MEETING THE CHALLENGES OF COMPLEXITY AND SCALE FOR MANUFACTURING WORKFLOWS

Michael Feldman

White paper

November 2014

MARKET DYNAMICS

Modern manufacturing increasingly relies on advanced computing technologies for digital modeling and simulation, applications used by engineers for product design, analysis and testing. By creating digital models, manufacturers are able to subject their designs to manufacturability analysis and performance testing, without the need for physical mockups. Benefits include improved product quality, shorter time to market, and reduced manufacturing costs.

The manufacturing segment is comprised of both large-product and consumer product manufacturing, which together encompass household products, consumer devices, automobiles, ships, aircraft, military equipment, and construction equipment. The largest manufacturers have used high performance computing (HPC) technologies for decades, employing computer-aided engineering (CAE) and computational fluid dynamics (CFD) to conceive and test new products to as well as maintain older ones. Meanwhile, consumer product manufacturing and component manufacturers in the large product tool chain, while underrepresented, are increasing employing HPC at small and medium-sized companies.

Today large product manufacturing is the biggest segment in the commercial HPC sector, while consumer product manufacturing is among the five fastest growing verticals[1]. Growth is being driven by the increasing affordability of powerful HPC systems, the rising number of engineers familiar with advanced computational methods enabled by the latest CAE and CFD applications, and better ease-of-use for hardware, applications, and system software.

[1] High Performance Forecast for 2011 through 2015: Economic Sectors and Vertical Markets, Intersect360, Research, June 2011

It is also being driven by competition. In a study of manufacturers, 72 percent reported that they believed increased adoption of advanced computational methods would lead to competitive advantage for their companies[2].

The Workflow Challenge

Digital manufacturing entails one of the more complex workflows in high performance computing, which is a result of the interplay between applications involved in bringing a product from conception to market. Advanced computation can be brought to bear at each stage of product development:

1. Design (conceptualization, early prototyping, visual renderings)

2. Engineering (prototyping, component analysis, testing and certification)

3. Manufacturing (component integration/assembly, manufacturability testing and analysis)

During development, multiple iterations of each stage are common, as issues uncovered by engineering or manufacturing simulations necessitate revisiting designs. For example, for the wing of an airplane, vibration analysis during engineering simulations may uncover a design problem. Likewise, stress testing of the entire airplane during a virtual wind tunnel simulation may precipitate redesign for one or more components. This circular development cycle is streamlined by doing much of the work with digital models, but workload management can bog down product development if it's not managed efficiently on whatever compute resources are available.

Engineers working on different stages of product development may be simultaneously running computer applications at each stage. However, dependencies between jobs must be respected to provide a coherent workflow. At the same time shifting job priorities tied to unforgiving product schedules can complicate workload staging. As a result, for the manufacturing sector, workload resource managers must be robust enough to handle intricate workflows and nimble enough to manage dynamically changing job loads. As a consequence, workload managers are key to getting the most value out of the underlying HPC infrastructure.

Another challenging aspect of manufacturing workflows is managing the scale of the product simulations. Large firms, in particular, such as automobile or aircraft manufacturers, use

[2] Modeling and Simulation among U.S. Manufacturers: The Case for Digital Manufacturing, Intersect360, Research, September 2010

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

systems with thousands of cores to queue up hundreds of thousands of jobs at a time. Each one, which may be modeling a particular feature of a particular product component, can take less than a minute to run. But with hundreds of thousands in flight at any one time, workflows can become an enormous logistical challenge for the job scheduler and a nightmare for the system administrator.

OPPORTUNITY FOR ADAPTIVE COMPUTING

Adaptive Computing, with its focus of intelligent workload management, has offered advanced resource management tools to HPC users for more than 10 years. According to our latest HPC Site Census Survey, the company is the top supplier of job management solutions, with 44 percent of sites citing Adaptive as a supplier[3]. Likewise, Moab was the most oft-mentioned middleware package by those same respondents.

The company's Moab-branded products, which can operate across traditional clusters, grids, private clouds, and public clouds, have been designed to provide automated job scheduling and resource management for many types of complex workflows, such as those found in manufacturing. They are geared for complex workloads in both technical and business environments, where application performance, throughput, scalability, and dynamic job prioritization, are paramount to the organization's mission.

Moab Brings Big Workflow to Workload Management

Adaptive has continually refined its product set, most recently to incorporate the notion of "Big Workflow," a term coined by the company that provides a model for integrating data-intensive and compute-intensive workflows under a single workload management scheme[4]. Essentially, the model draws together HPC simulations with big data analytics workloads so that customers can manage them in a unified manner.

In the latest rendition of the Moab HPC Suite (Moab 8.1) and the Moab Cloud Suite, Big Workflow services deliver dynamic scheduling, provisioning and management of multiple applications across HPC, traditional cloud and big data environments. By doing so, Moab is able to automate much of the workflow management, which otherwise would need to be performed manually with scripts and administrator intervention. An important component of this is data staging, which ensures that any application data is located advantageously prior

[3] HPC User Site Census: Middleware, Intersect360 Research, April 2014
[4] Big Workflow: More than Just Intelligent Workload Management for Big Data, Intersect360 Research, February 2014

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

to the start of the job – the idea being to maximize system utilization and minimize resource contention.

For product development, this streamlines the complexities of job management and increases application throughput, which, in turn, reduces time to market. As a consequence, manufacturers employing Moab are more easily able to meet product market schedules, while maintaining or even improving quality.

A Tool for Accelerated Throughput

Adaptive also has devoted a good deal of effort to speeding application throughput, especially for those workload environments, as in manufacturing, where hundreds of thousands of small jobs can be running simultaneously. In these situations, scheduling overhead can become so burdensome as to make high throughput workloads impractical. To address this challenge, Adaptive developed Nitro, a tool that can reduce scheduling overhead significantly. It is able to do this with a very high-speed message queue and aggregating numerous small jobs into a single batch unit. This optimization method enables Moab to launch tens to hundreds of jobs per second on a node, increasing throughput up to 100 times.

Accounting for Resources

In those same large-scale manufacturing workloads, resource accounting is often used to track usage of all the engineering teams involved in product development. In some cases, the goal is to turn in-house IT into a profit center, by charging engineers for the resources they use. Users can even determine the importance of their workload by selecting high priority queues that cost more while saving funds on the low priority workload. In other cases, it is used to keep track of CPU, memory, or even power usage, in order to help guide resource management and planning. For both scenarios, Adaptive's Moab Accounting Manager (MAM) provides the necessary features to track this usage across all types of infrastructure – in-house clusters, remote clusters, or full-fledged clouds. MAM was derived from the Gold Allocation Manager, an open source resource management tool developed by the US Department of Energy, but has been enhanced significantly to provide a more full-featured tool for commercial usage. It supports the traditional accounting notions of deposits, withdrawals, transfers and refunds, but also provides customizable tracking based on the particular user demands.

The Moab toolset also includes other features critical to many manufacturing workloads environments. These include:

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

- A cloud bursting capability to offload peak workloads to resources managed as a utility
- An administrator portal to simplify job management
- An end user portal to simplify job submission (available in early 2015)
- Job templates to manage job with similar characteristics
- An advanced Service-Level Agreement (SLA) facility that enforces service delivery under widely varying conditions

INTERSECT360 RESEARCH ANALYSIS

Manufacturing represents a very large and growing domain for HPC applications. However, there is still a sizeable untapped market in this area, especially for small and medium-sized companies. Regardless of the size of the organization, the inherent challenges of procuring, deploying, and operating HPC infrastructure in commercial data centers makes managing those workflows optimally critical for manufacturing enterprises.

To serve those needs, manufacturers looks to solutions that can deliver the highest application throughput on its HPC infrastructure in order to maximize their return on investment. The complexities of manufacturing workflows, the extremely large number of jobs that often constitute them, and the diverse nature of HPC infrastructure that underlie these applications, represent additional challenges in this area.

Today only a limited of vendors offer the kind of advanced workload manager that can meet the demands of these complex workflows. Adaptive Computing is one such vendor, and, with its Moab HPC Suite and Moab Cloud Suite, offers a unique product set for this industry. In particular, the Big Workflow model used by the Moab software is able to integrate the simulations and analytics into a unified workflow that can deliver optimal utilization of the underlying compute and storage hardware. Further, it provides the flexibility to automatically scale out these workloads across clouds, grids, or other off-premise infrastructure.

The manufacturing industry's critical need for complex workflow management plays well into Adaptive Computing's product strength. The company has a 13-year track record of delivering intelligent workload management for HPC customers with widely varying application profiles and hardware platforms, and for commercial organizations that must rely on proven technology for critical business needs, Adaptive has demonstrated itself to be a reliable partner.

Intersect360 Research
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com