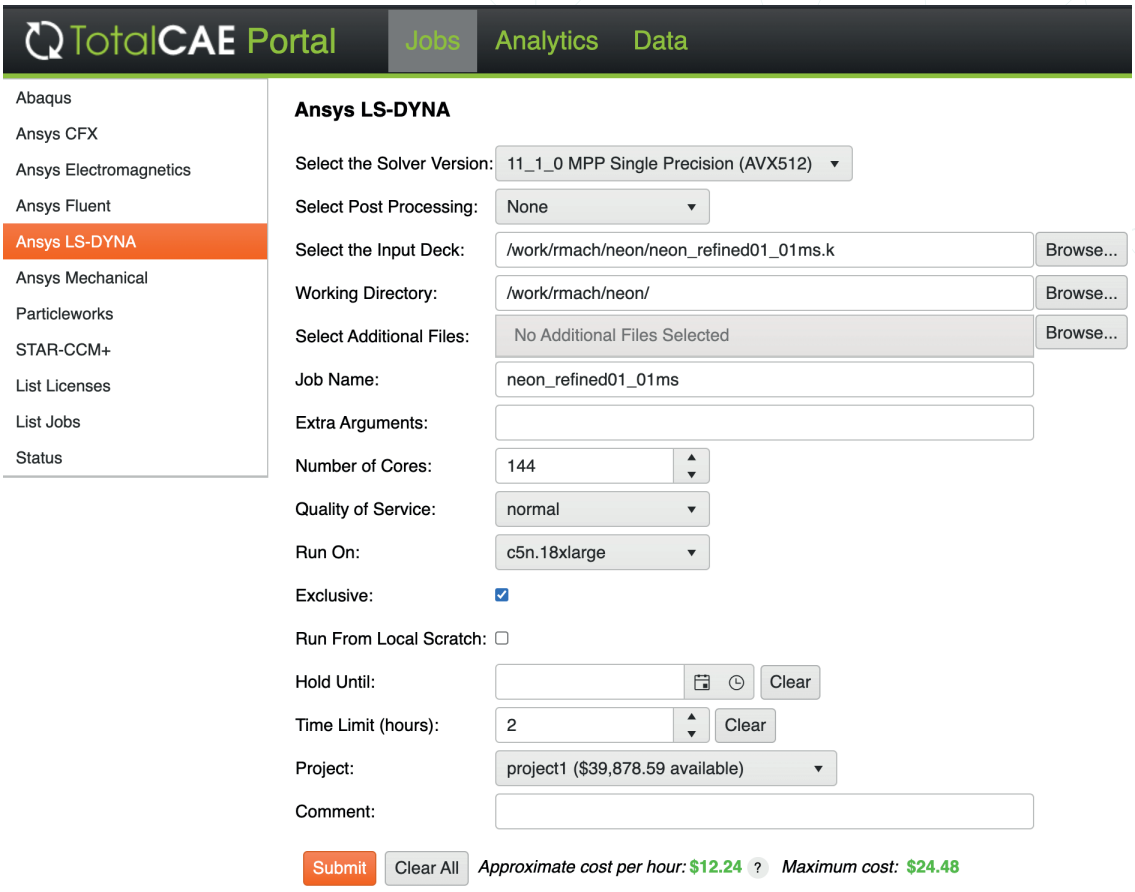

Figure 1: TotalCAE Portal.

COST MANAGEMENT

TotalCAE helps clients gain the best utilization of their simulation budget, monitor project costs, and avoid costly budget overruns that can occur with pay-as-you-go pricing for cloud computing. For example, TotalCAE offers a differentiated billing structure that estimates, manages, and reports cloud costs on a per-simulation basis. By understanding what a simulation costs, engineers can make more effective use of the simulation budget.

Clients use TotalCAE's cost controls to set maximum spending by job or project to manage the simulation budget more efficiently. Users can see available project budgets and are warned via email when projects are close to going over budget. TotalCAE billing reports are exported to Microsoft Excel and include all job details, allowing the CAE manager to see how simulation dollars were spent.

TotalCAE also makes it easy for clients to optimize for cost using various AWS compute resource consumption models. These models include On-Demand Instances, Reserved Instances, and Spot Instances.

With On-Demand Instances, you pay for compute capacity by the second with no long-term commitments. You have full control over when to start and stop your simulation and pay for what you use. On-Demand Instances are recommended for applications with short-term, irregular workloads that cannot be interrupted.

AWS offers significant savings over On-Demand Instances through AWS Savings Plans, Spot Instances, and Reserved Instances. Reserved Instances are not physical instances, but rather a billing discount applied to the use of On-Demand Instances in your account. Savings Plans also offer significant savings on Amazon EC2 costs compared to On-Demand Instance pricing. With Savings Plans, clients make a commitment to a consistent usage amount. This provides companies with the flexibility to use the instance configurations that best meet their needs and continue to save money.

Spot Instances are often the most inexpensive way to run on the Cloud with no commitment. Spot Instances are unused compute capacity available at up to a 90% discount compared to On-Demand prices. The Spot pricing fluctuates and is determined by supply and demand for Amazon EC2 spare capacity. If the Spot capacity is available, the job will run at a significant discount compared to the On-Demand pricing.

Examples of jobs good for Spot Instances include shorter running jobs that aren't time sensitive or can easily be checkpointed. TotalCAE enables you to automatically fall back to On-Demand after a predetermined amount of time waiting for Spot Instances, to make sure the job starts when you need it, if Spot capacity is not available.

LICENSE FLEXIBILITY – SAVE MONEY

On-premises HPC solutions often exist in concert with cloud-based solutions. TotalCAE enables clients to utilize their CAE licensing on-premises or in the cloud, ensuring maximum utilization of CAE licensing. Clients don't have to make changes to their existing license servers to work on the cloud in most cases. TotalCAE will also manage the CAE license server to alleviate the burden of updating and managing CAE licensing.

It is important to understand CAE licensing requirements when getting ready to run CAE applications on the cloud. The CAE application licenses should allow users to solve the model on multiple cores, more than can be done with existing local resources. Otherwise, the team will likely not realize the benefit from virtually unlimited cloud computing resources.

TotalCAE includes built-in analytics on all CAE license usage, both on-premises and in the cloud, for better insights into the amount of CAE licensing being used. This is helpful when it comes time to renew your CAE application licensing, as the company has insight if limited CAE licensing is the bottleneck to more engineering throughput.

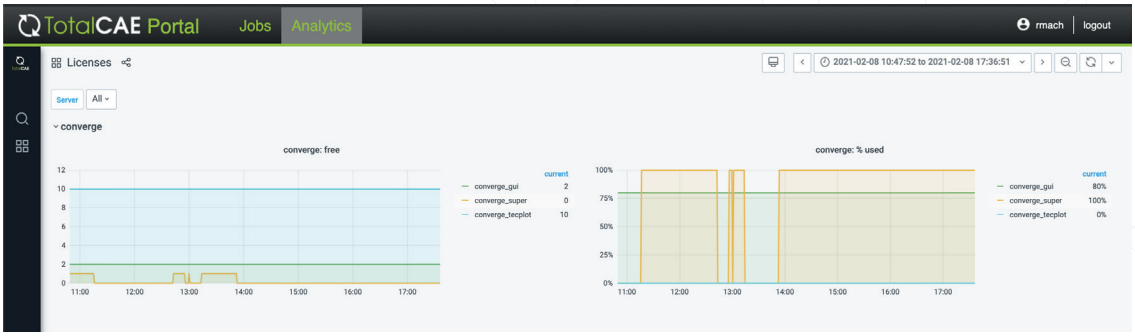

Figure 2: TotalCAE Analyzer Reports on Cloud Hardware, Storage and CAE License Usage.

A typical misconception among engineering firms is that existing CAE licenses will allow them to run anywhere. An on-premises-only license used for workstations may not work for the cloud without purchasing different licensing, as many CAE applications are license limited by CPU core, or node locked to a specific computer or range of computers. It is also possible a company may be prohibited by the End User License Agreement (EULA) to run on third-party systems and will need to update their license agreement to use cloud resources. TotalCAE works with clients to help them avoid common licensing pitfalls when moving to a cloud computing environment.

HUNDREDS OF PRE-INTEGRATED APPLICATIONS – SAVE TIME

Several ISV providers such as Ansys and Dassault Systèmes offer their own cloud solutions. TotalCAE's differentiator is enabling users to utilize any application in their own AWS accounts with TotalCAE's large application catalog.

Engineering teams typically use multiple engineering simulation software products. Using a cloud solution like TotalCAE, which enables users to use any of their CAE applications using the same workflow, saves engineers time compared to juggling multiple cloud solutions separately to get their work done.

APPLICATION SPECIFIC MONITORING – SAVE TIME AND MONEY

One common engineering challenge is getting models to solve properly. Engineers are well aware that models have many potential error sources, including incorrect boundary conditions, numerical error, meshing problems, and other issues that can prevent the software from converging on a solution. It is difficult for an engineer to monitor a simulation around-the-clock for potential issues that might cause them to miss an important deadline from a failed simulation.

To address this challenge, TotalCAE created TotalCAE Watchdog. TotalCAE Watchdog is a tool which automatically monitors solution progress to identify errors in real time. For example, if an application has not written any data in a few hours, Watchdog would alert the user. Under regular operation, the application would write output progress periodically. If the simulation

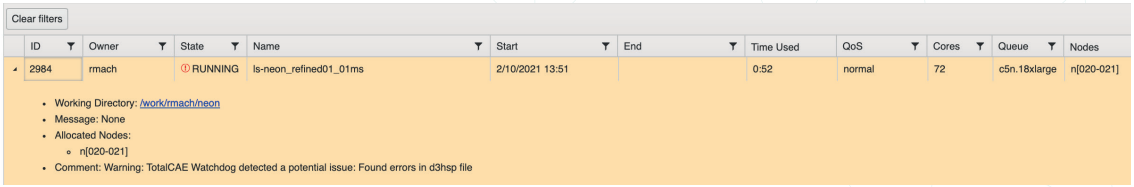

Figure 3: TotalCAE Watchdog.

were stuck in an infinite loop, Watchdog would recognize this and email the engineer, allowing them to either ignore the alert or cancel the run to save time and money.

Though Watchdog can't know every possible error, it can identify issues it has seen over the last twenty years of supporting a wide variety of CAE applications. TotalCAE also partners with its clients' engineers to customize alert emails to match model specific issues the client is interested in detecting. The sooner an engineer is aware of an error, the quicker they can get to the desired result. This is an additional way TotalCAE helps users save money on pay-as-you go cloud-based simulation.

DATA MANAGEMENT – SAVE TIME

Some engineering applications require a significant amount of time to prepare the model to send to the cloud. For example, dealing with model 'include files' is a common source of pain when trying to upload models to solve on a cloud system.

"One potential client was spending hours per day just preparing models to send to their cloud for simulation as they had a large number of 'include files' scattered across multiple on-premises network drives," says TotalCAE's Rod Mach. The TotalCAE cloud solution includes a tool to reduce this cloud prep time to zero. TotalCAE automatically locates your model files on your local PC and network drives, then reformats them to remove local network references, compresses and zips them in preparation for analysis. You then upload the zip file to TotalCAE on AWS, where the model is automatically solved.

Download time is a common barrier for new clients moving to a cloud-based simulation environment. CAE application output can be in the order of hundreds of gigabytes, which can take hours to download, depending on the network bandwidth available to the user. To minimize download delays, TotalCAE offers a Windows auto-downloader client. Models will start downloading to the user's local storage environment as soon as the simulation finishes. A simulation submitted at the end of the day can solve overnight with the results ready to post-process when the engineer arrives at work in the morning — even if the download took hours. Even for clients using a dedicated high-speed network — such as AWS Direct Connect — engineers can benefit from automatic downloading. Users can also choose to do some pre/post-processing remotely using TotalCAE

remote visualization capabilities. TotalCAE offers a full Linux desktop in your browser that runs most popular CAE applications. This reduces the need for downloading every model, or bringing models back to make simple model changes and then needing to re-upload them.

In addition to optimizing simulation data download/upload, TotalCAE offers safe and secure simulation data management for all jobs submitted through the TotalCAE portal. TotalCAE will store and archive all simulation inputs and results data into a virtual filing cabinet hosted by Amazon Simple Storage Service (Amazon S3), offering virtually unlimited storage capability. Engineers can use TotalCAE's built-in search feature to find simulation files that were run months or years earlier, using job properties such as materials, job and project names, or other client-specific search criteria. No additional steps are required to store the data to make it discoverable.

END TO END SUPPORT WITH ONE NUMBER TO CALL

TotalCAE helps clients meet deadlines by having just one number to call for any issues. This eliminates having to decide which vendor or in-house IT staff member to call. TotalCAE will determine if issues are cloud issues, application version issues, or model problems — with a one-hour response time guarantee.

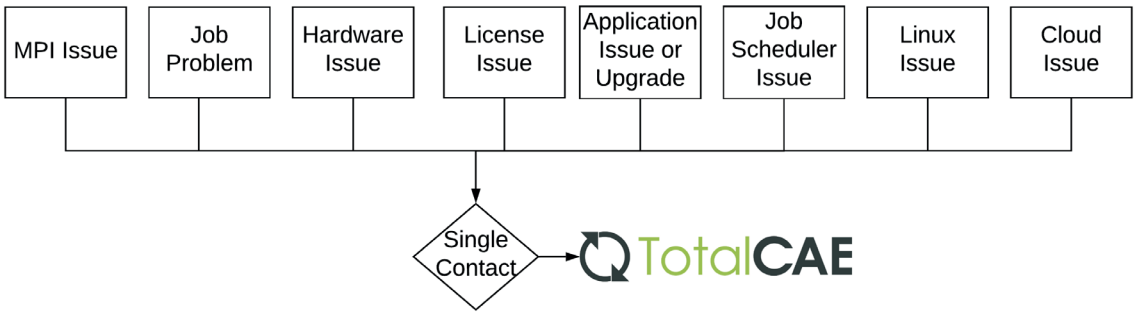

By removing this burden off the engineer to determine whether issues are caused by the cloud hardware, the software, or the model, users can stay focused on solving their engineering problems. TotalCAE's support team can help guide the user to a potential workaround. TotalCAE can also implement and deploy custom fixes provided by the ISV to the client's cloud if the issues are application related.

If TotalCAE determines there is an issue with the CAE application, TotalCAE will help open a support case with the CAE vendor to identify a workaround. Solving CAE applications with HPC is not always straightforward due to the complex nature of distributed computing.

"TotalCAE is very fast answering questions which is a big help in keeping project deadlines." -

Lead CAE Engineer - Medical Device Manufacturer.

WHEN IS TOTALCAE ON AWS NOT A GOOD FIT?

TotalCAE on AWS may not be a good fit in the following two situations where cloud computing may not provide any added benefits:

1. If you only run small simulation jobs that turn around quickly on a workstation then the extra compute power provided by the cloud is overkill.
2. Clients who have limited CAE per-core solver licenses which are already fully utilized on available workstations may also not benefit from the unlimited compute power that the cloud offers.

WHAT TYPE OF AWS INSTANCES DOES TOTALCAE USE FOR CAE?

TotalCAE recommends Intel-based instances on AWS for CAE workloads. Intel has been supporting cloud HPC for many years. Most engineering solvers take advantage of Intel Advanced Vector Extensions (Intel AVX), Intel MPI Library, and Intel Math Kernel Library (MKL) to speed up their calculations. Intel engineers collaborate with prominent software vendors developing engineering simulations, to optimize their code on Intel processors, ensuring peak performance and reliability.

TotalCAE supports robust, time-tested, Intel solutions to avoid debugging issues with applications which have not been extensively tested on other CPU alternatives. Using Intel-based instances combined with the scale of AWS infrastructure offer real performance, reliability, and stability advantages.

TAKEAWAY

TotalCAE on AWS has helped clients reduce time waiting on simulation results from weeks to hours. Clients have reduced time to market and increased their CAE agility to meet unexpected demands. It can be difficult for clients to build out a reliable and easy to use cloud HPC environment for engineers. Challenges include making HPC job submission simple to use, accessing and managing CAE licenses, managing job simulation costs, optimizing the entire workflow from submission to solution, managing the variety of CAE applications, and debugging issues when they arise.

Using TotalCAE on AWS eliminates the complexity associated with building, managing, and maintaining engineering HPC in the cloud. Take the next step and email info@totalcae.com today to learn about our free TotalCAE on AWS trial to experience these benefits first-hand.

intel

TotalCAE

This report is submitted by TotalCAE.

engineering.com