



Computer-aided
engineering (CAE)
on AWS empowers
engineers to innovate

- 3** Engineer a New Frontier in R&D
- 4** Most Common CAE Applications
- 5** Aerodynamics with CFD Simulation
- 6** Modern Engineering Demands Higher Performance
- 7** Run Your CAE Workloads in the Cloud,
and Unleash Unlimited and On-Demand HPC
- 8** Faster Innovation Means Faster to Production
- 9** Why CAE on AWS?
- 12** Case Studies
- 16** Get Started With CAE on AWS



Engineer a New Frontier in R&D

Unleash productivity with CAE

Today, computer-aided engineering (CAE) is standard in any industry that uses design software to develop new products or to create new features for existing products. CAE allows engineers to perform simulations of a product's physical properties quickly by eliminating the need to build a physical prototype. Virtually all industries are using CAE in product development processes, including: automotive, aerospace, plant engineering, electronics, energy, and consumer goods.

Simulating with CAE takes only a few hours at most, compared to the days or weeks it would require to build a physical prototype. Products that can be test simulated range from extremely small components inside electronic devices or machine parts, to very big and complex structures such as race cars, bridges, skyscrapers, or even power plants.

Engineers need simulation models to quickly and accurately ensure the effective functioning of a product in all environments and at different periods of time in a product's lifecycle. However, realistically simulating a product's complex geometries and testing those in all conditions requires an immense amount of computing power. CAE simulations require storing and processing large amounts of data, and this is only possible with "clusters" or groups of really powerful high performance and reliable computers connected through high speed communication channels.

Modern simulation tools, software as a service (SaaS) solvers, and cloud computing, have removed barriers to streamlining the difficult endeavor of CAE. Engineers can now leverage the power and speed of high performance computing (HPC) in the cloud to speed up the standard CAE workflow and significantly reduce the cost and time required for each design iteration cycle.

Most Common CAE Applications

The most computationally intensive applications in CAE for R&D include computational fluid dynamics (CFD), finite element analysis (FEA), and computational electromagnetics (CEM).

CFD:

Any time that fluids are tested in the environment of a new product, the complexity multiplies because of the enormity of the challenge to quantitatively describe all physical phenomenon involved in a real fluid and determine a prediction model for how it will interact with a product. Modern CAE applications allow engineers to simulate fluid motion and use numerical approaches to examine its physical properties such as velocity, pressure, temperature, density, and viscosity. In order to virtually generate a solution for a physical phenomenon associated with fluid flow, those properties have to be considered simultaneously in a simulation using multiple integrated numerical solvers and technologies to analyze and solve a wide range of problems related to laminar and turbulent flows, incompressible and compressible fluids, multiphase flows, and more.

FEA:

Engineers use FEA to simulate any given physical phenomenon, such as structural, heat transfer, wave propagation, etc. using mathematical techniques to reduce the number of physical prototypes required in the product design phase and develop better products, faster. Computers are used to calculate highly complex mathematical expressions that need to be solved to compute quantities of a structure such as stresses and strains, and in order to predict the behavior of a structure under a given load.

CEM:

This is a young and fast growing discipline used for the design of electrical devices like television, radio, internet, mobile telephones, satellite communication systems, generators, lasers, medical imaging systems, and much more. CEM previously required the building and testing of prototypes; however, today CEM can be completed faster and cheaper with numerical computation.

With the availability of powerful computing resources it is more attractive to use analytical methods to simulate the interactions between electric charges at rest and in motion. Newer application areas for CEM opened up by high powered computing workloads include designing microwave networks, antenna placement, the optimization of radar cross sections for stealth devices, and studying electromagnetic signal interference on 5G devices, which require simulations at 28GHZ.

Industries that use common CAE applications will find many CFD/FEA scale out solvers available that use mathematical algorithms to simulate responses to FEA structural loads, thermal loads, buoyancy, wireless electromagnetic signals, and much more.^{iv}

Aerodynamics with CFD Simulation



Formula 1® Motorsports runs CFD simulation applications on AWS

Aerodynamics is the science of moving air and the interaction between the air and solid bodies moving through it. Aerodynamic simulations play an important role in the design of almost all the transportation vehicles and big structures in the world today and such aerodynamic flows can be accurately simulated with Computational Fluid Dynamics (CFD).

For Formula 1 cars, aerodynamic downforce is a very significant performance factor, improving tire grip and helping a car travel faster through corners. As a result of the downforce produced, the current generation of F1 cars suffer when they are running in the wake of another, and the ability to sustain close racing and overtaking is reduced. Currently, an F1 car running 10 meters behind another typically loses up to 50 percent of its downforce.

With increased excitement for racing fans at the core of the project, F1 used AWS to look closely at how cars interact when racing together and the results have driven the recently announced 2021 F1 Technical Regulations. With the insights gained from these simulations, F1 has been able to design a car with only 10 to 15 percent downforce loss when running 10 meters behind another, which means better racing and more overtaking. The F1 team correlates the CFD on AWS simulation data with 50 percent scale model tests in the Sauber Motorsport wind tunnel.



“Through our Amazon Web Services (AWS) connections, we’re using their EC2 resource to greatly increase our aerodynamic testing capabilities. In CFD, we split the environment up into millions of tiny cells and we solve equations for each one of them. In a typical team environment, you might use 192 cores to solve 95 million cell models. In our current configuration, we’re using 1152 cores and we’re solving up to 550 million cell models, enabling us to develop state-of-the-art unsteady multicar CFD models to Formula 1 timescales.”

Pat Symonds, Chief Technical Officer, Formula 1

Modern Engineering Demands Higher Performance

Limitless HPC resources in the cloud empowers engineers

CAE enables engineers with a way to solve the most complex problems quickly, but only when they are free to run simulations at optimal scale. However, compute-intensive simulations can get constrained by limited on-premises IT infrastructure capacity, high capital expenditures on hardware, and the persistent need for traditional licensing models and technology refreshes. HPC on AWS unlocks essentially unlimited amounts of computing power and data storage when engineers need them.

- IDC reports that for every \$1 spent on HPC, **businesses see \$463 in incremental revenues** and **\$44 in incremental profit.**ⁱ
- **72.8 percent of organizations** using on-premises HPC reported delayed or cancelled HPC jobs due to lack of resources.ⁱⁱ

Hidden costs of on-premises HPCⁱⁱⁱ

- **Lost productivity** – as utilization increases, engineers have to wait for their CAE applications to start, which leads to downtime and lost productivity.
- **Missed innovation** – questions are left unasked, simulations left undone, and potential revenue is left on the table.
- **Technology refresh cost** – customers are locked into a technology decision for five years or more when storage must be refreshed with a complete data migration. This in a time when new silicon and software is being released at an accelerating pace. Cloud-based HPC solutions are continuously refreshed.
- **Increasing technical debt** – organizations with high utilization of on-premises HPC infrastructure often find fixed configuration, on-premises HPC clusters logical and attractive. But evolving engineering practices and research concepts are often overlooked. Retooling and adapting these engineering applications to meet the requirements of an existing infrastructure (instead of the other way around) results in delays and below-par performance. With cloud-based HPC infrastructure, engineers have easy access to the latest technologies, so they can employ the latest applications to innovate faster.
- **Longer time to results** – on-premises HPC solutions offer a limited number of cores available at any one time. When quick results from CAE applications are highly desirable, cloud-based HPC is the best choice and can help organizations make innovations available more quickly to gain first-to-market advantage.

Faster Innovation Means Faster to Production

A better alternative—leverage the latest technologies on AWS

By using the AWS Cloud, engineers across all disciplines can access nearly infinite HPC compute infrastructure capacity and resources to run complex simulations, for a fraction of the time and cost it would take to run locally. AWS delivers an integrated suite of services that provides everything needed to quickly and easily build and manage HPC clusters in the cloud to run the most compute-intensive CAE workloads to support your organization or project. AWS removes the long wait times and lost productivity often associated with on-premises HPC clusters. Flexible configuration and virtually unlimited scalability allow you to grow and shrink your infrastructure as your workloads dictate, not the other way around. All the while, you only pay for the resources you consume.

Your organization may be at the start of the journey with some initial CAE projects in the pipeline, or you may be an engineer working independently without a great deal of administrative support. Either way, AWS is here to help you integrate the value of CAE more deeply into your engineering process. If you aren't already exploring CAE now, you may need to play catch-up. The speed of innovation required to engineer new products and solutions is continuously challenging the processes that administrators need to develop internally to stay competitive. There is tremendous value for your organization and for your career as an engineer, developing skills in cloud-based CAE.

How It Works


Monitor performance and change your cluster as needed when your workloads evolve



Determine the compute, storage, and networking requirements for your HPC codes



Create and configure your HPC cluster using AWS ParallelCluster



Submit your HPC job using a job scheduler, or directly from the CLI



Visualize your results using NICE DCV or pipe your results into other AWS services such as machine learning (ML), deep learning (DL), or Athena for post processing

Why CAE on AWS?

Value Case #1 – Company using on-premises HPC



Stop worrying about infrastructure, and instead transform your R&D

Today's engineering professionals want to move faster and run more CAE simulations, more often. It is so important to provide your engineers with the infrastructure they require—when they require it. After all, the more quickly successful R&D is completed, the sooner your organization can begin marketing the latest products. Bring your legacy applications to AWS, and we can help you to easily set up the resources your engineers need—without breaking your budget.

"The ever-increasing interdisciplinary challenges of vehicle development require innovative, disruptive computational technologies to shorten development cycles to continue to satisfy customer needs."

Dr. Henry Bensler, Head of CAE Methods at Volkswagen Group Research

- **Breadth and depth.** More CAE engineering problem solvers are running their simulations on AWS than anywhere else due to the breadth of instance types. Companies such as Formula 1 Motorsports and Autodesk trust AWS to run their most critical HPC workloads.
- **Scale up fast.** Once you have experienced the scale of AWS for computationally intensive applications, see how you can also leverage the ability to modernize the product development lifecycle with IoT and ML using the data from your simulations.
- **It's easy to manage your budget.** AWS offers a range of pricing options so you can easily manage your budget efficiently with little or no upfront investment. You pay only for the resources you use, and it's easy to set up automated monitoring and notifications systems.
- **Strong partner ecosystem.** AWS partners can help accelerate your journey using secure, cloud-based HPC by offering cost-effective scalable solutions, engineering solvers, and CAE simulation tools on the AWS Cloud. Partners like Intel®, Rescale, and Altair provide you the engineering agility to quickly respond to changing business needs.

Find out how your engineering organization can profit by running your CAE applications on AWS and harnessing the power of Intel® Xeon® processors.

Now it's your turn

[Learn more](#)

Why CAE on AWS?

Value Case #2 – Small/Mid-size company or start-up



Want to change the world? Then step up your game

If you run a small business, the cost of acquiring on-premises HPC can be prohibitive, so the business case for cloud-based solutions is much more straightforward. Independent engineers who don't have a great deal of administrative support could gain tremendous benefit from access to more compute resources on the AWS Cloud, enabling them to run their CAE workloads faster and leverage the latest technologies to uncover deeper insights. AWS unlocks your creativity by making it easy to open an account and get the capacity you want—just for the time you want, and then scale up or tear it down as you wish. With AWS, you are in control of how much you want to spend, when, and for how long. Just imagine what you could get done with unlimited amounts of computing power and data storage at your fingertips.

- **Want a better machine? The AWS Cloud delivers.** Open an account on AWS and start doing your work on a new level. Get quicker turnaround, complete troubleshooting and optimizing faster, and reduce your overall time to market.
- **Whatever your budget, AWS can accommodate you.** On the AWS Cloud, you are in control of the amount of capacity you need to run. Stop anytime. Launch new CAE workloads right away and easily scale resources to meet your needs, which will lower your R&D costs.
- **AWS is ready when you are.** Get to work right away with novel cloud-native solvers and SaaS partners to help you model, develop, and test your products faster—enabling you to be more competitive.

“The more power you can give an engineer, the more complex problems they'll solve.”

Gerry Harvey, Vice President of Engineering, OnScale

What are you waiting for? The world is waiting for your new breakthrough in product design and there are industries to disrupt. Unleash your creativity and get more done.

[Open an account](#)

[Learn more](#)



Gaining competitive advantage for a business that is driven by innovation

Western Digital (WD) designs and manufactures hard disk drives (HDD) for the extremely competitive enterprise market with its ever-increasing demands for greater capacity and faster read/write times. WD engineers need to simulate millions of combinations of materials, energy levels, and rotational speeds enabling them to cram more 1's and 0's, or bits, into the same space and make sure the drives can handle the increasing demands for performance. The challenge was how to get the results of these electro-magnetic simulations more quickly allowing their engineers to make quicker decisions and get new products to market faster. WD found the best solution was to run their CAE simulations on a massive HPC cluster built on AWS using Amazon EC2 Spot Instances.

WD built a cloud-scale HPC cluster on AWS enabling the company to run 2.3 million simulation jobs on a single HPC cluster of 1 million vCPUs—built using Amazon EC2 Spot Instances, **reducing the time to results from 20 days to eight hours.**

“This successful collaboration with AWS shows the extreme scale, power, and agility of cloud-based HPC to help us run complex simulations for future storage architecture analysis and materials science explorations. Using AWS to easily shrink simulation time from 20 days to eight hours allows Western Digital R&D teams to explore new designs and innovations at a pace un-imaginable just a short time ago.”

Steve Phillpott, Chief Information Officer (CIO), Western Digital Corporation

AWS solution:

- Run Amazon EC2 instances, and pay only for what you use.
- Take advantage of spot pricing to further reduce cost for time-flexible workloads.
- The simulation tasks read and write data from Amazon Simple Storage Service (S3), taking advantage of S3's ability to store vast amounts of data and to handle any conceivable request rate.

Results delivered:

- The cluster grew to one million vCPUs in one hour and 32 minutes and ran full-bore for six hours at a cost that was approximately half the cost of making the simulation run on an in-house cluster, if they had one of that size.
- Achieved extreme scalability in terms of capacity and configurations—without a large upfront investment or lock-in to legacy IT infrastructure.

[READ MORE](#)



Minimize spending without compromising research

TLG Aerospace is an engineering services company that provides CFD analysis and conducts aerodynamic simulations on aircraft to predict the pressure and temperature surrounding airframes. The company wanted to reduce the costs associated with running simulations and gain scalability to take on larger simulations. The challenge was how to cut costs and accept projects that exceeded its on-premises capacity. The solution TLG found was to migrate CFD simulation application from its existing provider to AWS.

TLG takes advantage of Amazon EC2 Spot instances, a way to use unused Amazon Elastic Compute Cloud (Amazon EC2) computing capacity at a discounted price. **TLG also uses Amazon Simple Storage Service (Amazon S3) buckets to store multiple terabytes of simulation data** on the cloud and Amazon Elastic Block Store (Amazon EBS), which offers persistent block-level storage volumes that can be used with Amazon EC2 instances. **TLG uses the Amazon CloudWatch monitoring service to track the progress of CFD simulations in an online AWS console.**

“Running a typical simulation was costing us hundreds of dollars per case, and there may be hundreds of cases per project. We saw a 75 percent reduction in the cost per CFD simulation as soon as we started using Amazon EC2 Spot instances. We are able to pass those savings along to our customers—and be more competitive.”

Andrew McComas, Engineering Manager, TLG Aerospace

AWS solution:

- Choose from a range of AWS services and Intel powered Amazon EC2 instances, and pay only for what you use.
- Take advantage of spot pricing to further reduce cost for time-flexible workloads.
- Avoid the capacity limitations of many other cloud providers.

Results delivered:

- Reduced cost per simulation by 75 percent, on average.
- Gained ability to check simulation status from anywhere to catch problems early and reduce costly rework.
- Removed limit of 1,000 nodes imposed by previous service provider and is now able to bid on projects that have a larger scope.

[READ MORE](#)



Accelerating innovation in modern vehicle design and development

Nissan Motor Corporation is a global, full-line vehicle manufacturer that sells models under the Nissan, INFINITI, and Datsun brands, and in 2018 sold 5.52 million vehicles worldwide. Nissan engineers require flexibility to run hundreds of simulations with on-demand agility, without exhausting capacity, but found they were limited by fundamental aspects of their on-premises computing infrastructure, such as: limited electric power, high total cost, and data center utilization challenges (Nissan averaged 40 to 80 percent utilization). The company recognized that these on-premises problems threatened its innovation, market leadership, agility, and time-to-market. Therefore, **Nissan decided to shift its technical computing to AWS via the Rescale ScaleX platform to gain an agile, cloud-enabled, platform-based solution** on a pay-per-use model to minimize overall cost-per-simulation, gain instantaneous agility, and continuously adopt the latest technologies.

“The challenges we faced at Nissan revolved around managing our in-house HPC system alongside keeping up with the latest technology innovations needed to meet market demands.”

Seiji Kawachiya, General Manager of Engineering and Quality System Department, Nissan Motor Corporation

AWS-Rescale Solution:

- Scale a diverse range of the latest hardware and applications instantly.
- Access controls for administering simulation users, projects, software, hardware, and budgets.
- Allocate the right resources at the right time.
- Avoid over-provisioning and wasting resources.
- Avoid under-provisioning and creating product development delays.

Results delivered:

- 18 percent cost optimization of applications and productivity.
- 50 percent cost reduction for HPC expenses.
- Reduced queue time from two to three days to zero days.

[READ MORE](#)

How CAE on AWS Benefits Engineering



Go limitless. Answer more questions with your simulations than ever before, without the worry of computation resource limits. Get the power to exceed your on-premises capacity.



Get answers more quickly to accelerate development. Launch simulation workloads, and run CAE workflows faster. Scale capacity in seconds, and automate the design process with workflow orchestration.



Control the process and don't be controlled by it. Create, operate, and tear down secure, well-optimized HPC clusters in minutes.



Securely run your compute-intensive simulations. Maintain safety of your intellectual property using AWS identity, authentication, and access control services.



Shorten time to answers. Leverage the latest technology, and gain a competitive advantage.



Unleash productivity. Free engineers to launch compute-intensive simulations to speed new product development without worrying about compute time by leveraging HPC in the cloud to circumvent internal capacity constraints.



Unleash creativity. Start compute-intensive simulation workloads as soon as engineers are ready, avoiding the queue. Start new CAE applications to match the specific requirements of each simulation, and gain immediate access to the latest technology upgrades without stalling the process.



Avoid schedule overruns. Add more capacity as deadlines approach.



Take on new business. Gain confidence in your team's ability to deliver.



Getting Started with CAE on AWS

Getting started with CAE on AWS is easy. You can set up an account with just a few clicks and without any up-front commitment. When you choose AWS, you gain instant access to HPC resources including Intel Xeon processor technologies and a full spectrum of storage, computing, and analytics. Start experimenting with AWS today with a sample simulation or tutorial, gain deeper insight through whitepapers and videos, or find a partner and get hands-on guidance.

Learn more

<https://aws.amazon.com/hpc>

- i Hyperion Research, 2018.
- ii Economic Models for Financial ROI and Innovation from HPC Investments, Hyperion, 2018.
- iii What a TCO Analysis won't tell you, White Paper, AWS, 2018.
- iv Some popular CFD solvers include Ansys Fluent, CFD++, LS-DYNA, and OpenFOAM.