ALTAIR

EBOOK

SUPERIOR CAE PRODUCTIVITY WITH HPC APPLIANCES FROM ALTAIR AND INTEL

CONTENTS

1. INTRODUCTION 2
2. THE CRITICAL ROLE OF CAE IN PRODUCT DESIGN AND DEVELOPMENT 2
   - Diverse CAE Applications 2
   - Increasing Demand for Simulation Capacity 3
3. COMPLEXITY IS BARRIER TO HPC ADOPTION 4
4. HOW ENGINEERING TEAMS DEAL WITH THESE CHALLENGES TODAY 4
5. THE ALTAIR UNLIMITED APPLIANCE 6
   - Altair HyperWorks Software 7
   - Pre-integrated with Everything You Need 7
   - Built for performance with Intel® Xeon® Technology 8
   - Intel-optimized Altair Solvers 9
   - GPU Nodes 9
   - Reduce Costs and Boost Performance with Efficient Resource Sharing 9
6. TAILORED TO YOUR ENVIRONMENT 10
   - A Turnkey Managed Service 11
   - Built with Security in Mind 11
   - Flexibility and Scalability built-in 11
   - Scale to The Cloud with Ease 11
7. COMPELLING BUSINESS VALUE 12
8. CONCLUSION 12

INTRODUCTION

High-performance computing (HPC) has emerged as an essential technology in modern design and development environments. Driven by increased competitive pressures, manufacturers increasingly compete based on their computer-aided engineering (CAE) and HPC infrastructure environments.

Deploying and managing HPC infrastructure is challenging, however. HPC clusters are complex and costly to deploy and require specialized skills across multiple domains to manage and maintain them effectively. As a result, engineering teams can spend much of their time managing the HPC environment rather than focusing on more productive work. This can lead to higher costs and longer design cycles due to reduced productivity.

This paper will explore this dilemma and explain how turnkey HPC appliances from Altair and Intel can dramatically simplify HPC cluster deployment and management. It will show how the Altair Unlimited™ appliance powered by state-of-the-art Intel® Xeon® processors can reduce simulation time, reduce costs, and improve overall productivity and efficiency.

THE CRITICAL ROLE OF CAE IN PRODUCT DESIGN AND DEVELOPMENT

CAE plays a critical role in product design, development, testing and manufacturing across multiple industries. Simulating and validating product designs in software enables engineers to improve product quality and durability and to reduce the need for physical prototyping. They can also optimize the use of materials and improve processes, resulting in more reliable products that are easier and less costly to manufacture. This translates into faster time to market, reduced manufacturing costs, and lower downstream service and warranty costs, directly impacting the manufacturer's bottom line.

Diverse CAE Applications

CAE is used in multiple industries up and down the supply chain, from automotive and aerospace to electronic design to consumer goods to pharmaceuticals. The steps involved in a CAE process typically include preprocessing, model solving, post-processing and optimization.

In the preprocessing stage, engineers use design tools to model a design's geometry and physical properties. Models are typically "meshed" into small, discrete components. Parameters are assigned describing material properties, the environment, and various boundary conditions, loads, and constraints.

Next, the mathematical model representing the design is "solved" by simulators (also called solvers) that predict how the design will behave under real-world conditions.

Finally, in the post-processing phase, the simulation results are presented for review and analysis. Some of the techniques used in CAE to address the underlying physics include:

- Finite element analysis (FEA)
- Dynamic analysis using implicit and explicit finite element modeling techniques
- Computational fluid dynamics (CFD) and thermal analysis
- Multi-body simulations (MBS)
- Electromagnetism

Engineers frequently used simulation-based techniques such as designs of experiment (DoE) and numerical optimization to explore designs systematically. Optimization workflows are used to analyze and optimize attributes across multiple areas, including topology, strength, vibrations and fatigue. With automated optimizations, organizations can more easily validate model performance leading to lighter weight, more structurally efficient designs.

Increasing Demand for Simulation Capacity

As the quality of simulation tools has improved and HPC infrastructure has become more powerful and cost-effective, engineers have more incentives to simulate designs in software. Simulations have also become more compute and data-intensive for reasons that include:

- Larger, more complex designs
- Finer-grained mesh sizes
- Increased use of composite, non-uniform materials
- Increased use of multiphysics simulators
- Use of photo-realistic digital prototypes

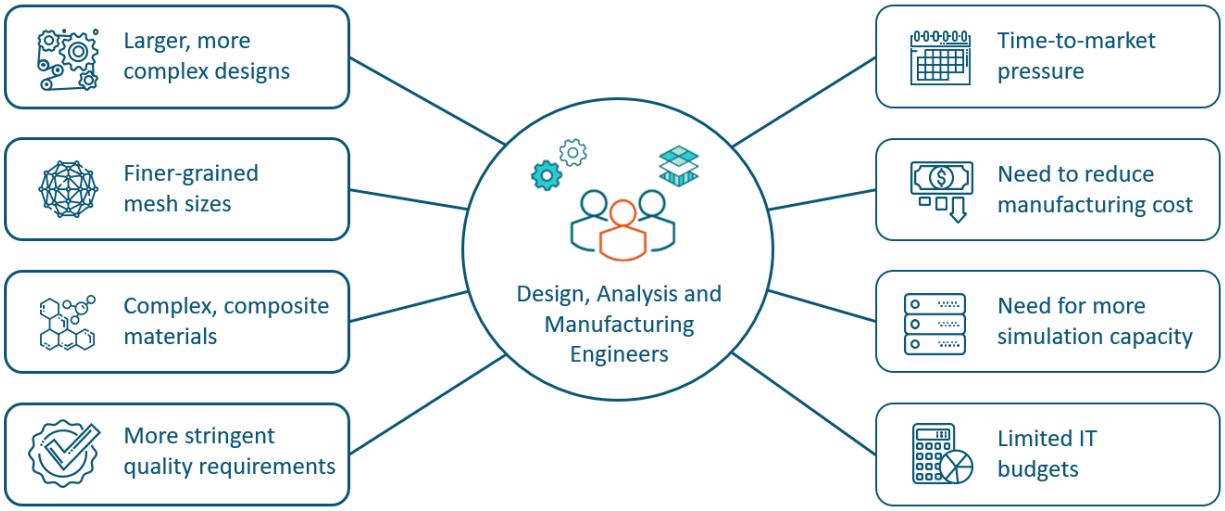

Figure 1 – Product design and development teams are challenged to do more with less

Design and development teams are under pressure to increase capacity to meet project deadlines and stringent quality requirements despite several constraints, as illustrated in Figure 1. Challenges include tight project deadlines, the need to contain manufacturing costs, and limited design and IT budgets.

COMPLEXITY IS BARRIER TO HPC ADOPTION

While design and development teams can benefit significantly from HPC, HPC clusters are complex to install and manage. This is especially true for small and mid-sized design teams that may lack dedicated IT support or have limited access to skilled system administrators. Customers face several barriers when deploying in-house HPC environments, including:

- **Infrastructure management** – The cost and complexity of acquiring, deploying, and managing on-premises infrastructure such as compute clusters, storage subsystems, and interconnects.
- **Complex HPC middleware** – HPC clusters typically involve multiple software prerequisites such as shared or distributed file systems, workload managers, MPI libraries, and drivers for GPUs, host bus adapters, and other specialized hardware.
- **Software licensing** – License management can be a headache for design teams. Organizations need to purchase licenses for various solvers, implement policies to avoid oversubscription (that can lead to job failures at runtime), and share licenses fairly among team members.
- **Optimizing performance and utilization** – Even with an HPC cluster in place, optimizing resource utilization can be complex. Workload management policies need to be created that maximize job throughput, share resources fairly, and keep expensive hardware, software and personnel fully utilized and productive.

The more time that design engineers spend on the above issues, the less time they have available to focus on their core job responsibilities — designing better, higher-quality products.

HOW ENGINEERING TEAMS DEAL WITH THESE CHALLENGES TODAY

Some large organizations may have dedicated teams to manage HPC applications and infrastructure. However, most small and mid-sized organizations experience the challenges described above. How organizations approach HPC infrastructure management depends on their size, in-house capabilities, and available budgets. Figure 2 illustrates some common approaches to dealing with HPC infrastructure challenges described below:

- **Shared, dedicated servers** – Small design firms may run powerful shared servers in their CAE environment instead of a cluster to offload simulations. While this approach frees up desktop cycles and improves simulation capacity somewhat, it is still limiting. System administrators still need to manage servers and applications. With this approach, resource sharing is manual, limiting efficiency, scalability, and productivity.
- **On-premises HPC clusters** – Organizations with more internal IT support capacity may elect to deploy an on-premises HPC cluster. HPC clusters typically include a head node, multiple compute nodes, a shared storage array, and workload management software. An HPC cluster can dramatically improve throughput, but they are complicated to deploy and manage, and teams may experience wait times for resources shared across teams within their organization.
- **Cloud-deployed HPC clusters** – Another solution is to deploy HPC clusters on public cloud infrastructure-as-a-service (IaaS) platforms. While cloud deployments can reduce the need for on-premises infrastructure, they can bring additional challenges. System administrators still need to install and manage applications and HPC software. Also, they need to be proficient in managing cloud services via cloud toolkits or CLIs and worry about securing the cloud environment (configuring IAM[1], inbound

[1] IAM refers to Identity and Access Management. Implementations vary by cloud provider, but setting up IAM in accordance with best practices is critical to ensuring that cloud accounts remain secure and that control is maintained over spending. An example of best practices here - https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

and outbound security group rules, VPNs, etc.) Administrators and users also have the added burden of managing cloud spending to avoid overshooting IaaS budgets.

- **HPC SaaS offerings** – For some customers who do not want to manage software and infrastructure, cloud-based software-as-a-service (SaaS) may be a good solution. However, users who need access to multiple tools may find themselves running across multiple clouds, complicating workflow and data management. Also, HPC SaaS offerings frequently deploy resources dynamically for each application separately. This can lead to long turnaround times and inefficiencies given the need to provision cloud infrastructure for each simulation, leading to higher costs in the cloud.

Figure 2 - How design teams solve HPC infrastructure challenges today

Unfortunately, none of these solutions described above are ideal. On-premises clusters can impose a significant system administration burden, making them difficult for small workgroups to manage effectively.

Cloud-based IaaS or SaaS solutions help in some cases. However, they do not always reduce the management burdens related to applications and middleware and bring their own additional challenges related to security, cloud administration, data management, and potential issues around lock-in, cost transparency, and cost management.

Ideally, design and development teams need the convenience and ease-of-use of a turnkey SaaS solution, but with the performance, flexibility, and cost-efficiency of a secure, on-premises HPC cluster. Fortunately, there is a good solution that addresses these concerns — a managed, on-premises HPC appliance from Altair and Intel that fits seamlessly into a customer's private cloud.

THE ALTAIR UNLIMITED APPLIANCE

Altair Unlimited is a turnkey, state-of-the-art physical appliance optimized for FEA, CFD, and other simulation workloads. It is available in various standard rack configurations for on-premises deployments and can scale to thousands of cores. The appliance is easily tailored to individual customer applications and business requirements. Altair Unlimited bundles unlimited-use Altair HyperWorks™ solver software for simulating mechanics, fluids, electromagnetic, and more, as well as integrated software tools for modeling, visualization, and design exploration and optimization.

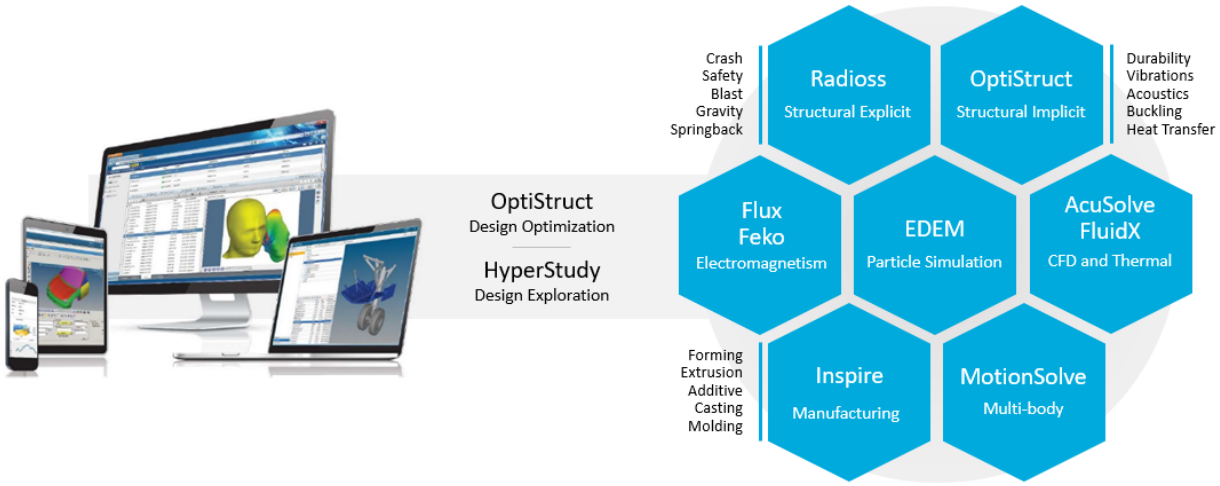

Figure 3 – Altair Unlimited provides seamless access to leading multiphysics solvers

Unlike traditional on-premises HPC cluster deployments, Altair Unlimited is a fully managed service. The cluster is built using state-of-the-art Intel Xeon-based server technology. It includes Altair PBS Works™ software for workload management, scheduling, remote visualization, notification, and collaboration. All software is pre-installed, making for a package that is exceptionally easy to use and manage.

For busy design teams, the Altair Unlimited appliance can dramatically boost engineering productivity and efficiency. It does so by improving simulation performance and increasing capacity for simulation-intensive workloads such as design of experiments (DoE) and numerical optimizations. With Altair Unlimited, designs teams can:

- Dramatically reduce job setup time and turnaround time
- Run more, higher-fidelity simulations to explore the design space thoroughly
- Easily collaborate with other users, both local and remote
- Avoid concerns about software licenses limiting simulation capacity
- Spend less time managing infrastructure and more time engineering products

Access from desktops, tablets, and mobile phones is provided via Altair Access™, a simple, consistent interface for submitting and monitoring jobs pre-integrated with the appliance. Users interact with integrated desktop tools or an intuitive, user-friendly web portal. Engineers can submit and manage simulations, monitor progress, and use built-in remote visualization services to view results.

Altair HyperWorks Software

At the heart of the appliance is HyperWorks®, a comprehensive suite of integrated tools for design, simulation, and visualization. HyperWorks provides intuitive direct modeling and editing, supports various multiphysics solvers, and delivers a complete post-processing environment to visualize, query, and process data.

HyperWorks provides easy-to-learn workflows that leverage domain knowledge and increase team productivity for faster, more efficient design. HyperWorks tools free engineers to maintain a model-centric view of their designs while quickly moving between multiple physics disciplines. HyperWorks produces robust designs that accurately model structures, mechanisms, fluids, electromagnetics, electrical, embedded software, systems design, and manufacturing processes. HyperWorks desktop tools such as Altair HyperMesh™, HyperGraph™, and HyperView™ work seamlessly with the scalable Altair Unlimited appliance.

The Altair Unlimited open architecture also makes it possible to use third-party software in a bring-your-own-license model.

Pre-integrated with Everything You Need

The appliance frees system administrators from tedious details such as OS installation, patching, scheduling policy, license management, and various other HPC system administration tasks. HyperWorks solvers are pre-integrated with Access application templates, providing users with an intuitive, customizable interface for job submission. Altair® PBS Professional® is preconfigured with appropriate policies to ensure that simulations run optimally, exploiting single- or multi-node parallelism to complete simulations as quickly as possible.

The appliance illustrated in Figure 4 is comprised of standard server building blocks to deliver optimum performance, throughput, and value for compute-intensive HPC workloads. Node types can include:

- Cluster head node – The head node runs the Access web portal and PBS Professional workload management software. It typically provides 85 TB (usable) of redundantly configured shared storage to support multiple HPC application workloads.
- Visualization nodes – Users can remotely visualize graphics-intensive simulation on visualization nodes equipped with NVIDIA® Quadro RTX™ 4000 graphics cards and remote visualization software. Visualization nodes can double as simulation nodes during busy periods or overnight hours when remote visualization services are not required.
- Explicit nodes – The appliance typically contains dedicated nodes optimized for explicit mechanical solvers such as Altair Radioss™ or CFD solvers like Altair AcuSolve™, with large numbers of cores, high clock speeds, and fast, local, SSD-based scratch storage to speed simulations.
- Implicit nodes – Nodes optimized for memory-intensive implicit simulators such as Altair OptiStruct™ may also be included. The workload manager ensures that HyperWorks and third-party simulators are dispatched to optimal node types to maximize throughput and turnaround time.

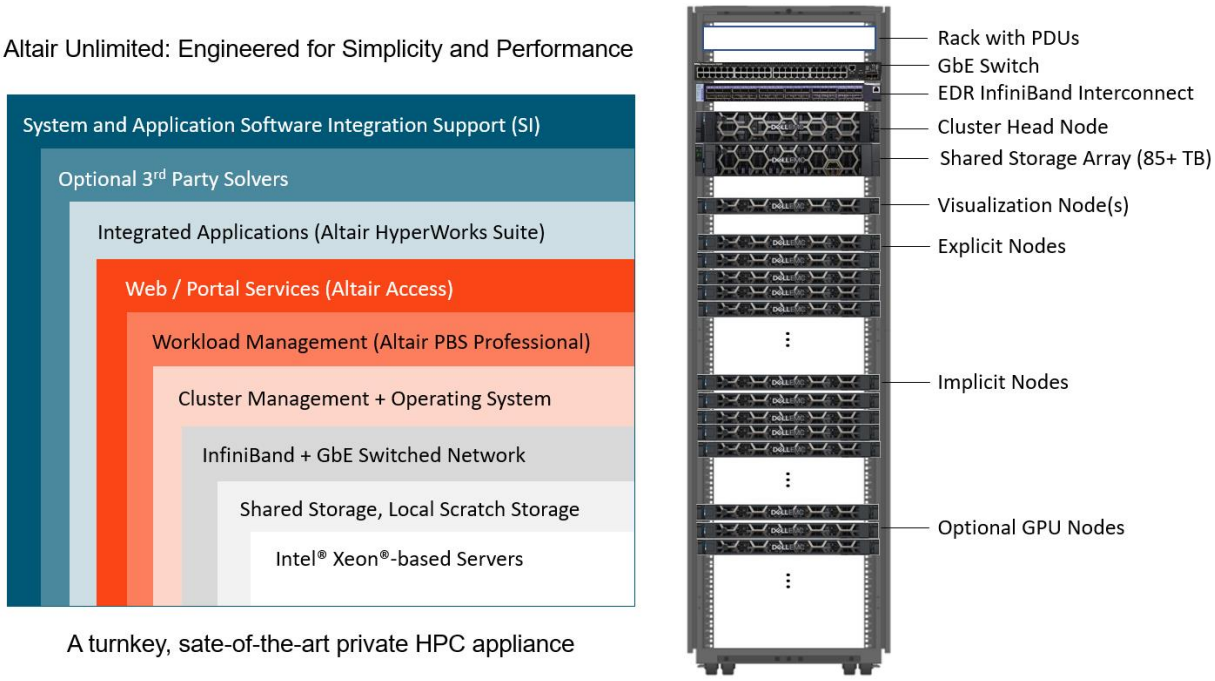

Figure 4 - The Altair Unlimited physical appliance

The appliance includes Gigabit Ethernet (GbE) and EDR InfiniBand (IB) switches connected to each rack mount server node. The InfiniBand interconnect is used by distributed parallel solvers with Intel® MPI to accelerate parallel simulations. File sharing, workload management, and remote visualization run over the separate GbE network. All of the bundled software shown in Figure 4 is preconfigured and ready to run optimally on the node types installed in the appliance.

Built for performance with Intel® Xeon® Technology

Altair Unlimited™ appliances use state-of-the Intel Xeon Scalable processor shown in Figure 5 to deliver best-in-class performance and value for a wide range of HPC and engineering workloads. The latest Intel processors deliver higher core/thread counts and industry-leading clock frequencies with Intel® Turbo Boost Technology. They also provide large amounts of cache L2/L3 cache, six or eight DDR-4 memory channels per socket, and Intel's blazing fast Ultra Path Interconnect (UPI) technology to deliver maximum performance in an economical, power-efficient package[2]. These improvements combine to deliver up to 36% better performance and 42% better price/performance than previous generation processors[3].

[2] Altair includes pre-qualified Intel Xeon processor SKUs. DDR-4 memory channels per socket vary depending on the processor. The processor technology used in each appliance node type may vary with time. Contact Altair for additional details on the particular SKUs used for various nodes types.

[3] https://wccftech.com/intel-2nd-gen-xeon-scalable-cascade-lake-refresh-server-cpu-family/

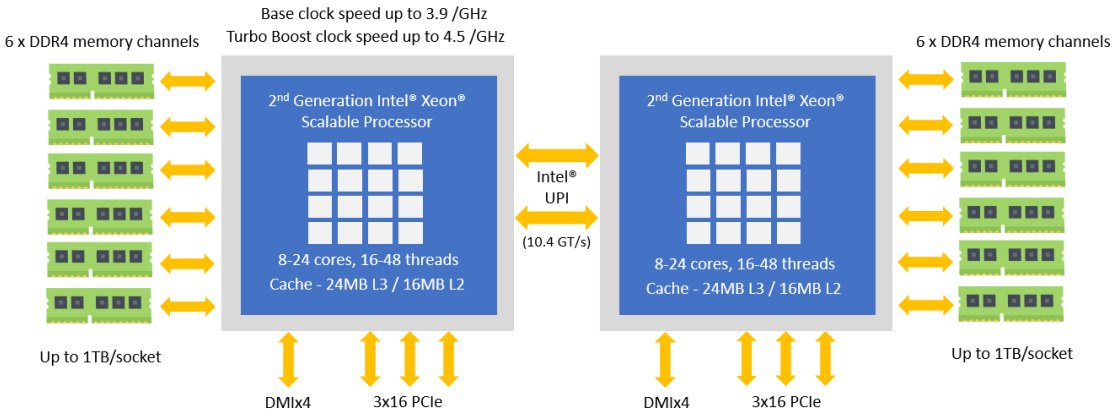

Figure 5 – Altair Unlimited appliances use state-of-the-art Intel Xeon processors[4]

Intel-optimized Altair Solvers

HyperWorks solvers benefit from several advanced architectural features of Intel Xeon Scalable processors. Altair solvers are built using Intel compilers, and Altair Unlimited ships with Intel MPI. This high-performance, multi-fabric message passing library delivers optimum performance on HPC clusters based on Intel processors. Compute-intensive solvers such as Radioss are designed to take advantage of Intel® Advanced Vector Extensions 512 (Intel® AVX-512) to accelerate the floating-point calculations prevalent in Radioss simulations.

GPU Nodes

Increasingly, HyperWorks and third-party solvers leverage high-performance general-purpose GPUs to accelerate calculations. Examples include Altair AcuSolve™, Altair ultraFluidX™, and Altair nanoFluidX™ for CFD, aerodynamics, and particle-based fluid dynamics simulations. Altair EDEM also takes advantage of GPUs for bulk and granular material simulation. Depending on the application, GPU-based acceleration can reduce simulation time by a factor of 4x or more[5].

Altair offers multiple optimized GPU appliance node types with configurations that vary depending on the simulator's demands. GPU nodes are typically preconfigured with one or more NVIDIA® V100 or A100 Tensor Core GPUs connected via PCIe or optional NVLink interconnects.

As HPC administrators know, setting up GPU-capable clusters can be demanding. GPU clusters involve prerequisite software such as CUDA libraries and tools and the complex workload management policies to schedule parallel GPU solvers across multi-host, multi-GPU environments. With Altair Unlimited, engineers simply submit their workloads as usual and the appliance takes care of the rest. Built-in PBS Professional scheduling policies ensure that the host and GPU portions of CUDA-based workloads are placed optimally. The appliance automatically takes care of core affinity, bus topologies, Linux cgroups, and other complex details associated with GPU workloads.

Reduce Costs and Boost Performance with Efficient Resource Sharing

A common problem, particularly in small or medium-sized design environments, is resource sharing. Without workload management, access to shared server resources and licenses needs to be coordinated manually. Suppose multiple users attempt to run simulations on the same servers.

[4] The Xeon processor depicted in Figure 5 is a second generation Intel® Xeon® Gold Scalable Processor. Actual processor SKUs used may vary by appliance nodes

[5] Full Throttle: Altair Accelerates Engineering Simulations with NVIDIA GPUs https://blogs.nvidia.com/blog/2020/03/26/altair-gpu-support/

Simulations can hang or fail for reasons that include insufficient memory or scratch space, license availability, or attempts by multiple applications to access the same physical GPU device. This results in delays, failed jobs that need to be restarted, and poor resource utilization, as illustrated in the top portion of Figure 6.

With built-in PBS Works software and preconfigured PBS Professional scheduling policies, the Altair Unlimited appliance schedules jobs automatically, enabling users to achieve consistently higher job throughput and higher resource utilization at a lower total cost of ownership (TCO).

With efficient resource-sharing policies, jobs are queued and dispatched by a central scheduler to the node(s) best able to support the workload. This means that jobs complete sooner and hardware is used more efficiently.

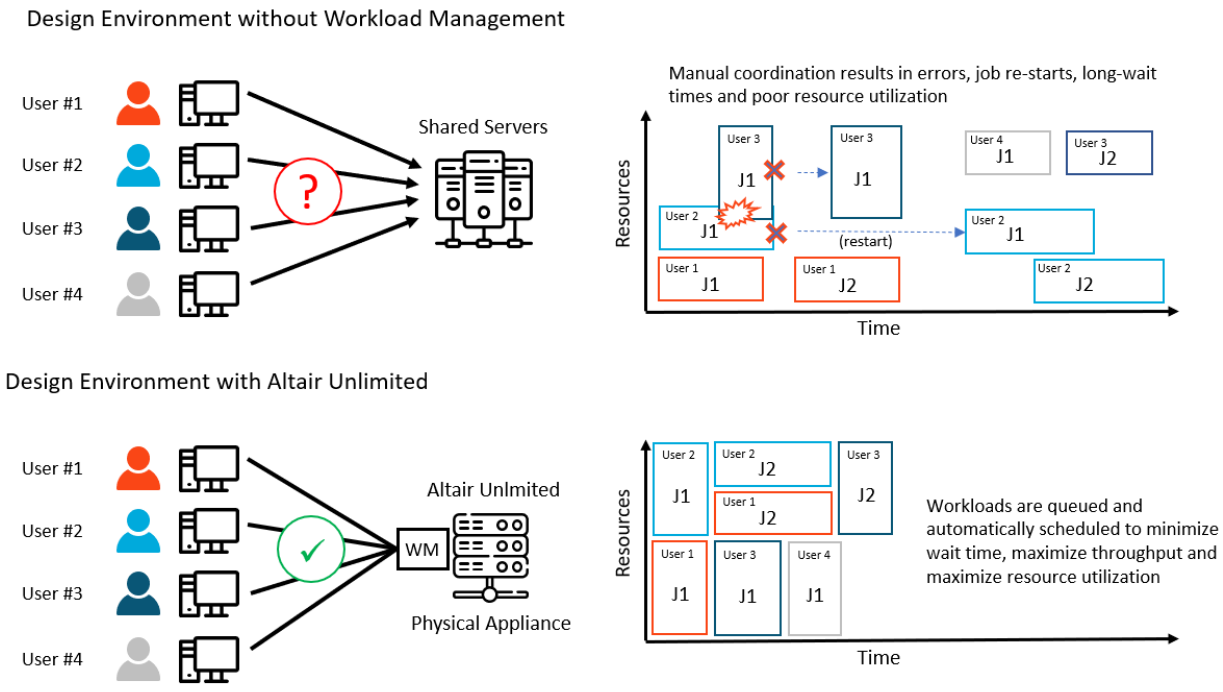

Figure 6 - Workload management dramatically improves throughput and resource utilization

TAILORED TO YOUR ENVIRONMENT

Altair Unlimited is available in a wide variety of standard configurations, from a single-node starter kit to a full rack with clustered servers optimized for specific workloads. Altair works with customers before installation to determine the proper mix of standard server building blocks. Appropriate power distribution units (PDUs) are installed by Altair depending on server configurations and on-premises power options. Datacenter environments with air cooling are preferred, but the appliance may also operate in standard office environments depending on the appliance configuration.

From a network perspective, the appliance requires just two physical RJ-45 CAT 6 network connections. Altair preconfigures three externally exposed IP addresses based on the customer's IP address space. Two addresses are used to access head node services so that the cluster can be used immediately. One IP address is reserved to optionally allow Altair access for remote system, cluster, and application-level support and administration.

A Turnkey Managed Service

A key advantage of the Altair appliance is that it includes system, cluster, and application-level support. Authorized appliance users can log various requests with Altair through a support portal. Requests can include system-level requests (updates to operating systems and infrastructure), workload management requests (configuring HPC middleware and scheduling policies), and application requests for assistance with HyperWorks and third-party solvers. In many respects, this is better than having a dedicated internal support team. Support requests are handled promptly by Altair personnel expert on the infrastructure, middleware, and applications included with the appliance.

Altair takes responsibility for routine tasks such as OS and application-level patching to address security vulnerabilities and ensure optimal functionality. Requests such as updating application versions, changing resource sharing policies, and updating queue configurations are all handled by Altair. With this level of service, engineers can focus on their work rather than becoming bogged down in details of HPC cluster administration.

Built with Security in Mind

Different customers have different security needs. Altair offers a spectrum of remote support options ranging from dedicated VPNs to support via web meetings to on-premises support.

Altair will work with the customer in most environments to open a single port in the customer's external firewall. Traffic will be routed from a pre-authorized Altair IP address range to a single IP address on the customer's subnet. A trusted contractor login provides authorized Altair personnel from designated locations with secure access to the appliance. Customers with elevated security concerns may leave ports closed by default and allow remote access only in the event of a service request.

Flexibility and Scalability built-in

The Altair Unlimited appliance provides a flexible, scalable alternative to in-house HPC clusters or cloud-based environments. The appliance provides scalability in multiple dimensions. Customers can easily add new Altair or third-party solvers, add additional users or workgroups, and scale up compute nodes as requirements evolve. For convenience, hardware can be financed over flexible 3 to 5-year terms with the option to expand or upgrade configurations mid-term or buy out the lease and capitalize hardware.

Scale to The Cloud with Ease

For Altair Unlimited users that need assistance scaling to public cloud services, Altair can help configure PBS Professional to tap public cloud resources during periods of peak demand.

Administrators retain full control over what users and applications are shifted to the cloud using PBS Professional features such as quotas, limits, and access control lists. PBS Professional supports all major cloud providers, including Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud.

Customers can also extend their on-premises Altair Unlimited appliance using the Altair Unlimited virtual appliance, providing similar functionality in the cloud.

COMPELLING BUSINESS VALUE

For design firms across all industries, the Altair Unlimited appliance offers compelling business value. It is a fully managed, pre-engineered, and pre-integrated solution designed and supported by experts.

With Altair Unlimited, customers can effectively take the risk out of decisions related to HPC infrastructure. Users can securely access HyperWorks and third-party applications from existing engineering desktops or remote locations. They will realize immediate productivity gains without the hassles of managing on-premises or cloud infrastructure and applications.

With the Altair Unlimited appliance, design and development teams can:

- Engineer better, higher quality products
- Dramatically reduce management and IT support costs
- Reduce infrastructure costs and improve operational efficiency
- Improve collaboration
- Enjoy top-quality support from a single vendor

CONCLUSION

Design engineers increasingly rely on HPC to design better, higher-quality products faster and more cost-effectively. For small and mid-sized engineering teams, installing and managing HPC clusters is a significant challenge. Organizations often lack sufficient IT support personnel with the necessary expertise.

For many CAE customers, the Altair Unlimited appliance provides an ideal solution. Altair Unlimited is a turnkey, state-of-the-art private appliance based on industry-leading Intel Xeon processor technology. The appliance includes the HyperWorks software suite with unlimited licenses, enabling simulations across multiple engineering disciplines. With a user-friendly web portal, design engineers can be productive immediately. They can focus on their work rather than be concerned with details about the underlying HPC cluster.

The Altair Unlimited appliance helps improve productivity and collaboration and dramatically improves simulation throughput. It also helps reduce overall TCO by ensuring that hardware and software assets are used efficiently.

Learn more about the Altair Unlimited appliance by visiting https://www.altair.com/altair-unlimited/

