A TotalCAE White Paper

IMPROVING HPC ROI

Benchmarking the value of an on-premise HPC
cluster appliance for engineering applications.

TotalCAE DELL Technologies intel

Engineering teams are under pressure to create better designs faster, and to perform more frequent simulations throughout the product development process. At the same time, model size and complexity are rapidly expanding, which is taxing the capabilities of even the most powerful workstations.

That's because per-core clock rates have changed very little over the past five years. To achieve any measurable speedup in solve times requires the utilization of GPU compute power, scaling out jobs to several CPUs, or increasing throughput by using a scheduler to improve efficiency.

Because upgrading workstations to address the need for additional compute power is often cost prohibitive, engineers frequently turn to workarounds (such as limiting model size or detail), which can reduce the fidelity of results. Engineers need on-demand access to greater computing power.

An on-premise high-performance computing (HPC) cluster can provide that engineering horsepower by giving engineers the ability to offload complex simulations to the cluster while remaining productive on their personal workstations.

While an HPC cluster can provide clear benefits when it comes to productivity, it can be challenging to justify the capital expenditure and support costs of deploying such a solution to management.

A turnkey, on-premise HPC cluster appliance based on 2nd generation Intel® Xeon scalable processors can reduce the cost and complexity of these deployments by shifting much of the burden to a third-party appliance provider. But organizations still need hard data on the potential return on investment (ROI) before making this transition — and fortunately, generating that data has gotten easier as well.

THE VALUE OF AN ON-PREMISE HPC CLUSTER

The benefits of an on-premise HPC solution are well known—by distributing simulation workloads across a cluster, engineers can run higher fidelity models and still get results faster. That means organizations can complete more iterations in less time, allowing them to determine optimal designs quicker and accelerate time-to-market for new products.

While an on-premise cluster is built on the same type of architecture as a cloud HPC solution, there are distinct benefits that an on-premise system offers. For organizations with a sufficient, ongoing simulation workload, an on-premise cluster can alleviate bottlenecks and improve performance.

Utilization rates determine the "rent-versus-buy" decision when it comes to HPC. For higher volume users, on-premise clusters are less expensive than cloud-based solutions in most cases. While there is some value in tapping into cloud resources to obtain burst capacity when there is peak demand for compute resources, an on-premise cluster (which operates much like a personal cloud) can be significantly less expensive than a public cloud for steady state usage. In industries such as defense or healthcare, on-premise clusters also provide a solution without the compliance and regulatory headaches that cloud computing can present in those environments.

An analysis for typical clients finds that HPC can be as much as 300% less expensive than the cloud at an average of roughly 1 cent per core hour when fully managed. When a single job is run across multiple nodes, an HPC cluster can also provide a 100% to 250% performance improvement compared to using a single workstation.

CLOUD VS. ON-PREMISE PRICING

OPTION	MINIMUM TIME TO PAY	UPGRADE HARDWARE?	SAVINGS OVER BASE ON DEMAND
On demand	**1** hour	YES	Base
Pre-pay up front	**1-3** years	NO	**40%-60%**
Pre-pay with upgrade option	**1-3** years	YES	**30%-50%**
On premise	**3** years	NO	**200%-300%**

Utilization drives the "rent vs. buy" decision. On-demand HPC resources can provide additional compute capacity, but they can be costly for high-volume users. As shown in the chart above, an on-premise HPC cluster can provide significant savings through reducing recurring costs and improved productivity.

BARRIERS TO ON-PREMISE HPC DEPLOYMENTS

While these benefits make sense from a productivity standpoint, cost and complexity have traditionally served as barriers to on-premise cluster installations. The cost of purchasing the hardware is high, but the ongoing costs of configuring, managing, and supporting the cluster over time can be even higher. Most engineering organizations lack in-house HPC expertise, and software licensing restrictions may further complicate the deployment. So even though there are clear operational benefits, they are often outweighed by support costs.

There are appliances available, such as the TotalCAE HPC Cluster Appliance, that can reduce both those costs and additional complexity by shifting this burden to the third-party provider, making the transition to an on-premise HPC cluster more cost effective.

Proving out the business case can still be a challenge for engineering teams that are not well-versed in typical HPC overhead and the ways in which IT support costs can be reduced via a managed solution.

A new free benchmarking program initiated by TotalCAE, Dell and Intel can help clearly define the return on investment in this type of solution. The benchmarking program can provide hard data on potential performance improvements and cost savings enabled by an on-premise HPC cluster appliance, using company-specific models to calculate benefits.

BENEFITS OF AN HPC CLUSTER APPLIANCE

An HPC Cluster Appliance based on 2nd generation Intel® Xeon scalable processors can reduce or eliminate barriers to on-premise HPC deployments. The TotalCAE software and services combined with Dell EMC HPC cluster hardware can reduce 83% of HPC TCO costs identified by research from IDC.*

An HPC cluster appliance combines software and hardware in a turnkey solution that is externally managed. This shifts the burden of tasks like storage, remote visualization, job scheduling and workload management from the engineering team to an appliance provider like TotalCAE. This enables organizations to enjoy the benefits of additional compute power without additional IT responsibilities and costs. With an easily installed and managed on-premise solution, engineering teams can more rapidly achieve benefits.

The improvements enabled by these appliances include:

EASE OF USE: With a single, web-based interface, an HPC appliance makes it easy for engineers to tap into the cluster. Engineers can submit a model for analysis, select the number of CPUs required, and click submit. The rest of the process is fully automated. You can see how the process works in [this video](#).
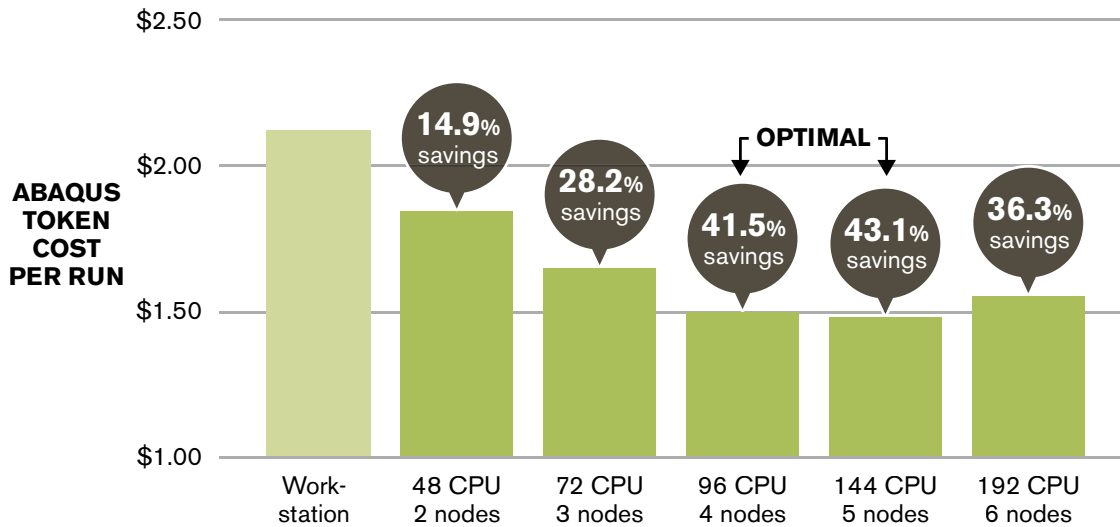
REDUCED IT EXPENSES: TotalCAE takes on all management and maintenance responsibilities for the cluster and associated CAE applications, which reduces the burden on the IT staff. Internal IT resources can be better optimized to focus on value-added projects, while the HPC cluster remains fully supported and updated.

RAPID DEPLOYMENT: Configuring and deploying an HPC cluster is no easy task, and internal staff may not have the expertise to do so. With an HPC cluster appliance, the solution can be set up in as little as 30 minutes thanks to pre-integration with more than 100 popular engineering applications.

CENTRALIZED LICENSE MANAGEMENT: Engineering firms often use multiple tools from a variety of vendors, and staying on top of licensing can be a drain on resources. Because the software is pre-integrated, license management can be centralized and simplified via an HPC cluster appliance.

RELIABLE SUPPORT: TotalCAE provides white glove support services and rapid, one-hour response times when problems arise. This helps keep engineering teams online and productive.

---

*IDC, Three Year Server TCO. Based on more than 300 interviews conducted across numerous platforms. Presented in composite form, 2007.

TOTALCAE HIGH-PERFORMANCE COMPUTING (HPC) BENCHMARK STUDY

This automotive supplier was able to reduce per-job costs using a managed HPC cluster with Intel Xeon Processors.

CASE STUDY: AUTOMOTIVE SUPPLIER

*This major Tier 1 automotive supplier
manufactures a variety of different vehicle components. The company was under
pressure to complete more design experiments per day, but faced a number of challenges.
The engineering team could not control or prioritize the simulation workload and was
not able to run simulations on a 24x7 schedule. The company also needed to
more efficiently utilize its existing Abaqus solver tokens.*

*An HPC cluster was needed to improve capacity and workload management,
but the company lacked internal expertise to successfully transition from a
workstation environment to an HPC environment.*

*Deploying the TotalCAE HPC appliance improved productivity, increased iterations,
and made the integration of their CAE applications seamless. The company can
run jobs around the clock, and TotalCAE was able to reduce run times by identifying
model optimizations. As a result, the company has experienced a 41% per-job license cost
savings compared to the previous workstation configuration running four nodes per job,
while simultaneously getting results back in less time (see chart on page 2).*

CASE STUDY: MEDICAL DEVICE MANUFACTURER

This company was running ANSYS jobs in a workstation environment, and wanted to reduce lengthy simulation wait times in order to reduce time to market. The firm also wanted a solution to manage workloads based on available licensing instead of self-scheduling jobs via inefficient phone calls.

The firm identified an HPC cluster and scheduler as a potential solution, but lacked internal expertise to deploy and support a cluster and the associated applications on it. TotalCAE deployed a solution that achieved 300% speedup for the client's models by utilizing multiple nodes per job, and enabled batch scheduling around their limited application licensing to reduce the burden on the engineer to manage simulations around license availability. The entire solution was managed by TotalCAE, enabling the client to focus on engineering, and not IT complexity.

BENCHMARKING HPC SUCCESS

The use of an HPC Cluster Appliance can reduce the cost to deploy an on-premise cluster and help improve engineering productivity. Quantifying those benefits can be difficult for engineering departments that need to justify the expense of the technology, but there are now tools available that can help engineering teams generate return on investment (ROI) data to help prove out the business case.

TotalCAE has partnered with Dell and Intel to offer a free performance benchmark program that will enable engineers and their managers to calculate the benefits and the advantages of moving from a workstation-centric environment to a Dell HPC cluster based on 2nd generation Intel® Xeon scalable processors managed by TotalCAE.

By accessing the benchmarking tools on this new, dedicated website, engineers can provide a model for analysis, and then receive data comparing the runtimes of the HPC solution with the performance of their current workstations.

In addition to providing the benchmark results, TotalCAE will provide a no-obligation quote for a TotalCAE HPC appliance with Dell hardware based on 2nd generation Intel® Xeon scalable processors that TotalCAE will manage for you in your data center. The TotalCAE HPC appliance makes it easy to submit simulations to the HPC system in a few clicks, and is completely managed by TotalCAE. End users can focus on their engineering tasks, without having to worry about IT support for the solution.

To get started, visit the [TotalCAE HPC Performance Benchmark website.](#)