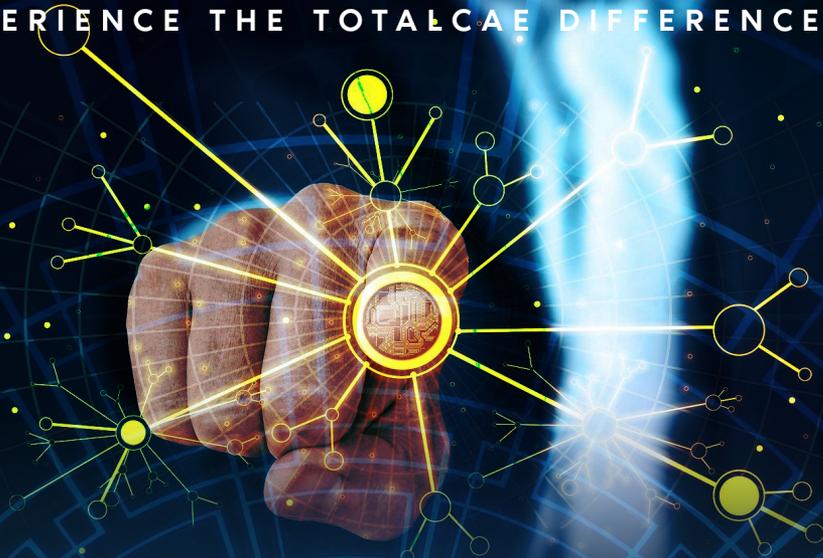


EXPERIENCE THE TOTALCAE DIFFERENCE!



MANAGED HPC CLUSTERS AND CLOUD FOR ENGINEERS - HPC EVERYWHERE

TotalCAE
Blueprints for
Abaqus

This white paper will discuss the options you have to help you speed up your Abaqus models and get more throughput from your existing Abaqus tokens with High Performance Computing (HPC). TotalCAE offers HPC for engineers including a managed on-prem cluster appliance for Abaqus and hundred of other CAE applications, and on-demand public cloud solutions.

Why HPC?

HPC is used today to take runtimes from weeks to days, and days to hours. HPC enables engineers to achieve higher engineering throughput to get more engineering done per day. A typical result we achieve is below, going from once a day engineering results to many times a day.

Workstation Runtime	HPC Runtime	Runtime Reduction
7 hours	50 minutes	> 7X

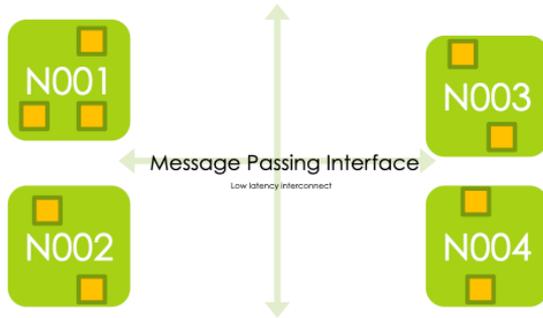
Bigger Workstation vs. HPC ?



Most engineers have noticed just buying a newer PC with updated but similar number of CPUs does not provide a game changing boost in performance. This is because per-CPU speeds have remained relatively flat over the last few years.

Speedup for CAE Applications is now achieved by spreading the bits of the calculation out across many CPUs across many

computers. These “compute nodes” exchange data through through software middleware called Message Passing Interface (MPI) that utilize a very low latency InfiniBand network to move messages between these computers in under 1 microsecond to enable the application to speedup as more CPUs are added.



Why isn't HPC Ubiquitous?

Why it is clear HPC provides great benefits to speeding up Abaqus, HPC is often not been ubiquitous due to the complexity of procuring, managing, and maintaining HPC systems has made them non-attainable for many mid market clients. In other cases the proper hardware was purchased, but the systems were difficult to use, unreliable or otherwise not able to keep up to date with the business making them unusable for the engineering team.

Most corporate IT departments are not staffed up to support Linux HPC clusters, Linux, HPC batch schedulers, and supporting all of the CAE applications on top of that due to the niche aspect of HPC for engineering. Hiring an in-house IT team to support HPC is often not cost effective, which leads clients to believe only cloud

solution options are a viable option (which TotalCAE offers, but cloud is not the lowest cost option).

Solution: TotalCAE Managed HPC



On-prem HPC Cluster
& Cloud Management

+



Application
Management

+



Efficient Interface
Three Clicks to Solve

TotalCAE solves this HPC challenge by outsourcing the HPC Cluster management and upkeep to TotalCAE. TotalCAE managed cluster appliance for Abaqus is an HPC cluster that is pre-configured for Abaqus that TotalCAE completely manages for you in your company's data center. TotalCAE includes everything needed pre-configured out of the box including an easy to use job submission web and command line portal, monitoring and analytics, data management, job scheduler, and hundreds of applications pre-integrated.

TotalCAE also offers the identical solution in a pay as you go offering on the cloud (utilizing AWS or Azure hardware) in your cloud subscription, or hosted by TotalCAE.

Running on TotalCAE

The screenshot shows the TotalCAE Portal interface. At the top, there is a navigation bar with 'TotalCAE Portal' and tabs for 'Jobs', 'Analytics', 'Data', 'Billing', and 'Terminal'. On the left, a sidebar menu is open to the 'Abaqus' section, which includes links for 'List Licenses', 'List Jobs', and 'Status'. The main content area is titled 'Abaqus' and contains a form for job configuration. The form includes the following fields and options:

- Select the Solver Version:** 2021 (dropdown)
- Select the Input Deck:** /work/mach/flattened/e6.inp (text input with 'Browse...' button)
- Working Directory:** /work/mach/flattened/ (text input with 'Browse...' button)
- Select Additional Files:** No Additional Files Selected (text input with 'Browse...' button)
- Job Name:** e6 (text input)
- Extra Arguments:** double (text input)
- Number of Cores:** 12 (spinner)
- Quality of Service:** normal (dropdown)
- Run On:** c5n.18xlarge (dropdown)
- Exclusive:**
- Run From Local Scratch:**
- Hold Until:** (calendar icon, refresh icon, and 'Clear' button)
- Time Limit (hours):** 2 (spinner with 'Clear' button)
- Project:** default_project (\$3,821.55 available) (dropdown)
- Comment:** (text input)

At the bottom of the form, there are three buttons: 'Submit' (orange), 'Clear All' (grey), and a status line: 'Approximate cost per hour: \$4.02 ? Maximum cost: \$8.04'.

Running Abaqus on TotalCAE takes just a few clicks. You point your browser to a persistent URL set up for your company and you are presented with a simple to use interface shown above to submit all your engineering applications.

The steps to submit are:

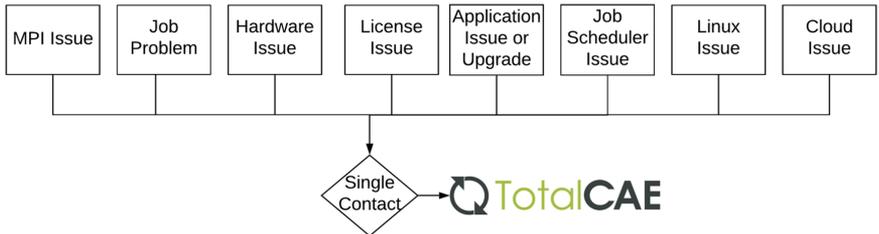
1. Upload your model.
2. Select the number of CPUs.
3. Push Submit.

Just three clicks to solve. There is no complicated setup or wizards to click through so you can focus on engineering and less on job setup. To see a video of using our cloud platform visit

<https://totalcae.com/abaqus>. For power users that prefer the command line, the full power of the Linux command line is available through the portal Linux terminal. This is useful for submitting multiple load cases that can benefit from the automation of the command line version of our portal.

TotalCAE 1 Hour Response for Any Issue - One Number to Call

If you have any issues, the entire solution is backed by TotalCAE 1-hour response, where a member of our team will help you debug any issues you are having with submitting or running your job. You don't have to determine if this is a cluster issue, MPI issue, hardware problem, scheduler, Linux, or cloud issue, TotalCAE will take this burden off the engineer so they can keep their engineering deadlines



TotalCAE HPC On-Demand Cloud

In addition to TotalCAE managed HPC Clusters that we remotely manage in your data center, TotalCAE also offers on-demand public cloud computing. Cloud computing has two main benefits.

1. Providing on-demand burst capacity – Public cloud allows users to tap into additional computing capability outside the user’s existing HPC cluster or workstation to meet peaks in demand or short term burst projects. These types of projects don’t require adding persistent capacity, or may be time sensitive and can’t wait for a procurement cycle.
2. Run optimization or large jobs that you don’t run on a normal basis – If you want to run a very large optimization, or even a special large job, public cloud can be a way to enable this type of workload that may otherwise may not be possible.

TotalCAE on demand public cloud is available for both our existing managed on-premises customers, and for standalone on-demand cloud only usage. TotalCAE will host and manage the solution in your Azure or AWS subscription so you pay the lowest possible cloud costs.

Controlling Cloud Job Costs

When using on-demand clouds, a challenge clients face is how to manage and control cloud costs to stay in budget. TotalCAE portal on cloud has many features that enable engineers to help manage their cloud spending and make sure they are optimizing their cloud budget for their simulations.

TotalCAE will give you an estimate of the job cost prior to pushing the submit button so the engineer can gauge how much a particular simulation job run will cost, and which project to bill the job usage to. The cost of a particular job will depends on how many CPUS you need to rent, dialing up and down the CPUS will show you the job cost impact.

The TotalCAE scheduler will automatically route your job to the appropriate node, power them on if there is enough Abaqus tokens to run, and finally run your job. The nodes will be powered off when your job completes to save money as the compute resources are “pay as you go”.

TotalCAE Billing tab includes a project based billing system that enables clients to put users in one or more projects, set budget limits by user or project, set budget warning limits, and easily generate one-click Excel job-based cost accounting reports of all usage for chargeback or record keeping.

The screenshot shows the TotalCAE Portal interface. The top navigation bar includes 'TotalCAE Portal', 'Jobs', 'Analytics', 'Data', 'Billing', and 'Terminal'. The user 'rmach' is logged in. The 'Billing' tab is active, showing a 'Projects' view. A table lists several projects with their details.

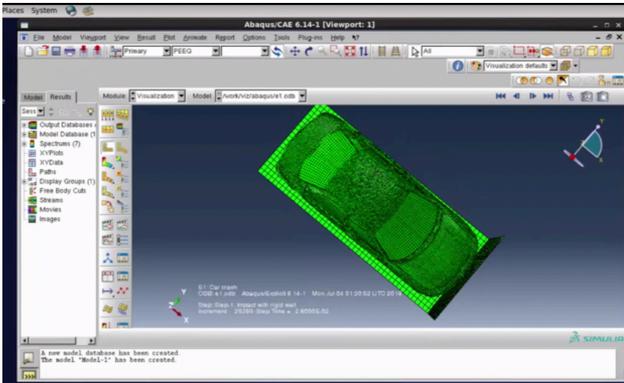
Name	Description	Status	My Available Funds	My Total Spent	
benchmarks	benchmarks	Active	\$4,902.98	\$97.02	⚙️
default_project	Default Project	Active	\$3,699.45	\$2,189.32	⚙️
demo_project	Demos	Active	\$1,000.00	\$0.00	⚙️
RITM1234567		Active	\$18,609.45	\$0.00	⚙️
sales	sales	Active	\$18,609.45	\$0.00	⚙️
test	test	Active	\$18,609.45	\$0.00	⚙️

At the bottom of the table, there is a pagination control: 'Page 1 of 1' and '10 items per page'. The total number of items is '1 - 6 of 6 items'.

Post Processing Results

Post processing differs slightly between on-prem and cloud. On most on-prem HPC environments the user will mount the cluster drive on their Windows workstation, for example as a “Z:\” drive, and directly post process the job after the run, or view the results as the job is running.

On the cloud, the data center is often far away from the engineer, so there is a challenge in public cloud downloading large ODB results from the cloud back to on-prem. One solution is to avoid the downloading by post processing on the cloud instead of the local workstation.



TotalCAE TotalViz is the TotalCAE remote visualization solution that lets you run Abaqus/CAE on a "cloud workstation" and avoid downloading every single result.

TotalCAE enables 1-click to turn on your cloud desktop to perform your post processing work, and can automatically stop the desktop to ensure the billing stops. The costs of running a cloud desktop are around 1 dollar per hour, and you can choose a project to bill the cloud desktop against.

TotalCAE Automatic Download [Public Cloud]

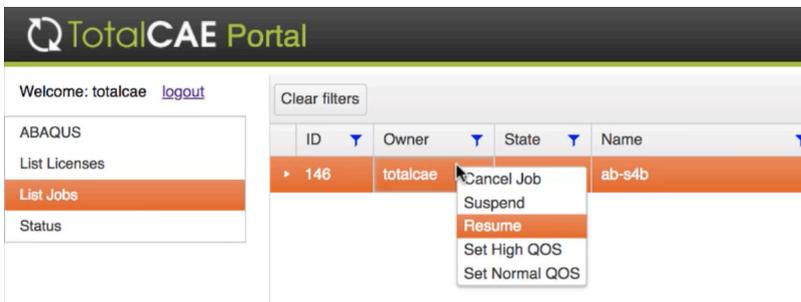
Even with remote visualization, engineers will want to bring back some result data from the internet back to on-premises. TotalCAE has an auto-download tool that will start streaming the job results back on-premises automatically when the job completes with NO user intervention.

For example, if your simulation finishes on the cloud at midnight while you are asleep, the results automatically start streaming back on-prem so that when you arrive to work in the morning, the results are downloaded and waiting for you to post process.

Monitoring and Managing Abaqus Tokens

Jobs can be monitored and managed through the TotalCAE web portal, the TotalCAE remote cloud workstation, or the command line.

TotalCAE makes it easy to suspend and resume Abaqus tokens on long running Abaqus jobs so you can run quick data checks or or run urgent short running jobs without having to cancel the existing job or do a tedious full checkpoint/restart.



With TotalCAE, just right click on a job in the TotalCAE portal to suspend/resume the tokens. It is that easy!

Users can see if jobs are running, successful, failed, or pending due to lack of licenses. TotalCAE automatically calculates the requires number of Abaqus licensing and will make sure the jobs will not run until sufficient Abaqus tokens are available. Users can also choose to re-prioritize jobs to have “hot” jobs utilize limited Abaqus licensing before low priority jobs.

Once a job is completed, users are emailed that the job is completed with relevant information from the STA file.

Detecting Abaqus Problems with TotalCAE CAEWatchdog

With TotalCAE CAEWatchdog you can be immediately notified something is looking amiss with the running Abaqus job so the engineer can decide to cancel the job, or take some remediation to save time, and real money when running on the cloud. It is also handy for on-prem to alert you to job issues without constantly monitoring the job to see if it still running properly.

For example, if the STA file is not being updated but Abaqus is still running, or if Abaqus appears running but is queued for licensing, and various other Abaqus specific issues can be detected and alert you via the portal and email to avoid you having to constantly check in on the run to see if it is iterating correctly and generating useful results.

Clear filters						
ID	Owner	State	Name	Start	End	
▶ 6	totalcae	RUNNING	myjob3	11/9/2018 21:21		
▶ 5	totalcae	RUNNING	myjob2	11/9/2018 21:20		
▶ 3	totalcae	❗ RUNNING	myjob	11/9/2018 20:50		

Warning: detected errors in the model output, this job may not be behaving correctly

Abaqus Application Licensing

The primary method of accessing Abaqus application licensing is bring your own license, TotalCAE will install a utility on your license server that will enable cloud access. TotalCAE will take care of making sure the TotalCAE HPC Cluster, TotalCAE cloud, and your end user workstations can all utilize your Abaqus licensing.

TotalCAE does not sell or rent the Abaqus licensing, we just manage it for you to take the license management, monitoring, and reporting burden off of IT.

TotalCAE HPC Everywhere

Most TotalCAE clients are running a mix of on-prem HPC for the day to day engineering work, and Public cloud for extra on-demand capacity. This is primarily due to the fact that renting or buying HPC compute time on cloud is more expensive for constant day to day usage when compared to on-prem. An example of the relative pricing between on-prem and public cloud pricing is shown below for HPC hardware.

Hardware Option	Cost	Vendor	Instance Info
56-core 3rd Gen Intel Xeon Scalable Processor	\$420 per month	TotalCAE On-Pem	3-Year with InfiniBand
64-core 3rd Gen Intel Xeon Scalable Processor	\$1,843 per month	AWS	3-Year RHEL Reserved c6i.32xlarge Instance with EFA
96-core AMD EPYC Milan 7R13	\$1,230 per month	AWS	3-Year RHEL Reserved Hpc6a.48xlarge Instance with EFA
64-core 3rd Gen Intel Xeon Scalable Processor	\$5.57 per Node Hour	AWS	Hourly RHEL On Demand c6i.32xlarge Instance with EFA
96-core AMD EPYC Milan 7R13	\$2.88 per Node Hour	AWS	Hourly RHEL On Demand Hpc6a.48xlarge with EFA
120 core AMD EPYC Milan 7003	\$1,445 per month	Azure	3-year RHEL Reserved HB120rs V3 with InfiniBand
120 core AMD EPYC Milan 7003	\$3.96 per Node Hour	Azure	Hourly On Demand RHEL HB120rs V3

TotalCAE provides both on-prem TotalCAE HPC cluster appliances AND public cloud managed by our staff so our clients have the best of both worlds. Run with low-cost on-prem HPC simulation as a service managed by TotalCAE, coupled with on-demand TotalCAE public cloud capacity for unexpected capacity or projects to maximize your CAE budget and increase your flexibility and agility to respond to new projects.

HPC Compute Hardware for Abaqus



Abaqus implicit and explicit workloads have different performance characteristics that impact the type of hardware to use to reduce the runtime of your models. In addition, the amount of time your model will take on HPC is highly dependent on the model itself. Abaqus/Standard linear and nonlinear implicit solver time is dependent on degrees of freedom (DOF) and iterations. If you increase the DOF, elements and iteration count for a simulation it will increase the runtime.

Similarly for Abaqus/Explicit the solve time is dependent on the number of elements and time duration of events. If you increase the number of elements or the time duration of the simulation, it will take longer to solve. A basic summary with some general hardware recommendations are below

Implicit - Linear Static Analysis

These types of workloads benefit from:

1. Higher clock rate processors with lower core counts.
2. Larger memory with 16GB of memory per core.

3. Extraction of natural frequencies and mode shapes utilize dedicated SSD scratch disks

Implicit - Non-Linear Static Analysis

For runs greater than 2M DOF we recommend the same hardware as the Linear Static Analysis. Smaller jobs from 500K - 2M DOF can run with 2-4GB of memory per core.

Explicit - Non-Linear Static Analysis

For small to medium-sized jobs of 100K to 400K elements, 2GB Memory/Core is common and these models can run on 1 node.

For large jobs of 400K elements and above, 3.5-4GB of memory per core is recommended.

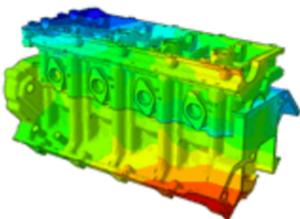
Note that memory not used by the Abaqus solver is not wasted, having additional memory available allows the file I/O to be cached by Linux to allow better disk I/O performance.

For complex models, InfiniBand or similar low latency fabric is required for multi-node Abaqus/Explicit performance scaling due to messaging passing between nodes.

The table below summarizes the compute hardware included in the TotalCAE recommended on-prem HPC Cluster appliance for Abaqus:

Type	Memory	CPU	RDMA	Other
Explicit	256GB DDR4 3200 MHz Memory	Two Intel® Xeon® Gold 6348 - 28 cores at 2.6 GHz Base/3.5 GHz Turbo - 56 cores per node	HDR100	
Standard	512 GB DDR4 3200 MHz Memory	Two Intel® Xeon® Gold 6346 - 16 cores at 3.1.GHz Base/3.6 GHz Turbo - 32 cores per node	HDR100	Dedicated SSD Scratch Drive

Example Benefits of 3rd Generation Intel® Xeon® Scalable Processors with s4b



The s4b benchmark is a nonlinear static analysis that simulates bolting a cylinder

head onto an engine block. s4b is a 5 million Degree of Freedom (DOF) model.

This was run on 32 cores, 16 cores per socket on both systems.

Processor	Elapsed Time
Intel® 2nd Generation Intel Xeon® Scalable processors	309 Seconds
Intel® 3rd Generation Intel Xeon® Scalable processors	203 Seconds

Note that the new Intel® 3rd Generation Intel Xeon® Scalable Processors have a 34% improvement over the previous generation processors on s4b.

Visit <https://caebench.org> for the latest Abaqus benchmarks on Intel 3rd Generation Intel Xeon Scalable Processors.

What about GPUs?

Abaqus can transparently take advantage of GPU's, and the use of the GPU can reduce the amount of tokens required to achieve speedup, though as you add more CPUs the relative benefit of the GPU diminishes. Since the benefit of a GPU on real world client models varies widely, TotalCAE recommends a benchmark analysis to determine if a GPU will reduce the amount of Abaqus tokens for your model to achieve the desired turn around time.

To register for this analysis visit:

<https://www.totalcae.com/free-performance-benchmark-program>

TotalCAE White Glove Support

TotalCAE HPC Clusters and public cloud comes with our 1-hour response white glove support via phone (where you talk to a human in your language), email, chat, text, or web with a full ticket system tracking. TotalCAE proactively manages all your on-prem and cloud HPC environments, all of your CAE application and license servers so you can focus on engineering, and not IT.

Every ticket gives you the opportunity to rate the TotalCAE Team member, if you do not give them a good or great rating then you will receive a call from the TotalCAE CEO to ensure we are exceeding your expectations on every interaction.

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

- BD CAE Engineer

Users also have access to our knowledge base with common questions for Abaqus with our self-service help.

Why You Need the TotalCAE

The TotalCAE HPC solution does all the following for your business:

- One vendor to manage all your HPC environment, CAE applications, and license servers with 1 hour response. No more finger pointing or figuring out who to call with an issue.

- Choose On-prem, cloud, or both with a single vendor.
- Simplified TotalCAE Portal interface to eliminate the complexity of HPC simulation, three clicks to solve hundreds of CAE supported applications.
- High speed download and remote visualization tools to overcome cloud download challenges
- Full command line access for power users.
- CAEWatchdog to monitor and alert to Abaqus job issues.
- Minimize the usage of rented public cloud resources by intelligent and automatic smart power on/off of resources based on job demand.
- Simple pricing model. Pay as you go flat-fee subscription.
- TotalCAE Public Cloud looks and feels identical to an on-premises cluster, no workflow changes. Utilize your own cloud subscription, or TotalCAE can host it.
- Use popular Abaqus integrations such as user subroutines with Intel compilers, PolyUMod, and other 3rd party software that you need.
- Access to TotalCAE IT professional experts that speak your language via phone, email, and web with support for nearly every FEA/CFD solver and engineering application on the market.

- Complete billing and cost control system to assist in controlling, monitoring, and maximizing your CAE budget on cloud.

Your Next Step

TotalCAE offers free consultations to help you find the best HPC strategy for your business needs and budget, including helpful calculators to determine the best mix of on-prem and cloud.

There is no risk and no obligation.

Discover how TotalCAE can enhance your engineering department capabilities. Don't wait: claim your FREE consultation today by emailing info@totalcae.com