

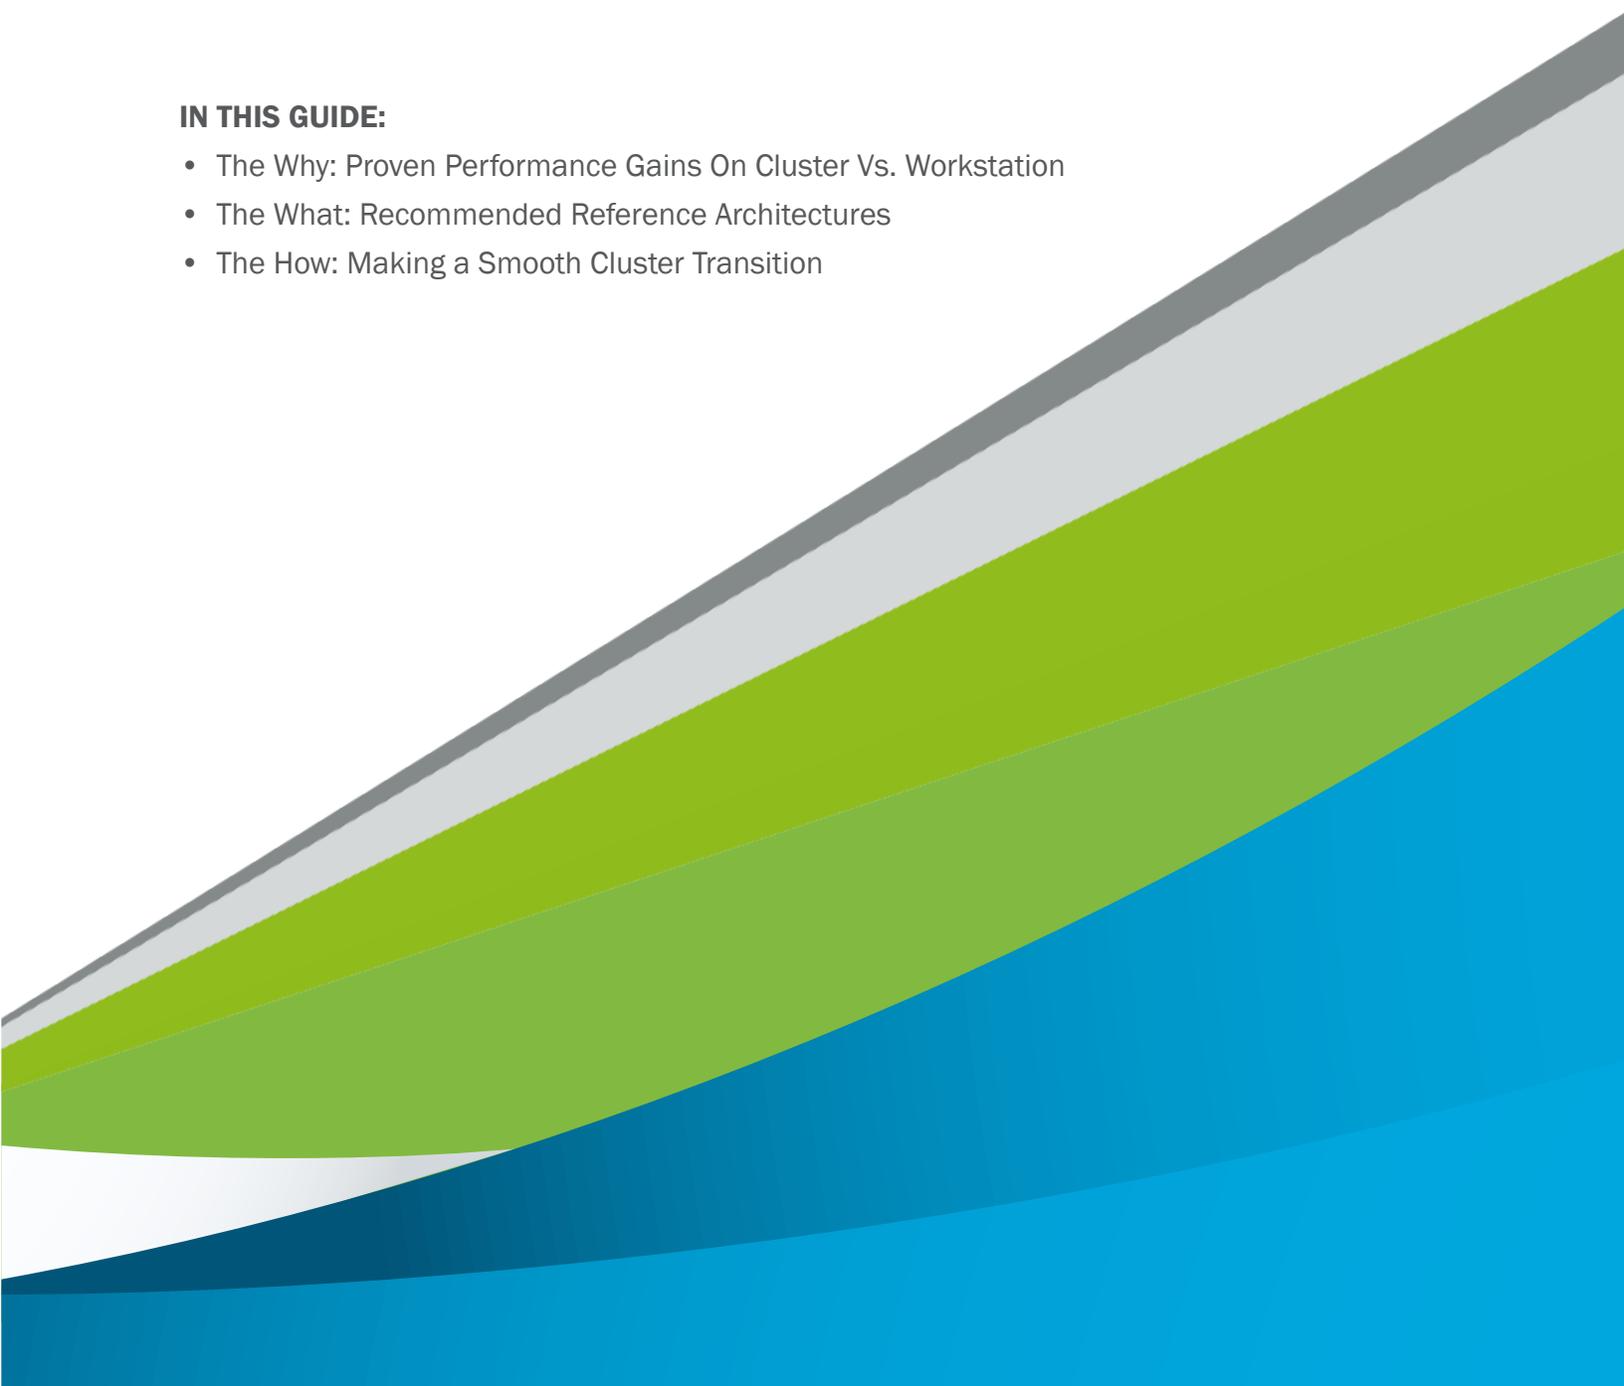


Altair
Innovation Intelligence®

Cluster Transition Guide: Scaling from Workstation to Cluster for Compute-Intensive Applications

IN THIS GUIDE:

- The Why: Proven Performance Gains On Cluster Vs. Workstation
- The What: Recommended Reference Architectures
- The How: Making a Smooth Cluster Transition



INTRODUCTION

It's now an established truth that cluster computing represents the best, most cost-effective way to scale performance for professional applications. In this guide, Altair and Acer will present the business case for switching to a cluster environment, including documented performance gains when running a sample commercial application on both workstations and clusters. We'll also present recommended configurations and outline key considerations and steps for choosing and deploying a professional-grade cluster.

WHY CLUSTER? TO ACHIEVE MORE

Though new to many organizations, clusters have been around for over 30 years and have replaced much of the mainframe market and proprietary supercomputing designs that previously existed. A quick look at the "Top500 Supercomputers" website (www.top500.org) – a list of the world's most powerful high-performance computing (HPC) machines) offers a glimpse of the prevalence of clusters in the marketplace.

But that's just the tip of the iceberg. Top500 does not take into account the hundreds of thousands of smaller clusters – many as small as just a few nodes – that provide meaningful results to researchers, engineers, technicians and others on a daily basis. Considering this data, it's clear that clusters dominate technical computing – so gains in cluster performance stand to impact application performance (and thus productivity and innovation) across a wide range of areas.

In short, good computing means good clustering means good business.

Newcomers to cluster computing could at first be daunted by the seemingly complex nature of clusters. Previously unfamiliar technologies such as message passing interface (MPI), complex compilers, Linux, and other such tools are intrinsic in over 90 percent of existing clusters. However, we are long past the early days of cluster computing – professional grade software such as Altair HyperWorks® has long since been optimized

"In the past dozen years, standards-based clusters have become the dominant species of high-performance computing (HPC) systems.

"Clusters represent about two-thirds of HPC servers sold today... IDC predicts that in 2016, the market for large-scale clusters in the supercomputer and divisional segments will reach \$7.3 billion."

"Clusters of all sizes will perform most of the work in the rapidly emerging server market for high-performance data analysis (HPDA), which IDC forecasts will grow robustly (13.3% CAGR) from \$743.8 million in 2012 to about \$1.3 billion in 2016."

– IDC, 2013

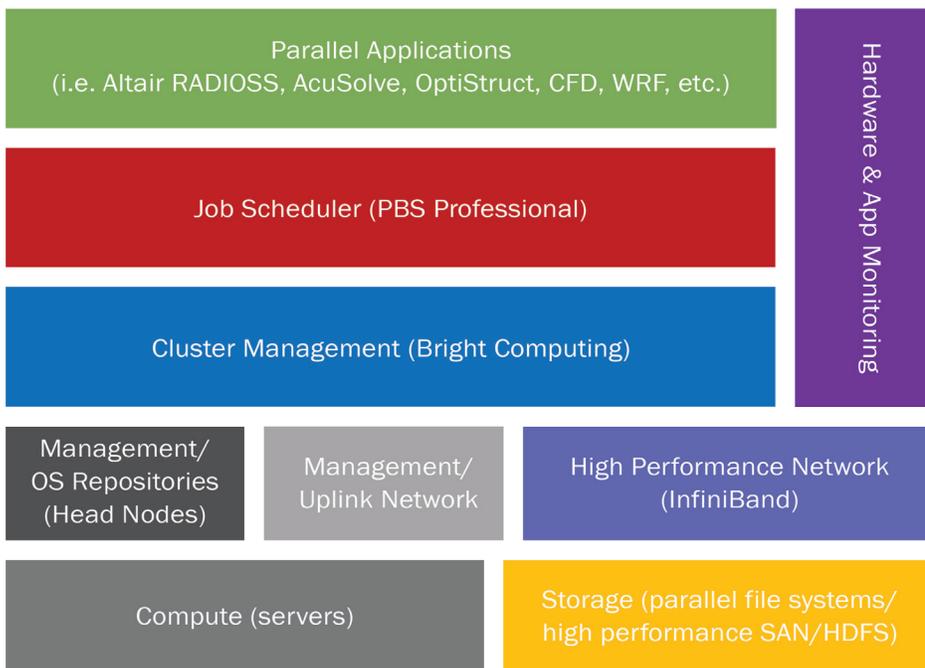


Figure 1: Cluster architecture and software components

A short list of top cluster benefits includes:

- Complete more jobs faster
- Create better products by running more simulations
- Increase workplace productivity by minimizing wait time
- Reduce risk by improving understanding of products before they are released
- Lower overall costs through consolidation

to take advantage of these very tools and can quickly be deployed and running, even for users unfamiliar with a cluster setup. Thus, most commercial applications are ready to take advantage of clusters now.

Detailed in Figure 1 (see previous page), the elements necessary to build a cluster include a mix of hardware, networking, management software and job scheduling / management tools that work in tandem to allow multiple users to run their workloads uninterrupted on a single cluster, starting with as few as five nodes. (For details on some of the key software elements in this stack, see pages 5-6).

PROOF POINTS: PERFORMANCE GAINS ON CLUSTER VS. WORKSTATION

Application users know that clusters will accelerate their processing and help get better results faster, which accounts for the dominance of cluster computing in HPC (see IDC quote on previous page). So it's no surprise that benchmarking in general for user applications on clusters tends to show substantial benefits.

For the purpose of this guide, the benchmark data provided shows the results between a single Windows 7 workstation versus a Linux-based cluster running a short simulation. The application used was Altair's RADIOSS solver, an engineering simulation application used for stress analysis including mechanical analysis of an impact – for example a crash test. This benchmark test leveraged a Neon car structure in the initial 8 milliseconds (ms) of a crash.



Figure 2: Performance improvements when scaling to clusters

- (1) Veriton P530 F2 - E5-2670, 8 x 8 GB, Windows 7
- (2) Altos R380 F2 - E5-2670, 8 x 4GB, Linux
- (3) AW2000h-AW170hf F2 - E5-2670, 8 x 4GB, GbE, Linux
- (4) AW2000h-AW170hf F2 - E5-2670, 8 x 4GB, QDR IB, Linux
- (5) 2 x AW2000h-AW170hf F2 - E5-2670, 8 x 4GB, GbE, Linux
- (6) 2 x AW2000h-AW170hf F2 - E5-2670, 8 x 4GB, QDR IB, Linux

RADIOSS is structural analysis software used to simulate real-world scenarios in order to test the performance of a product. A highly scalable solver, RADIOSS helps improve crashworthiness, safety, and manufacturability of structural designs. It can also use multiphysics simulation and advanced materials such as composites. Around for over 20 years, Altair's RADIOSS software exemplifies the advantages of cluster computing by demonstrating high performance at scale.

Result highlights:

1. 4 X increase in performance
2. 45% energy savings†
3. 50-70% employee time savings (meaning engineers spend more time engineering than waiting for results)
4. Better resource TCO (users can use terminals or lighter PCs for modeling, and can share the cluster resources for more compute-intensive work – evolving their work habits and ensuring all hardware investments are used to their greatest potential)

† Compares the maximum rated watts for the Veriton P530 F2 (900W) and AW2000h F2 (1600W) and the number of jobs that can be processed by each. The calculation is meaningful in two ways: a workstation will often equip a larger amount of additional components, such as graphics cards, which are not used in CPU calculation. Secondly, the ability of the AW2000h F2 to share redundant power supplies across four nodes helps decrease the amount of energy loss from the AC-DC conversion, having fewer power supplies in total can lead to significant conservation in aggregate power draw.

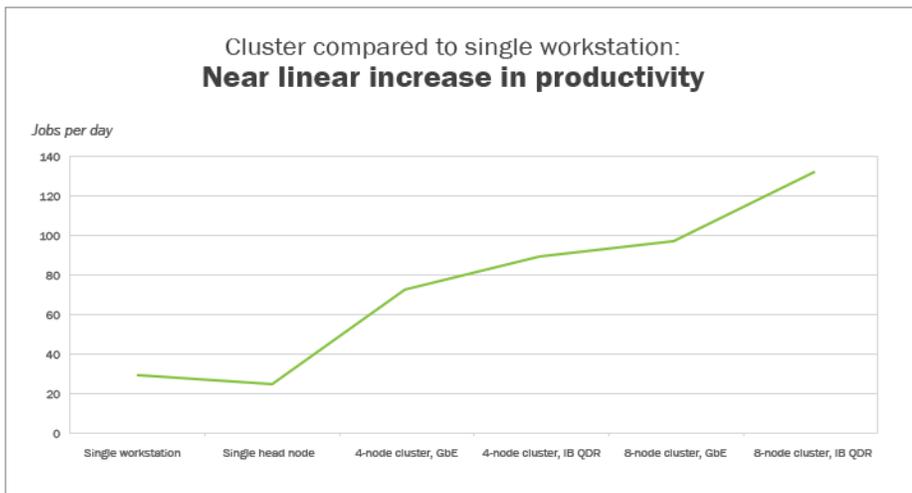


Figure 3: Increase in jobs per day (assumes 8 hour work day; more jobs can be completed if left to calculate overnight)

Proven results

The medium sized cluster will deliver 4x more jobs in a given day, and even more if jobs are queued to run overnight and on weekends.

REFERENCE ARCHITECTURES

For users who are unfamiliar with clusters – but looking to begin using them and enjoying the proven benefits – the configurations below offer a starting point for consideration.

	Single user / smaller job size	Multiple users / large job size	Departmental use / multiple large jobs
System	Veriton P530 F2	Altos R380 F2 AW2000h F2 (4 nodes)	Altos R380 F2 2 x AW2000h F2 (8 nodes)
Configuration	2 x Intel Xeon E5-2650 v2 64 GB DDR3 RAM NVIDIA Quadro K2000 1 TB 3.5" HDD Gb Ethernet dual port Windows 7 Pro	Per node 2 x Intel Xeon E5-2650 v2 64 GB DDR3 RAM 1 TB 3.5" HDD* Gb Ethernet dual port CentOS, Red Hat or SUSE Linux	Per node 2 x Intel Xeon E5-2650 v2 64 GB DDR3 RAM 1 TB 3.5" HDD* Gb Ethernet dual port 56 Gbps FDR InfiniBand single port CentOS, Red Hat or SUSE Linux
Cluster management	N/A	Bright Cluster Manager	Bright Cluster Manager
Job scheduler	N/A	PBS Professional	PBS Professional

Figure 4: Recommended small, medium and large cluster configurations

TRANSITION GUIDE: MAKING A SMOOTH CLUSTER TRANSITION

Clusters do not have to be complicated for the end user. Thankfully, clusters have been in use long enough that best practices have emerged and both hardware and software vendors have had years to focus on streamlining their products with ease-of-use in mind.

By leveraging standard Acer hardware with Bright Computing's Cluster Manager and Altair's PBS Professional for cluster management and job scheduling, users have all the building blocks they need to setup and run a professional-grade cluster.

Here are the high-level steps required to move to a cluster environment:

1. Identify the application

- Understand which application is providing the largest bottleneck for your user base
- Ensure this application is a cluster candidate (beneficial though they may be, clusters are designed to work with only a certain set of applications, and users must understand if their desired application can or cannot be run on a cluster)

2. Prepare the application to run on a cluster

- Inquire if the application offers a cluster version – for most HPC applications, this will be the case.
- If not, is the software portable to a cluster? Questions to ask:
 - Is it open source?
 - Can the code be changed?
 - Can it take advantage of parallel processing and MPI?

3. Determine the supported hardware for the application

- Many applications are CPU based, but some can leverage GPUs or coprocessing units (such as Intel Xeon Phi). Understanding which computing platform is best suited for your application will help determine the hardware plan and costs / benefits for scaling to a cluster.
- Altair PBS Professional supports Xeon Phi out of the box and offers a configuration tool to easily update an existing environment for Xeon Phi. Altair has been a key partner for Xeon Phi development and was among the first commercial software vendors to support the coprocessor both on its workload manager and also for commercial applications. For more information about Altair's collaboration with Intel, visit www.altair.com/partner-intel.

4. Understand the performance benefits desired

- Cluster optimized applications like RADIOSS can scale fairly linearly across several nodes, and for CPU-based applications, adding more nodes will result in an approximately even increase in performance.
 - As shown in our performance benchmarks, for smaller clusters Ethernet could be a sufficient interconnect – but as performance demands and/or usage increases, larger clusters will be able to take advantage of high-performance, low-latency interconnects such as Infiniband.
- Power requirements and scalability are also important. The Acer AW2000h F2 offers maximum 1/2 U density and solid performance. However, users looking for balanced CPU/GPU or coprocessor performance could prefer the Acer Altos R380 F2. The application's compute needs will ultimately determine the system design, and Acer can help offer recommendations based on the usage scenario.
- Manageability is critical for both cluster administrators and users.

What cluster users are saying about Altair and PBS Professional:

“PBS Professional offers fast, easy implementation so our cluster users could get up and running quickly.”

– *Czech National Supercomputing Center*

“Altair's expertise and dedication to implementation success is unbeatable...we get much better cluster utilization.”

– *Weizmann Institute*

“Everyone has benefited from this installation. Our scientists can now deliver research outcomes faster and more accurately – we're able to process much more data a lot more quickly. Researchers can also do things they couldn't do before, like running jobs multiple times or utilizing larger data sets, to achieve outcomes they couldn't otherwise get. It's pretty easy to see the value.”

– *QIMR*

- Acer and Bright Computing together offer a robust and proven solution through Bright Cluster Manager on Acer's high-performance systems. Offering both GUI and command line interfaces, Bright Cluster Manager removes the barriers to clustering and helps simplify complex tasks for software/hardware integration, monitoring, package management and more.
- For systems with larger jobs and multiple users / departments and software stacks, Altair's PBS Professional offers the perfect job scheduling toolset for managing multiple application requests across any number of nodes, from handful to thousands. Developed over 20 years ago in collaboration with NASA Ames research center, this robust job management system is customizable so experienced users can setup automatic job management systems based on specific needs. The system also provides backfilling, shrink-to-fit, and many other features to increase cluster usage and ensure users can get their results fast. PBS Professional is one of the most widely deployed cluster workload managers on the market.
- Compilers and debuggers are also often helpful for developers who are porting software to cluster environments. Not always necessary for commercial software, open source software users could make good use of software such as Intel Cluster Studio, and Allinea debugging toolset to identify sluggishness, and bad code areas when results seem to be taking too long.
- Professional grade, cluster ready applications. Never to be overlooked, the application is the heart of a cluster, and will determine the performance attainable more than any other part of the system. Users who are looking to cluster must first discuss what options exist with their software vendor, and can then refer to other hardware and cluster management options for their final planning.

5. Begin using and track usage and performance data

- It is important to understand how your cluster is being used, so the system can be fine-tuned and optimized for your specific environment. Acer and Altair recommend a cluster analytics tool such as PBS Analytics to access and visualize historical cluster usage, performance analytics and run reports to support data-driven planning and decision making. Easily extensible to meet unique requirements, data from multiple PBS Professional servers can be consolidated, providing a global view of cluster usage for chargeback, capacity planning, troubleshooting and project management.

WHY ALTAIR AND ACER?

Acer is the fourth largest provider of personal computers in the world, and offers a global presence for service and support across all of its platforms. Having focused on the high-performance computing industry for only a few years, Acer has proven its ability to scale performance and offer support for systems of tens of thousands of cores in size. The company's portfolio of high-performance workstations, servers and personal computers makes it an ideal partner for organizations seeking a single provider to fulfill all their IT needs.

Altair is a global software leader with a 27-year-plus track record for HPC workload management, advanced engineering software, on-demand computing technologies and enterprise analytics. Altair knows HPC: We are the only company to create HPC solvers and simulation applications, deliver HPC management tools, and use these applications and tools to solve real HPC issues. With more than 2000 employees in 20 countries worldwide, Altair experts serve over 3000 clients who use our market-leading products to improve their product quality, infrastructure performance and ROI, and time to market.

Together, Acer and Altair are among the most respected leaders in computing and offer the worldwide, world-class expertise needed to ease the transition to a cluster environment.

"The Acer-Altair integration offers a new option for HPC consumers who need robust, scalable, high-end systems to support their compute-intensive work."

– *PBS Works CTO*

"Thanks to this [Altair-Acer] partnership, Acer is able to offer professional-grade job scheduling to customers seeking robust performance in their clustered environment."

– *Acer Business Development Manager EMEA-Technical & Productivity Computing*

GETTING STARTED

Ready to begin the process of switching to a cluster – or just looking for advice on whether a cluster is a good fit for you? Let the experts at Altair and Acer help.

We can offer:

- Free initial cluster transition assessment for your environment
- Free trial licenses of cluster workload management
- Options for consulting and onsite training

To get started, contact us at:

- HPCsolutions@acer.com
- info@altair.com

You can also access more information online at:

- <http://www.pbsworks.com/acer-altair.html>