

Big Workflow: More than Just Intelligent Workload Management for Big Data

Michael Feldman

White Paper
February 2014

EXECUTIVE SUMMARY

Big data applications represent a fast-growing category of high-value applications that are increasingly employed by business and technical computing users. However, they have exposed an inconvenient dichotomy in the way resources are utilized in data centers.

Conventional enterprise and web-based applications can be executed efficiently in virtualized server environments, where resource management and scheduling is generally confined to a single server. By contrast, data-intensive analytics and technical simulations demand large aggregated resources, necessitating intelligent scheduling and resource management that spans a computer cluster, cloud, or entire data center. Although these tools exist in isolation, they are not available in a general-purpose framework that allows them to interoperate easily and automatically within existing IT infrastructure.

A new approach, known as “Big Workflow,” is being created by Adaptive Computing to address the needs of these applications. It is designed to unify public clouds, private clouds, Map Reduce-type clusters, and technical computing clusters. Specifically Big Workflow will:

- Schedule, optimize and enforce policies across the data center
- Enable data-aware workflow coordination across storage and compute silos
- Integrate with external workflow automation tools

Such a solution will provide a much-needed toolset for managing big data applications, shortening timelines, simplifying operations, and maximizing resource utilization, and preserving existing investments.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
TABLE OF CONTENTS	2
THE BIG DATA CHALLENGE	3
THE BIG DATA GOLDRUSH.....	3
ENTER BIG WORKFLOW: Processing BIG DATA workloads	4
CONCLUSION	6

THE BIG DATA CHALLENGE

Until recently, the majority of computing in the data center consisted of workloads that did not demand a very sophisticated resource allocation strategy. Their signature characteristic is that they were assumed to run forever – applications such as email serving, web hosting, and CRM systems with no definitive end to their lifetime. As a result of simple and permanent allocation, many of these applications could share the same server, without the need for scheduling or resource allocation beyond that of a single server node. The overriding goal for these applications was to maximize uptime.

Intense simulations and big data analysis, however, demand a different approach. They are resource demanding and thus each application often requires many servers – in some cases, thousands of them. Since their goal is to deliver a particular answer or insight in a specific timeframe, they do not run continuously. For these reasons, scheduling and resource allocation are critical, and to run these applications optimally across a compute cluster, data center, or cloud requires a level of sophistication that is difficult to automate.

Software tools and frameworks have been developed for compute- and data-intensive computing, but they tend to be deployed in an ad hoc manner. Workload schedulers, resource managers, workflow engines, and system monitors provide different aspects of the solution, but it is often left to the user or system integrator to combine them into a coherent environment. And for the reasons just stated, traditional IT has little to offer.

THE BIG DATA GOLDRUSH

Data collection and analysis has been a foundation of information technology since its inception. But with the rise of big data the application set is increasing in number and broadening in scope. More complex data sets – both structured and unstructured – are being amassed and analyzed to deliver critical intelligence to a growing set of users.

This expansion is happening across many domains – from financial services, scientific research, energy exploration, healthcare, manufacturing, security, and media/entertainment. In each field, data scientists and engineers are building applications to extract value from their growing datasets. Businesses want to better monetize this in a variety of ways: targeting the needs of new and existing customers, detecting fraud or security breaches, optimizing logistics and accounting, and streamlining product development, to name a few.

The enterprise, private sector and academic institutions use an analogous set of applications to perform data analysis and visualization for technical computing applications, principally digital simulations of natural phenomenon and engineering models. As a result,

organizations are investing in big data and HPC. Our research shows that most in the technical and business computing segments will spend more than 10 percent of their budget on big data hardware, software, and services¹. That is based on a survey of more than 300 organizations across the computing spectrum. Underpinning this is the expansion of HPC, again in both the business and technical segments. According to our analysis encapsulated in our latest market models², HPC revenue will grow 6.5 percent CAGR from 2012 to 2017, reaching \$39.9 billion at the end of the forecast period.

The propellant for all these applications are technologies such as larger and faster storage systems, more sophisticated analytics tools, and the availability of utility datacenters and clouds. Raw compute and storage are relatively cheap commodities; the magic that delivers it to the application is the system software and tool stack.

Collecting the data in vast storage arrays is the easy part; extracting useful results from it remains the more difficult challenge. In particular, extracting value from large-scale infrastructure (datacenters and clouds) is often the most daunting task, given the application challenges outlined in this paper. But as datasets grow and the analytics software becomes more complex, the ability to tap large-scale computing becomes more crucial.

The use of all these technologies and, especially, the orchestration between them is crucial to the success of these endeavors. Ensuring the pieces work together and deliver the needed performance is a challenge for workloads that meld compute- and data-intensive applications. In a recent study compiled by Intersect360 Research, it was found that even enterprise users recognized that greater performance and capacity were required when they considered their big data needs. That represents an important departure from the slower, more deliberate approach businesses traditionally take when defining their computing requirements.

Integrating the component parts – analytics software, high performance hardware, and scaled-out infrastructure – is especially challenging since there is such a wide array of hardware and software platforms to contend with. But if big data is to move beyond custom workflows, enabling technology must be developed to automate these processes.

ENTER BIG WORKFLOW: PROCESSING BIG DATA WORKLOADS

To integrate these disparate pieces, one model being considered is “Big Workflow,” a term coined by Adaptive Computing to encapsulate the workload needs of data-intensive and

¹ Intersect360 Research, Special Study: “The Big Data Opportunity, Intersect360 Research,” 2013

² Worldwide High Performance Computing 2012 Total Market Model and 2013-17 Forecast, Intersect360 Research, May 2013

compute-intensive applications in enterprise, government and academic environments. The goal is to streamline workflows that incorporate big data applications by automating the management of the various jobs. In doing so, execution timelines are compressed, SLAs are enforced, and resource utilization is maximized. The general idea is to provide a way for big data, HPC, and cloud environments to interoperate, and do so dynamically based on what applications are running.

Unlike generic workload automation, which generally just manages jobs and tracks resource usage, the Big Workflow model adds the context of data and resource awareness, enabling intelligent coordination of jobs based on data dependencies, as well as compute requirements. Input data is presented to the initial application, which processes it and passes its output data to the next job, and so on. Ultimately the user ends up with an actionable result, rather than just a completed job list. This virtual assembly line is designed to provide a high-level framework for data workflows.

The automation inherent in this model optimizes resource usage in a way that can't be achieved with workflows that are managed manually. With the knowledge of available resources and application dependencies encapsulated in the workload manager, compute-intensive simulations and data-intensive analytics can be scheduled intelligently across a cluster, datacenter, or cloud. The Big Workflow model also incorporates public clouds, such as that offered by Amazon Web Services (AWS).

An example of using AWS in this manner was recently documented by Cycle Computing³. In this case, a 156 thousand-core cluster was carved out of the Amazon cloud infrastructure for running a large compute- and data-intensive application: a materials science analysis to identify the best substrates for solar panels. The aggregated cluster represented over one petaflop of peak capacity and was able to execute 264 compute-years in a mere 18 hours. A custom-built workload scheduler, called Jupiter, was constructed to distribute tasks and data across the AWS cloud for this application.

Adaptive Computing's Big Workflow solution will leverage its long experience in HPC, big data, and cloud, incorporating Moab's scheduling and optimization capabilities to manage workloads across heterogeneous environments. Adaptive's latest integration endeavor on the OpenStack platform -- an open source operating system designed for private and public clouds and will likely span other environments such as technical computing -- provides a general-purpose platform that can process standard applications. In addition, Adaptive is developing a Moab/TORQUE integration with Intel Hadoop to manage big data. Moab runs on a multitude of platforms, to manage workflows that encompass simulations and data analytics for both virtualize machine (VM) and bare metal environments.

³ <http://arstechnica.com/information-technology/2013/11/18-hours-33k-and-156314-cores-amazon-cloud-hpc-hits-a-petaflop/>

Adaptive has defined a set of components that make up its Big Workflow platform. At its heart is a policy-based workload scheduler, based on the existing Moab technology, which is “data aware” and able to manage a pool of computing resources by application needs. A data expert component knows how the data is used and can coordinate its migration across storage silos so that it is available to the appropriate compute platform. Finally, Moab will integrate with popular business workflow engines to provide application task management on a system-wide basis.

Public clouds such as AWS and the HP Public Cloud will be supported, as well as private clouds and clusters -- both technical computing and Hadoop-style clusters. User and administration portals will be provided on top of Moab’s web services API to direct the workflow at the level of the workflow scheduler. The integrated platform will provide automated workflow management that can be applied to business and technical computing applications.

CONCLUSION

Simulation and big data analysis are creating a new set of workflow requirements for the data center. Like it or not, the data center and its users will have to deal with them. As users look for ways to integrate these workloads into traditional computing workflows, the Big Workflow concept offers a useful model for doing so within a general-purpose framework, and at scale by unifying all data center resources.

Given Moab’s long legacy in workload management and Adaptive Computing’s expertise in technical computing and cloud environments, the company is in a unique position to capitalize on this fusion of big data, HPC and cloud computing. Because of the extreme demands inherent in compute- and data-intensive environments, Moab was designed to extract the maximum amount of processor cycles from the underlying hardware. In addition, Moab’s various incarnations for HPC systems, enterprise clusters, and clouds provide a solid foundation for application workflows that operate in all of those environments.

With Big Workflow, Adaptive Computing is offering not just a vision, but a software suite that will allow users to meet the unique workflow requirements of big data applications. It does so in an evolutionary fashion that preserves investments in hardware, software, and staff skill. In a nutshell, Big Workflow is designed to compress timelines, simplify operation, and maximize existing resources, all of which accelerate insights that inspire data-driven decision as well as offer the necessary cost incentives relevant to this application set.

