


Speed Product Development via Virtual Workstation Clustering

How engineering departments can save time and money by tapping into their workstations' idle cores.

*Prepared by the editors of Desktop Engineering
on behalf of HP and Intel Corporation*



CONTENTS

Executive Summary	3
1. Searching for More HPC Power	4
2. Creating a Virtual Workstation Cluster	7
3. The Payoff	13
4. The Value of HPC to Engineers	15
5. Q&A with HP.....	17
6. Q&A with Intel.....	20

Appendix

A. Workstation Cluster Discussion Points.....	21
B. Additional Resources.....	22

EXECUTIVE SUMMARY

Parker Aerospace, a top-tier aircraft system manufacturer, was relying solely on its dedicated high-performance computing (HPC) cluster—a single HPC system—to process all its simulation and analysis jobs, no matter their size. The company is a heavy user of computational fluid dynamics (CFD) and finite element analysis (FEA) software to digitally test and optimize its designs. The HPC cluster was struggling to keep up with the ever-increasing workload. This created a long queue, forcing engineers to wait days or even weeks to get their turn, even for small and midsize jobs. The company realized it needed to bolster its HPC capacity to give its engineers more time to create high-quality products.

The traditional approach would have been to purchase and add more racks to the existing HPC infrastructure. Instead, Parker Aerospace adopted a powerful, less expensive solution. The company found a way to squeeze more computing horsepower out of its engineers' HP workstations—enough power to build a secondary HPC system—by assembling its workstations into a virtual cluster.

Powered by dual Intel® Xeon® 5600 series processors, each of the four HP Z800 Workstations in the cluster comprises 12 computing cores. Alone, they had more than enough processing power to meet the demands of the design engineers' daily tasks—a mix of 3D CAD modeling, document processing, emailing, and web browsing. With a modest investment in additional hardware and software for connectivity and clustering, Parker Aerospace was able to create a workstation-based HPC cluster to share the burden of its primary HPC system. Software providers' HPC-friendly licensing also made it possible for the company to deploy the secondary cluster at an acceptable cost.

Now, supplemented with a workstation-based HPC cluster, the company was able to dramatically reduce the job queue by moving small and midsize jobs to the workstation cluster. Once jobs get to the top of the queue, they are completed almost 60% faster.¹

Parker Aerospace's workstation-based cluster harnesses more computing power out of existing hardware, saving costs. But the most significant impact to the bottom line comes in the form of increased productivity. The company's engineers have the time and HPC resources they need to perform higher quality simulation and analysis as part of the up-front design process. They can now consider a greater number of design alternatives in their decision-making, ultimately leading to the creation of optimal products.

Workstation Cluster Benefits

- **Higher Quality:** Time for more design iterations using higher fidelity simulations. leads to optimal products.
- **Improved Workflow:** Eliminates days or weeks of waiting in the job queue of the dedicated HPC cluster.
- **Speed:** Cuts average analysis job runs by almost 60%.
- **Cost savings:** Uses computing power from existing hardware.
- **Scalability:** More workstations can be added to existing clusters, or new clusters can be created depending on network bandwidth.

1. Interview with Bob Deragisch, Group IT Manager, Engineering Services Manager, Parker Aerospace, 4 March 2011. To get the 60% figure, we divided the average time a job takes to run now (20 minutes) by the average time a job took to run before the workstation cluster (50 minutes), to get a savings of 60%.

1 SEARCHING FOR MORE HPC POWER

At Parker Aerospace, a division of Parker Hannifin, there are two high-performance computing (HPC) systems. One is a typical setup, a series of rack-mounted HP servers in an air-conditioned room. It occupies a sizable corner of the company's office in Irvine, CA. This system is reserved for the most demanding analysis jobs—or “big jobs,” as the company classifies them. The other system—reserved for small and midsize jobs—is a bit harder to detect. It is, in fact, invisible.

The second HPC system at Parker Aerospace is a virtual cluster, a distributed computing solution. It has its own job queue,

With most jobs taking days and sometimes even weeks to get their turns in the HPC queue, the wait time hampered Parker Aerospace's productivity.

swelling and shrinking as deadlines come and go. It has a software interface where people can submit jobs. But it doesn't live in a climate-controlled room like its larger cousin. It lives somewhere in the company's IT network, in a private cloud within the company's firewalls. To perform calculations, it uses idle computing cores in four of the engineers' Intel® Xeon® processor-based HP Z800 Workstations, and it does so without infringing on the engineers' daily computing tasks. Though invisible, it has significantly improved engineering productivity.²

Parker Aerospace's HPC Engineering Needs

Airplane component and system supplier Parker Aerospace designs, builds, and supports flight controls, hydraulics and fuel systems with thousands of parts. Expected to be in operation for 30 to 50 years, these mechanical assemblies have to be meticulously tested and validated for all conceivable flight conditions. The use of digital mockups—3D digital replicas of the aerospace systems that will be manufactured—has become the norm among top-tier suppliers like Parker Aerospace. Engineers in the company spend a considerable amount of time simulating, testing, and refining complex digital models using FEA and CFD software from ANSYS, LS-Dyna, and CFdesign.

Bob Deragisch, Engineering Services Manager for Parker Aerospace, is fond of saying, “We don't build airplanes, but we build just about everything that makes the planes fly.”

Deragisch's skills and responsibilities span both IT and design engineering. Though he began his career as a programmer, he has always worked with engineers and engineering applications. With an education in mathematics, he has an appreciation for the types of intense computations required to solve analysis scenarios, and he under-



Bob Deragisch, Parker Aerospace

2. Interview with Bob Deragisch, Group IT Manager, Engineering Services Manager, Parker Aerospace, 28 June 2011.

stands the engineer's need for the fastest possible turnaround time.

But Deragisch had a problem. The volume of data involved and the amount of calculations required made many of the company's simulation jobs unsuitable for individual workstations that must also be used for design. Consequently, Parker Aerospace had to rely on its HPC cluster to process them. As business and customer demands grew, so did the job queue's size. Unable to keep up with the ever-increasing workload, the HPC cluster became a bottleneck.

With most analysis jobs taking days and sometimes even weeks to get their turns, the wait time hampered productivity and tried the patience of engineers. That was when Deragisch and his colleagues realized they had to find a way to bolster Parker Aerospace's HPC resources.

A Creative Solution to the Bottleneck

The traditional approach would have been to purchase and add more racks to the existing HPC cluster. But in addition to the cost involved in upgrading the dedicated HPC cluster, Deragisch saw inefficiencies in that approach. The existing HPC cluster was meant to process large jobs, so it seemed counter-intuitive to expand it just to handle an increasing number of smaller jobs.



The multi-core HP Z Series Workstation can be equipped with two Intel® Xeon® processors.

“We made a conscious decision to identify and execute a new strategy in order to achieve greater value from our hardware and software investments,” noted Deragisch.

That idea led him—with the help of HP, Intel, and Microsoft—to a creative solution that lets him draw more computing power from the workstations already assigned to engineers.

Deragisch’s idea for a virtual cluster came from his recollection of the SETI@home project. The ongoing scientific experiment by UC Berkeley researchers aims to use leftover computing horsepower donated by home users to advance the search for extraterrestrial intelligence. By combining the powers of a large number of internet-connected home computers, SETI is able to create a virtual supercomputer powerful enough to digitally monitor and analyze radio signals from space. Deragisch realized that, using a similar method employed by SETI@home, Parker Aerospace could create a virtual HPC cluster out of idle computing cores in engineering workstations.

“ We made a conscious decision to identify and execute a new strategy in order to achieve greater value from our hardware and software investments. ”
— Bob Deragisch

The computing model—often called grid computing or cluster computing—is ideal for FEA and CFD applications because simulation software makers like ANSYS have invested significant time and effort to make parallel processing possible. “We did a lot of development work to optimize our software packages for HPC,” said Barbara Hutchings, ANSYS’ director of strategic partner-

ships. “So [HPC] is off-the-shelf capability. No extra work required from Parker Aerospace.”

Deragisch explained his idea to representatives from HP and Intel. As it turned out, the hardware companies already had a solution using multi-core HP workstations, on-chip virtualization technology from Intel and virtualization software from Parallels.

2 CREATING A VIRTUAL WORKSTATION CLUSTER


Deragisch had one main goal: to increase engineering productivity by breaking up the log jam of simulation and analysis jobs being run by ANSYS, LS-Dyna, and CFdesign on the company's dedicated HPC cluster. To do that with a workstation cluster, he figured he would need to cluster together at least two dozen computing cores, the equivalent of adding 12 dual-core workstations that were the norm just a few years ago.

The primary building block in Parker Aerospace's virtual cluster was the right hardware—in this case, HP Z800 Workstations with 32GB of memory and dual Intel® Xeon® 5600 series processors (each with six processor cores). These workstations had more than enough processing power to meet the daily demands of Parker Aerospace design engineers, who perform a mix of 3D CAD modeling, Microsoft Office document processing, emailing, and web browsing. More importantly, the workstations had enough unused power to enable the creation of a cluster.

Creating the virtual workstation cluster was simply a matter of deciding how many cores to dedicate to local vs. clustered tasks, networking the workstations together, providing sufficient power, and licensing the software.

Virtualization is Vital

The HP Z800 Workstations can come bundled with software from Parallels called Workstation Extreme, a virtualization solution that takes full advantage of Intel® Virtualization Technology for Directed I/O (VT-d). It creates the virtual machine on each individual workstation. This is the underlying infrastructure technology that made workstation clustering feasible at Parker Aerospace.



In most virtualization technologies, you use software to map the input and output among the virtual machines, the physical devices, and the computer ports. This is satisfactory for most usages; however, the high input/output (I/O) rate and performance required in many engineering applications is greater than what's possible with typical virtualization technologies. In most virtualized environments, drawing to the screen is significantly slower, and network speed is less than it is from the native operating system.

Workstation Cluster Requirements

- **Management:** After an initial setup to determine how many cores to dedicate to the local workstation, scheduling on the virtual cluster is managed by Microsoft Server 2008 HPC Edition.
- **Networking:** The faster the interconnects, the more efficient the cluster. Parker Aerospace created a private network via gigabit Ethernet for its four-workstation cluster.
- **Power requirements:** Parker Aerospace's cubicles weren't wired for the necessary 15-amp service needed by the workstations, so circuit breakers had to be upgraded.
- **Software licensing:** Licensing software for multicore execution could be cost-prohibitive, but many companies, such as ANSYS, have scalable licensing for clusters.

Parker Aerospace HP Z800 Workstation

Processor	2 Intel® Xeon® Processors X5650 (6 cores, 2.66 GHz, 12MB cache, 1333 MHz memory) with Intel® VT-d
Software	Genuine Windows® 7 Professional 64-bit OS. Parallels Workstation Extreme virtualization software.
Chipset	Dual Intel® 5520
Graphics	NVIDIA Quadro 4000 graphics card with 2GB cache memory
Form factor	Rackable minitower
Memory, standard	32GB 1333 MHz DDR3 ECC Registered RAM
Memory slots	12 DIMM slots
Internal drive bays	4 internal 3.5-in. bays
External drive bays	3 external 5.25-in. bays (up to 4 eSATA, up to 8 SAS)
Storage controller	LSI MegaRAID 9260-8i SAS 6Gb/s ROC RAID
Optical drive	SATA SuperMulti DVD writer
Hard drive	500GB 7200 rpm SATA NCQ
Ports	9 USB 2.0, 1 audio in, 1 audio out, 1 headphone out, 2 microphone in, 1 serial, 2 RJ-45 to integrated Gigabit LAN, 2 IEEE 1394a
Slots	1 PCI, 1 PCI Express Gen1 (x8 mechanically, x4 electrically), 1 PCI Express Gen2 (x8 mechanically, x4 electrically), 2 PCI Express Gen2 (x16 mechanically, x8 electrically), 2 PCI Express Gen2 x16 1 22-in-1 Media Card Reader
Network Interface	Integrated dual Broadcom 5764, TPM 1.2 Controller, optional Broadcom NIC, optional Intel NIC

In its Intel® Xeon® chipset, Intel provides for direct access from a virtual machine to I/O devices and ports, eliminating the need for a higher-level—and slower—software interface. Instead, devices can deliver I/O just as they would in the native operating system’s environment, so performance is comparable to the native environment. Such a setup is necessary in graphics display and network access, where high performance is imperative.

To produce robust performance, the virtualization software must take advantage of Intel VT-d. Parallels Workstation Extreme is the first to do so. But it goes beyond delivering high performance I/O. Parallels Workstation Extreme also provides the ability to tune resource utilization on every workstation. It’s possible, for example, to allocate four cores and 8GB of memory on one workstation to a design engineer’s interactive tasks, and eight cores and 16GB of memory for the same purpose on another system in the same cluster.

“The historical problem with doing clustering on workstations is that the cluster job can be unconstrained. If it’s a complex computation, it can take all cores and all memory, essentially bringing the interactive user to a halt,” said Dan Bennett, an HP Technical Consultant for the HP Z Workstations series. “But with Parallels Workstation Extreme, you can configure the cluster portion of the computer to use no more than, say, six of the 12 cores. And that means that while the cluster-based portion is going flat-out, doing a great job at your analysis, there’re still plenty of cores and memory and graphics totally at the disposal of the engineer sitting at the workstation.”

The first step to determining how to divide the workstation cores and memory is to understand the processing characteristics of the work being done locally, according to Deragisch. Some engineering applications can make use of all of the cores and memory available, while others require a great deal of network bandwidth. Cluster administrators have to take into account the needs of the application and the type of work being done in order to appropriately configure each workstation.

Users of CAD software, for example, typically only use a small number of cores when running their applications. After Parker Aerospace allocated individual cores and memory for the desktop operating system, productivity applications, and any other applications, it was able to leave six or eight cores of each of the four 12-core workstations available for clustering.

This one-time setup of each individual application (e.g. ANSYS, Microsoft Office, LS-Dyna, CFD Design) on each clustered workstation took anywhere from a few minutes to a few days for Parker Aerospace, depending on the characteristics of the jobs and underlying applications. Some were

very predictable, while others varied a great deal. But this only had to be done once to determine how many cores to initially allocate. Once the computational behavior of a task was ascertained, Parker Aerospace could allocate the appropriate amount of computing resources to its purpose.

“The nice thing about Parallels Workstation Extreme is that you can configure it on the fly,” said Mike Diehl, the HP Z800 Workstation Project Manager. “So if an engineer finds that occasionally compute tasks in the virtual machine are treading on his work, he and the administrator can change the resources allocated to the cluster task to use fewer cores or less memory.”

The virtual cluster is managed by Windows Server 2008 R2 HPC Edition, which runs on the Parallels Workstation Extreme virtual machines. It manages memory, starts applications, interfaces with the I/O ports, and manages the workload, just as an operating system on a non-virtual machine would.

Once the virtual machines were up and running, they just needed to be networked together in a cluster configuration to become a comprehensive engineering resource for the entire workgroup.

Making the Network Work

A crucial aspect of Parker Aerospace’s workstation-based HPC solution was the network of interconnects linking the computing cores and memory. In order to pass data and instructions back and forth within the cluster without creating performance-killing delays, the clustered workstations needed to be connected with a



The HP Z800 Workstation can be easily customized with more memory, network interface options and processors, making it well-suited to building workstation clusters.

With greater access to HPC, design engineers can simultaneously study multiple design options, allowing them to select the best option based on performance, strength, endurance, and cost.

network pipeline that was significantly faster than most corporate local-area networks (LANs).

“Our workstations are connected to what I call the public network, which is shared by all of the other systems in the building,” Deragisch noted. That network provided the data communication needs of the engineers using the workstations for daily tasks. However, because of congestion caused by normal corporate network traffic, it was simply not feasible to run a workstation cluster on the same network.

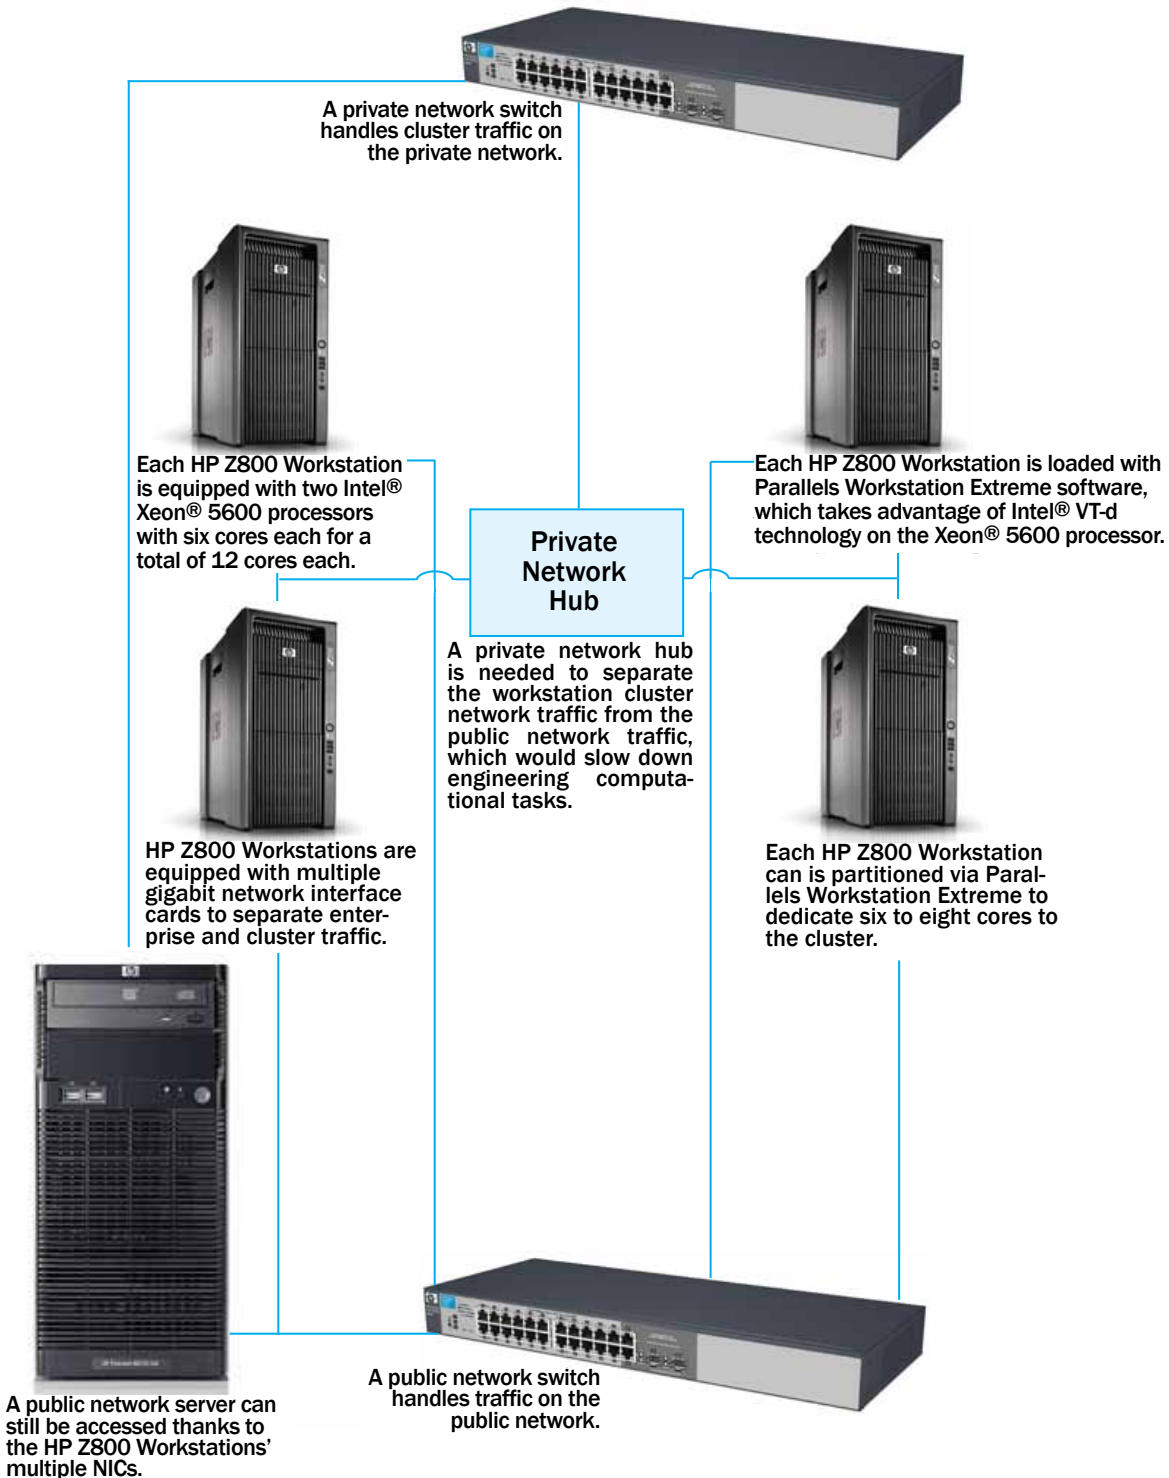
Therefore, Deragisch had to deploy the virtual cluster on a separate private network, confined to the clustered machines’ transactions only (see Exhibit A on page 11). This was possible on the HP Z800 Workstation, which came with dual network interfaces, both facilitating gigabit Ethernet connections. The second network was dedicated solely to workstation cluster traffic (again using the VT-d capability of Parallels’ software), making it possible to get data and instructions to and from different machines quicker with little contention for network bandwidth.

Once the systems are networked together in a cluster configuration, it becomes a comprehensive engineering resource for the entire workgroup.

Aside from the hardware setup, the software code’s ability to execute analysis and simulation jobs on multiple CPUs with efficiency also proved a critical factor. “Because you’re taking a single computation, dividing it into pieces, and running it on several processors, at some point those processors have to communicate with one another, because the pieces of the problem they’re working on are interdependent,” explained ANSYS’ Hutchings. “We have done a ton of software tuning to optimize the message-passing component of our products.”

Parker Aerospace already had the building wired for gigabit Ethernet, which would be more than adequate to handle the traffic of the 24- to 32-core virtual cluster Parker Aerospace was building. But Deragisch ran into an obstacle when it came to plugging the workstations into the power outlet. The Z800 requires a 15-amp circuit. The company’s power panels were not wired to provide 15 amps to the cubicles, so they had to upgrade the circuit breakers and make sure the power loads were adequate.

Exhibit A: Parker Aerospace's Workstation Cluster



The Parker Aerospace workstation cluster is connected by a private gigabit Ethernet interconnect that provides fast I/O directly to the cluster. It is also connected via a second gigabit Ethernet to the public enterprise network, which provides all normal network services such as Web, email, and backup.


Multicore Software Licensing Solutions

Another challenge to workstation clustering can be software licensing. Because some software makers have not yet caught up with the evolving use of HPC, a software package's licensing policy can sometimes hinder those who'd like to deploy the software on tens, hundreds, or thousands of cores.

"Some software we use becomes prohibitively expensive when running on dozens of cores, because we'll need a license for each core," Deragisch explained. "Getting ANSYS HPC Pack offered us a significant advantage because it is not priced per core."

At Parker Aerospace, engineers use ANSYS Mechanical, ANSYS CFX, and ANSYS Icepak for structural, fluids, and electronics thermal management, respectively. ANSYS' Hutchings noted, "We really want to encourage our customers to take advantage of HPC so they can examine their designs from the system level, not at the component level. ANSYS HPC pack licensing allows extreme scalability at an incremental cost. It can take advantage of hundreds or even thousands of cores."

The software licensing was the last piece of the puzzle. Deragisch and his team had linked four HP Z800 Workstations together and assigned each of their Intel® Xeon® processors with a total of 12 cores to create a virtual HPC cluster that could fluctuate between 24 and 32 cores. And the benefits were immediate.



Simulation and analysis software providers, such as ANSYS, are offering HPC licensing packages that allow engineers to cost-effectively take advantage of cluster computing.

3 THE PAYOFF

The workstation-based cluster enabled Parker Aerospace to use its dedicated HPC cluster more efficiently and tap into the power of its workstations' idle cores. But, most importantly, it allowed the company's engineers to work faster and smarter.

More Time for Quality

With the days and weeks spent waiting in the long queue for their turn a thing of the past, Parker Aerospace engineers use their extra time and resources to refine designs iteratively. They can create an initial design, analyze it on the workstation cluster, then make changes interactively based on analysis results. They are able to quickly go through multiple cycles of design and analysis using higher fidelity models (those with greater mesh density and more geometric details), yielding more accurate results. This added productivity ultimately gets them closer to the best design option based on performance, strength, endurance, and cost.

"We're not finishing our work more quickly," Deragisch is quick to point out. "We have contractual dates, and we meet those dates. What we are doing is delivering better designs—designs that, in many cases, exceed requirements."

By executing small and midsize jobs submitted by design engineers, the workstation-based cluster leaves the HPC server cluster free to concentrate almost entirely on larger, more complex projects. That not only reduces the productivity-stealing HPC queue, it means more jobs can be processed simultaneously by both the workstation cluster and the dedicated HPC cluster, which contributes to the almost 60% time reduction in analysis jobs.³

“What we are doing is delivering better designs—designs that, in many cases, exceed requirements.”
— Bob Deragisch

"Analysis jobs used to run 40 to 60 minutes," noted Deragisch. "Now the average runs are 15 to 25 minutes. The queue length is dramatically shortened."

Previous workstation cluster models could only run jobs "after hours" when the machines were not being used for other work. But being able to design while running analysis and simulation concurrently is a game changer.

Cost Savings

In studying cost-benefit ratios, Deragisch and his colleagues concluded, "[Adding] a second six-core (or more) processor on a work-

3. Interview with Bob Deragisch, Group IT Manager, Engineering Services Manager, Parker Aerospace, 4 March 2011. To get the 60% figure, we divided the average time a job takes to run now (20 minutes) by the average time a job took to run before the workstation cluster (50 minutes), to get a savings of 60%.

station is far less expensive than expanding the rack or just adding an additional ‘similar-core’ processor on the HPC server.”

If an engineering organization has performed a hardware upgrade in the last two or three years, it may have already invested in a fair number of multi-core workstations. Since complex simulation and analysis jobs have become standard practice in the industry, creating an HPC cluster to process these tasks out of existing workstations gives an organization the benefit of an HPC server without the cost and expense of purchasing and maintaining a dedicated HPC server.

“The results speak for themselves,” said Deragisch. “We have significantly increased our throughput and we are using all the available technologies at Parker Aerospace to expedite design decisions from applications like ANSYS, CFdesign, and LS-Dyna.”

For Deragisch, besides the cost savings and his professional pride, there’s another motivation—a personal one—that drove him to tirelessly seek ways to improve Parker Aerospace’s designs. With a son who has just been promoted from first officer to airline captain, he has good reasons to make sure he delivers the most robust internal components for aircraft in operation.

“ [Adding] a second six-core (or more) processor on a workstation is far less expensive than expanding the rack or just adding an additional ‘similar-core’ processor on the HPC server. ”
— Bob Deragisch

4 THE VALUE OF HPC TO ENGINEERS

All of Parker's results have been achieved with just four workstations. However, more workstation cores could be added to a cluster, or additional workstation clusters could be created to multiply Parker's results. Because more cores dedicated to the cluster equals more data transmission, network bandwidth is the only limitation to scalability.

Even with gigabit Ethernet connectors (a fast network by today's enterprise standards), data transfer could quickly reach the network's maximum limit. Anything beyond eight workstations would cause too much contention for network bandwidth via gigabit Ethernet, hurting the performance of many applications.

"We're looking closely at other interconnect alternatives for the future," said Deragisch. Such alternatives include Fibre Channel and InfiniBand, both of which have theoretical bandwidths reaching well into tens of gigabits per second. As communication technology improves, 10 gigabit Ethernet may also become a possibility.

It's easy to imagine the benefits of scaling up a cluster to 64, 128 or more cores. For now, the four-workstation cluster has been so successful that Parker Aerospace is in the process of implementing two more such clusters, one of which will use eight workstations. The goal is to enable all design engineers to iterate their design and analysis work more quickly.

Improved Workflow

At present, the cost and effort involved in setting up and deploying HPC resources prevent some small and midsize manufacturers from taking advantage of sophisticated computer-aided analysis and simulation. Even among many who have discovered the benefits of HPC, the use of the technology is often limited to validating a concept, or proving that a product would perform as intended in practice.

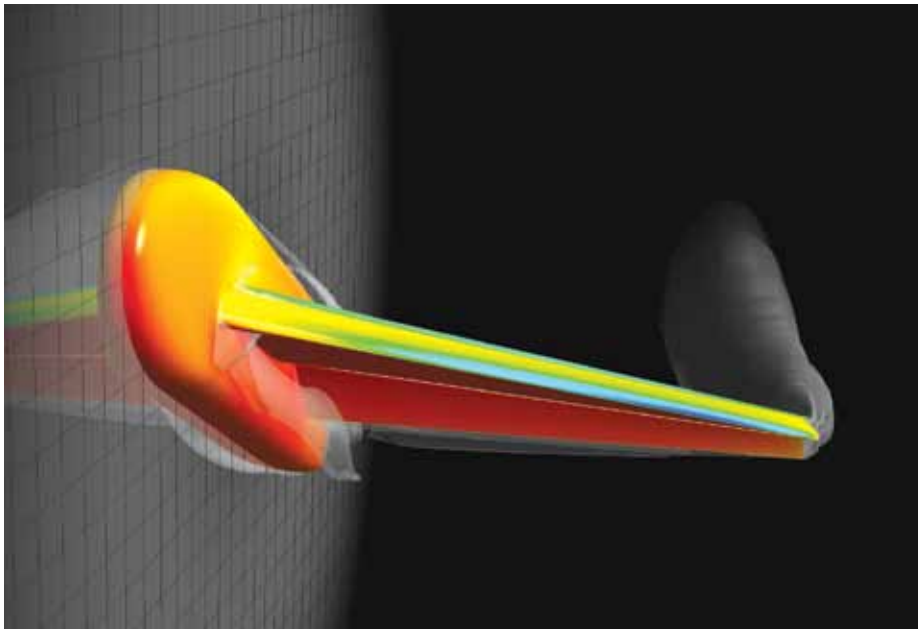
HPC solutions that are affordable and easier to implement—like Parker Aerospace's workstation-based cluster—are expected to change user attitudes. Simulation software makers' evolving licensing practices, as exemplified by ANSYS's HPC licensing, will also make cluster-driven FEA and CFD workflows more affordable and easier to manage. Instead of restricting HPC-enabled analysis and simulation to a few authorized individuals, manufacturers may encourage more users to employ the technology throughout the entire design cycle. Users would, for example, begin identifying the most promising concepts in early phase, then spend the rest of the development cycle perfecting the design.

"It's that idea generation that you want to accelerate because it gives engineers the opportunity to test more 'what ifs,'" says Wes Shi-

manek, Strategic Marketing Manager of Technical Computing at Intel Corporation. “If engineers can simulate more of those earlier in the design, they can get closer to that perfect product.”

Tom Salomone, HP segment manager for CAD and Design Engineering Industries, agrees.

“It fundamentally changes the way engineers work, because you can now be designing, running an analysis locally, looking at the results, incorporating that to refine the design, then running more analysis on the workstation cluster to refine it still more,” he says. “Analysis and simulation becomes a fundamental part of the design process, as opposed to doing your analysis post-design. That’s a huge step forward for the engineering community. It’s going to change cycle times for the design process, and it’s going to change how engineers understand their designs. There’s a significant amount of change ahead with this concept.”



Workstation-based clusters support making simulation and analysis an even more fundamental part of the design process.

5 Q&A WITH HP

The workstation-based HPC cluster at Parker Aerospace was made possible by the nature of the company's principle hardware—the HP Z800 Workstation. *DE* contributing editor Peter Varhol interviews HP executives to discuss the characteristics of the specific workstation model that made it ideal for clustering. Responses provided by Tom Salomone, HP Segment Manager for CAD and Design Engineering Industries; Dan Bennett, an HP Technical Consultant for the Z Workstations series; and Mike Diehl, the HP Z800 Workstation Product Manager.

Peter Varhol: What do you consider the key design characteristics of the HP Z Workstation series?

HP: *The key things we want to deliver to our customers are a very high-performance, highly reliable workstation that can stand up to the demands of their work. We want them to be high quality, and easy to deploy and easy to service, with little or no use of tools, and rack-mountable.*

PV: What are the primary characteristics of the HP Z Workstations that make them adaptable for workstation-based cluster computing?

HP: *The HP Z series family provides great computing power, and a great interactive experience for someone sitting at the workstation. At the same time, you are using otherwise unused resources on the workstation to contribute to a distributed compute solution. The HP Z800 Workstation is especially appropriate for this type of usage pattern, in particular because it's a dual-socket design; that is, it can have two Intel Xeon processors, each of which can have up to six cores. Another key, and a very important point in terms of the technology included in these workstations, is the use of Intel's Xeon processors, because they support a technology Intel calls VT-d, short for Intel® Virtualization Technology for Directed I/O.*



HP's Tom Salomone, Dan Bennett and Mike Diehl.

PV: Why is the support for Intel VT-d so important?

HP: The way we're implementing this workstation cluster computing model is to use Parallels Workstation Extreme, and it uniquely takes very good advantage of Intel's VT-d technology. That allows their software to implement virtual machines on the workstations, which can be directed to use only certain parts of the workstation.

PV: The traditional way of clustering workstations is often known as cycle stealing, but that's not at all what Parker Aerospace is doing here, correct?

HP: Correct. Any time you're running multiple things on a workstation, you're competing for resources, and unless there's an intelligent way of putting bounds on one of those tasks, so that it does not steal the resources of the workstation, it's going to slow down the engineer interacting with the computer. And that's what Parallels Workstation Extreme does. It takes advantage of some of the design characteristics of the HP Z800 Workstation. For example, the dual network interconnect provides a means for clustering on a fast private network, while also giving the interactive engineer a separate connection on the enterprise local network. That really helps in terms of keeping workloads separate.

PV: And that's also important because the cluster workload is dependent upon a fast network interconnect, right?

HP: That's correct, although it depends on the workload. Some applications and jobs have to move a lot of data, and would really benefit from this configuration. For example, the HP Z800 Workstation has enough power and available slots to be able to put two high-end graphics cards into the workstation. The engineer sitting at the workstation has access to one of those graphics cards for work, while the other card could incorporate a GPU that could be used to accelerate the execution of the cluster job, if the application supports GPU computing. That would be very useful for jobs that have a high ratio of floating point operations. And GPU companies like NVIDIA are reaching out to application companies to support GPU computing as a part of their solver engines. Combined with the HP Z800 Workstation running Parallels Workstation Extreme virtualization solution, this could be a very powerful solution.

PV: Do engineers give up any performance to lend a part of their workstation to cluster computing?

HP: That depends on the application being used by the engineer. There are applications, such as ray tracing, that can use all available cores, and reserving some of those cores for a cluster will necessarily have a performance impact. But if you're building an assembly in a CAD application, or modeling a component, multiple cores don't matter.

PV: What engineering software vendors have you collaborated with in the Z Workstation series design?

HP: *When we go to software vendors that produce applications, like an ANSYS or an Autodesk, or PTC or Siemens PLM software, we work with them extensively. We have engineers who are assigned to them, and are often on site at their facilities. We do testing, we run benchmarks and look at how those benchmarks exercise various pieces of the equipment. So we understand that when someone runs an application just what that application does to stress the system, and where those stresses are, so we can design appropriately. We also take our newest workstations and test them through what we call a certification process, which sometimes takes as much as three weeks of testing, to determine that the hardware is perfect for their application. We also talk to these engineering software vendors about their roadmaps and make sure that our products support their future needs. If their offerings in the future will require a lot of memory, for example, or more power at the network level, we work to understand and incorporate their needs in our future designs.*



HP Z800 Workstation with dual six-core Xeon 5600 processors.

6 Q&A WITH INTEL

In building its workstation-based cluster, Parker Aerospace relied heavily on Intel's technology to configure and facilitate a parallel-processing workflow among disparate HP Z800 Workstations. DE contributing editor Peter Varhol interviews Wes Shimanek, Strategic Marketing Manager of Technical Computing at Intel Corporation, to gather insights on the technologies working behind the scenes.



Intel's Wes Shimanek.

PV: How has the Intel Cluster-Ready program influenced the ability of vendors to create hardware and software for workstation clusters?

WS: *Intel Cluster-Ready is an initiative by Intel to work with the software and hardware community to help businesses building clusters to quickly load and go. Everyone thinks building and managing a cluster is highly complex. Intel makes it easy. We work with an ISV (independent software vendor)—ANSYS is an example—to validate certain configurations, and if you use those configurations, it's just that simple.*

PV: Can you explain the importance of Intel VT-d technology (Intel® Virtualization Technology for Directed I/O)?

WS: *Intel VT-d enables you to create hard partitions between resources. It is these hard partitions that help users get the most out of their workstation for CAD and simulation at the same time. Both run at near-native performance. With Intel VT-d, users will see terrific CAD performance while concurrently processing CAE simulations. This will help engineers create, test, and modify in a very fast innovation loop.*

PV: And there's no performance degradation for the user, and the cluster is pretty high performance?

WS: *That is correct. Interactive CAD and CAE simulation processing requirements are both running at near-native performance with the resources they are using. But remember that Intel isn't operating alone here. There is a company called Parallels that has software that does the partitioning and takes advantage of Intel VT-d. With Parallels Workstation Extreme software, users can have near-native performance on both the interactive partition and the cluster partition. And it doesn't just have to be Windows. You can be running Windows interactively and Linux on the cluster.*

PV: How do clusters add value to engineering?

WS: *There are three important assets you have to optimize.*

The first is the engineer. If you give them the tools they need to iterate through more ideas in less time, you're getting more value.

The second is when you take a dual-processor workstation, and software from Parallels or Microsoft, you can partition that machine to take advantage of all of its processing power. So now you've not only optimized the engineer, you've optimized the workstation, taking advantage of all of its cycles.

The last asset is any investment in HPC resources. When you use workstations, you can remove the burden of small and midsize jobs from your HPC resources. By removing these small and medium jobs, you're letting your fastest HPC systems and clusters do what they were designed to do—run the really big jobs really fast. More importantly, as Bob Deragisch points out, you get out of the HPC queue, and the net result can be that you actually get small and medium jobs back amazingly faster. You have now turbo boosted your company's pace of innovation.

Appendix **A**

WORKSTATION CLUSTER DISCUSSION POINTS

Use the discussion points below to begin a conversation with your managers and information technology colleagues about the benefits workstation clusters could have for your organization.

- **Higher quality:** The ability to use the power of clusters while engineers work, as opposed to using clusters only after hours, allows workgroups to be more efficient. Analysis and simulation become part of the design process, instead of being relegated to a post-design process. Running more simulations earlier in the design process can result in better products by allowing engineers to refine designs. It can also encourage use of higher-fidelity simulation and analysis to be run for improved accuracy.
- **A good investment:** If you have multi-core workstations, or plan to invest in them, they will likely have unused cores at times. For example, using CAD software often does not max out multi-core workstations, while running multi-threaded simulations can. Clustering makes the most of your multi-core workstation investment by sharing those unused cores with the workgroup.
- **Improved efficiency:** Clusters allow companies to remove the burden of small and mid-sized jobs from their high-performance computing resources, allowing them to dedicate HPC systems to running large jobs more quickly.
- **Ease of setup:** Building and managing clusters has been made simple by software and hardware manufacturers, who have come together to create standard configurations of their products that are designed for easy clustering.
- **Virtual cluster benefits:** Modern clusters can be set up to not steal cycles from other engineering tasks, as clusters have in the past. For example, Parallels Workstation Extreme, combined with the functionality of the HP Z800 Workstation and Intel's Xeon and VT-d technology, can be used to create virtual machines that only use a set amount of the workstation's resources. The virtual partitions don't have a noticeable impact on local performance while sharing resources with the cluster, and don't have to only run jobs after hours.
- **Ease of management:** Parallels Workstation Extreme can easily be configured to use fewer cores on the workstation for those times an engineer may need more cores. When the engineer does not need as many cores, that compute power can be released to the cluster to benefit the workgroup as a whole. Partitioning workstation resources also enables the IT department to easily manage and update those resources as needed.
- **Scalability:** Additional workstations can be added to a workstation cluster. The only limitation is the speed of the interconnect. If that is a concern, multiple separate clusters can be created to achieve many of the same clustering benefits for workgroups.

Appendix **B**

ADDITIONAL RESOURCES

- HP Z800 Workstations
<http://h10010.www1.hp.com/wwpc/pscmisc/vac/us/en/sm/workstations/z800.html>
- HP Workstation Clusters and ANSYS
www.hp.com/go/wsansys
- HP Workstations
www.hp.com/go/zworkstations
- Intel® Workstation Processors
www.intel.com/go/workstation
- Intel® Workstation Virtualization
<http://www.intel.com/technology/virtualization/>
- Intel® Cluster Ready
<http://software.intel.com/en-us/cluster-ready/>
- High Performance Virtualization on the HP Z800 Workstation
http://download.parallels.com/doc/pwe/en/HP_and_Parallels_Workstation_Extreme_High_Performance_WP.pdf
- Detailed information from Parallels
www.parallels.com/products/extreme/partners/hp/
- Free trial download from Parallels
<https://sites.google.com/a/myballincollig.com/parallels-workstation-6-extreme/>
- COE Ask The Experts Webcast Aug. 30, 2011 with Dan Bennett: High Performance Workstation Compute Cluster enabled with Parallels® Workstation Extreme — found in the “Ask the Experts Archive.”
www.coe.org/EventsEducation/AsktheExpertsWebcasts/tabid/227/Default.aspx



This whitepaper is sponsored by Hewlett-Packard and Intel® Corporation. It has been produced by *Desktop Engineering* as an educational resource for the engineering community.

© 2011 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. Intel, Xeon, and Core, are trademarks of Intel Corporation in the U.S. and other countries. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Parallels Workstation Extreme is a trademark of Parallels Holdings Ltd. Images on pages 12 and 16, and screen image on page 19 are courtesy of ANSYS, Inc. Screen image on page 5 is courtesy of Phoenix Analysis & Design Technologies, Inc.