



Workload Automation: Helping Cloud Computing Take Flight

By Mike Moser, product management director and program executive for Mainframe Service Management, and Tony Morelli, global product manager, BMC Software

TABLE OF CONTENTS

- EXECUTIVE SUMMARY 1
- WHY BUSINESSES ARE MOVING TO THE CLOUD 2
- AN OPPORTUNITY AND A CHALLENGE 3
- EFFECTIVE WORKLOAD AUTOMATION FOR THE CLOUD 4
 - » A Business-Level View of Workloads 4
 - » Dynamic Provisioning and Control. 5
 - » Security 5
 - » Auditing 5
- A PRODUCTIVE WAY TO TAKE ADVANTAGE OF THE CLOUD. 6

EXECUTIVE SUMMARY

Businesses are moving workloads out of traditional data centers and into the cloud as they seek to become more agile and cost-effective. Cloud computing provides an alternative model for IT infrastructure and application sourcing and delivery. However, IT organizations that leverage cloud computing may also be overlooking another big opportunity. That's because many data centers are using the cloud to host applications and data storage but typically aren't using the cloud to process *batch computing workloads*. These workloads represent one of the more mission-critical and complex tasks within the data center. For the most part, batch workloads — 70 percent of all workloads — remain hard-wired to the data center.

The key to unlocking this opportunity already exists in the data center in the form of enterprise workload automation tools. This is similar to what an air traffic controller must do when directing multiple flights from all over the world that all have different departure times, landing times, and resource requirements. Effective enterprise workload automation tools must find the best available resources wherever they reside — locally or in the cloud. These solutions then orchestrate work in a sequence that gets the job done efficiently and on time.

With the right workload automation solutions — ones that can dynamically exploit virtual resources based on business priorities and policies — businesses can take much greater advantage of cloud computing. They can achieve faster time to market, reduced costs, and more flexible business operations. How? By sending the workload into the cloud where it can benefit from access to required resources and faster processing while also reducing capital equipment expenditures inside the data center.

WHY BUSINESSES ARE MOVING TO THE CLOUD

Cloud computing is a powerful computing model that enables data centers to dramatically improve their efficiency and service levels. For example, cloud computing eliminates over-provisioning of infrastructure to get the IT capacity the business needs. Instead, the IT department can “rent” the resources it needs when (and only when) they are needed. This can reduce capital equipment expenses and can potentially also reduce operating expenses, if automated management systems are properly applied.

Workload management can be applied to different cloud deployment models — public, private, or hybrid. Each cloud model has different purposes and benefits. *Public* clouds are typically owned by a service provider and shared by many companies. *Private* clouds are typically owned and operated by a single company inside the firewall or sometimes operated outside the firewall by a service provider. Amazon EC2, Force.com, and Google are examples of popular public clouds. Companies may also operate a hybrid cloud, which is a collection of services that run on a combination of public and private clouds.

IDC predicts that revenue from IT cloud services will grow from \$17.4 billion in 2009 to \$44.2 billion in 2013; this is a five-year annual growth rate of 26 percent, which is more than six times the rate of traditional IT offerings.¹ (These figures do not include spending for private cloud deployments. They only include information for public IT cloud services offerings.)

According to industry analyst firm Forrester Research, “Many enterprise infrastructure and operations professionals are taking this [cloud computing] concept in-house and building their own internal clouds ... Cloud computing is compelling, but bypassing IT Ops can be dangerous.” (“Deliver Cloud Benefits Inside Your Walls,” Forrester Research, Inc., April 2009)²

Cloud computing can bring data centers much closer to realizing their vision of the dynamic data center: IT infrastructure that expands (and contracts) with elastic efficiency to meet business needs. However, most data centers are overlooking a big opportunity related to running batch workloads. Today, data centers are using cloud computing mostly to get hardware efficiencies and energy savings by distributing their computing workloads over virtualized versions of their own computers.

If they use public clouds, it is usually for hosting online applications and data or for running online workloads. Data centers *do not* generally use public clouds to run batch workloads, which account for 70 percent of all workloads. Batch workloads are fixtures in every company, from fast-growing Web 2.0 companies to *Fortune 2000* multinationals. Batch workloads handle the back-end processing for many critical business applications, such as online trading, ATMs, hotel and airline reservations, and revenue-generating Web applications. They are mission-critical; if batch jobs don’t complete on time, the business can slow down or come to a halt. And they are resource-intensive, making them great candidates for the cloud.

Traditional overnight batch windows are shrinking or going away as business runs 24x7 on the Internet, making job scheduling difficult even under the best of circumstances. Data centers are also handling more batch workloads, requiring a constant balancing act to match capacity with workload demand.

If you aren’t using the cloud for processing batch workloads, you are missing a big opportunity for reducing capital and operational expenditures because batch workloads represent such a significant portion of all workloads. Cloud computing provides the flexibility to achieve greater hardware efficiencies and energy savings, as well as improve the delivery of business services. Resources for batch workloads can be available when you need them, which can help IT organizations meet their service levels.

Beyond the Buzzword: What Is Cloud Computing?

Cloud computing is the Internet-based, just-in-time delivery of data, applications, storage, and computing power as services, done in a way that completely shields consumers from the underlying technical details. It reduces or eliminates the need for businesses to invest in IT infrastructure. A good analogy is a hotel chain, which temporarily rents rooms and other services to guests, eliminating the need for them to own homes in every city they visit. Businesses, nonprofits, and government agencies may be both hotels (providers of services) and guests (consumers of services).

AN OPPORTUNITY AND A CHALLENGE

Imagine the efficiencies and service-delivery improvements possible if you could dynamically send eligible batch-related workloads into the cloud. Consider this typical data center scenario. You're already running 50,000 batch jobs per day. You have a new business unit coming online, and you don't have the in-house resources to handle it. Instead of scrambling to reprioritize or reallocate in-house resources — and possibly risk missing service level deadlines that can cost your business money *and* customers — what if you could automatically obtain the resources you need from a virtual pool that includes the cloud? What if you could dynamically send batch jobs up to the cloud and run them easily, efficiently, and securely — just as when you host your applications and manage your data on the cloud?

Workload automation could help you achieve this objective. The *challenge* is providing a strong and secure management infrastructure that leverages all available resources, including the ones you don't own and can't easily see.

Workload automation can help, and it's already a mainstay of the data center. Similar to an air traffic controller, effective workload automation tools find the best available resources for your needs. They orchestrate work in a sequence that gets the tasks done efficiently and on time. But effective workload automation *for the cloud* creates some new requirements. Data center staff must now manage resources that they don't own and can't easily see, working under the usual time pressures.

Traditional workload automation tools and older job scheduling solutions aren't flexible, intelligent, or insightful enough for cloud environments. These approaches can't help "land the plane" safely. There are two obstacles: resource definition and resource availability.

Resource Definition

- » How can you define, schedule, run, and track workloads on resources you don't own?
- » How can you maintain control without installing workload automation agents everywhere?

Resource Availability

- » How do you know what's *really* available to you in the cloud?
- » How do you know a resource will support your workload?

Resource definition: Data centers need the ability to run workload dynamically without the requirement to pre-provision eligible servers with workload automation tool components such as agents. The solution should let them define the key resources and then send workloads off to a certain environment based on business need and resource availability. In the cloud, the resources are constantly changing, so having the flexibility to do this without the requirement to maintain a workload automation software footprint is crucial. If not, your team could be spending too much time managing the environment and identifying changing resources, and the cost would be prohibitive in a private cloud. That obstacle, however, can be overcome with a solution that supports agentless scheduling, allowing you to send workload anywhere within the traditional data center or up into the cloud. You can gain visibility into the underlying hardware while still achieving a guaranteed level of compute power.

Resource availability: Without the ability to define and track resources you don't own, how do you *really* know what's available or what's being consumed? Or whether a resource can support a workload? Guessing without enough information could result in a workload that either doesn't complete on time or doesn't complete at all.

With the cloud, you don't know what the resource is or where it is located, but when you need that resource, you need it right away. Unfortunately, traditional approaches to scheduling batch workloads are based on static parameters: *date and time* or *alignment of workload to specific physical locations and hardware devices*. But many of the reasons you need the cloud will be based on events. For example, imagine you've hit a resource threshold with your internal servers, and you need to reach out to the cloud. You need to ensure resources are available, use the resources, and de-allocate them when the job completes. How can you manage resource availability effectively across your workloads, without dedicating staff to "baby-sit" every job?

Any workload automation solution must overcome these two obstacles in order to take full advantage of the cloud.

EFFECTIVE WORKLOAD AUTOMATION FOR THE CLOUD

To be confident enough to entrust batch workloads to the cloud, data centers need a single enterprise tool for workload automation — one that can manage a vast array of workloads across owned and non-owned environments using business rules and policies. An ideal workload automation solution would provide the following:

- » A business-level view of workloads
- » Dynamic provisioning and control
- » Security
- » Auditing

(See Figure 1.)

A BUSINESS-LEVEL VIEW OF WORKLOADS

The big advantage of the cloud is flexibility. But flexibility always comes at a cost. To get the most out of the cloud, data centers will want to empower multiple departments, clients, and business users to use cloud resources — and ideally let *them* determine when, where, and how to use the resources.

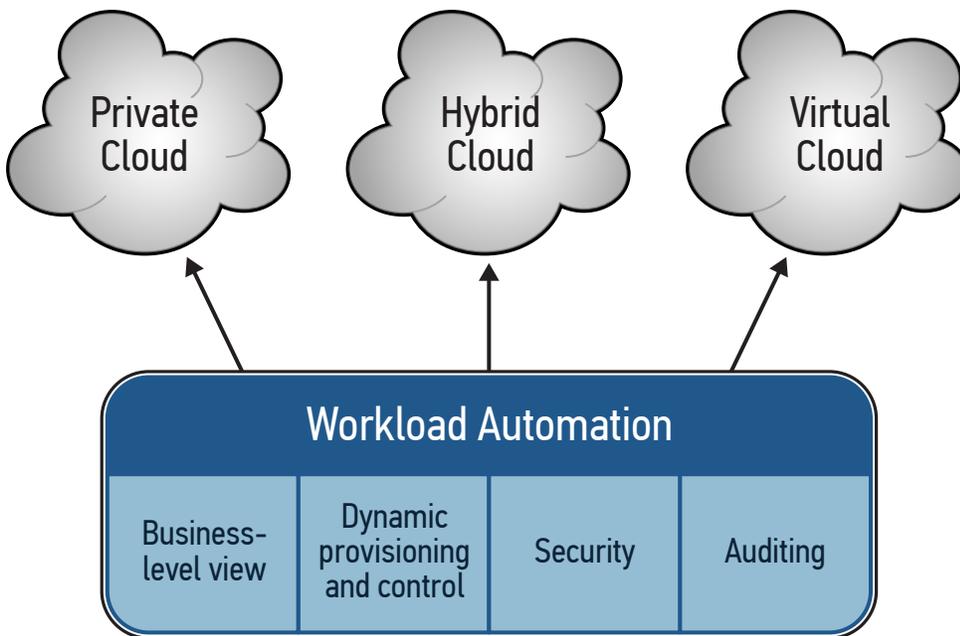


Figure 1: Ideal Workload Automation for the Cloud

A business-level view enables you to get the most flexibility at the right cost. It enables IT staff to orchestrate resource usage based on business policies, service-level agreements, and resource pricing, preventing misuse that can cause cost overruns or service problems. It also lets IT staff view workloads based on application or service status to ensure overall business service levels. It provides granular control over all virtualized resources: internal and cloud.

Furthermore, it allows your IT staff to manage workloads from different departments or business users differently, based on business requirements. Some workloads may be eligible for the cloud, but others will not be. You can ensure that one department or business user doesn't dominate others or receive a more expensive level of service than its agreement specifies. In addition, you can prevent your business users from sending *everything* to the cloud, thereby avoiding extra costs for outside services. Another advantage is that you can enable multiple departments or business users to run different workloads simultaneously without bumping into each other or even being aware that others are sharing resources with them.

DYNAMIC PROVISIONING AND CONTROL

To tap into dynamic resource pools, you need to be able to dynamically provision workloads — that is, to know when resources become available and to deploy workloads to those resources with as little “baggage” (dependencies) as possible. This means you can’t rely on software agents (cloud providers are unlikely to let customers install software on provider systems), and the workloads must be portable and be executed independently.

Two technologies — agentless scheduling and job-level virtualization — provide dynamic provisioning and control. Agentless scheduling eliminates the need to install software agents on target systems in the cloud. You can move workloads around based on resource availability at any point in time — whether the systems reside inside your firewall, outside the firewall in a private cloud, or outside the firewall in a public cloud. Job-level virtualization bundles complete execution instructions with every workload, so the workloads start, run, and complete automatically with full reporting and an audit trail.

Most use of virtualized resources will be based on thresholds and exceptions. Dynamic control enables you to schedule batch workloads on-the-fly based on events that create exceptions or exceed thresholds. For example, you can use business rules to specify that a batch workload go to a cloud if an unexpected spike in transactions maxes out the capacity of your internal servers. You control everything using business rules and policies.

Business Service Management (BSM) — a comprehensive and unified platform for running IT — plays a significant role in cloud computing. BSM solutions provide the automation that enables the high level of dynamism so essential to cloud computing, while still maintaining the control necessary to ensure high-quality service delivery, strong security, and strict compliance. In fact, BSM automates such processes as server provisioning and repurposing, event management, dynamic reconfiguration of IT infrastructure elements, and user self-service.

SECURITY

The workloads are secure no matter what they are or where they run. User-level policy management prevents unauthorized access or activities on workloads. When workloads are sent to the cloud, their contents are automatically encrypted and secured against outsiders. You can make the cloud an extension of your data center without worrying about security.

AUDITING

Built-in auditing tells you every step of the way what was done, how it was done, who did it, and what changes were made. You know where the workload ran, if it completed, how long it took, and, if it didn’t complete, why. This provides an audit trail for use in both operations management and governance.

BMC Control-M offers many of these capabilities today. This cross-platform workload automation solution provides a highly automated, consistent, and effective way to manage workloads across the enterprise.

When choosing a workload automation tool for the cloud, look for a solution that can manage workloads across virtualized environments that include UNIX, Windows, Linux, VMS, and z/OS. The workload automation solution should provide security, auditing, multi-tenancy and other technologies for managing workloads in virtualized environments. It should provide advanced workload automation to cross-platform environments, including those that span the enterprise and extended enterprises (supply chains, partners, and customers).

Many workloads today combine batch and distributed processing requirements. It is important to be able to deploy, view, and manage hybrid workloads on any combination of virtualized resources: internal mainframes and servers, external systems run by your service providers, and public clouds.

Workload Automation and the Cloud: A Web-Based Business Improves Its Service Levels

A fast-growing provider of consumer information services has used workload automation to improve service levels and reduce the costs of managing its Web application. The application runs on an internal Web server farm. It incorporates a database with more than six billion records; the database is continually updated with information from more than 32,000 sources and frequently queried by partners using batch processes.

Using BMC Control-M, the company was able to reduce manual processes by 80 percent, enabling it to manage job scheduling and workload automation with a staff of 2 instead of 20 to 30, while meeting service levels for its 24x7 business. The company now plans to exploit cloud resources for managing its batch operations, including billing and reporting. As a result, the company will be able to dispatch batch-workload jobs to the cloud when internal resource thresholds are reached, enabling the company to handle growing workloads without expanding its IT infrastructure or IT staff.

Agent-less scheduling raises the level of automation and intelligence of workload automation, job-level virtualization, and business policy management. Like a good air traffic controller, modern workload automation can make intelligent, rapid decisions for data centers. It can think, "We are going to miss a service level and don't have the internal resources, so let's send this batch workload to this cloud provider today." Or, "Let's look at balancing our thresholds and dispatch some work to other virtualized environments because those resources are cheaper at this time."

A PRODUCTIVE WAY TO TAKE ADVANTAGE OF THE CLOUD

Effective workload automation can help you fully exploit cloud computing, no matter how big your business is or what cloud model you are using. Workload automation can help fast-growing businesses scale quickly without adding IT infrastructure or staff. It can help multinationals cope with IT cost-cutting mandates, absorb acquisitions faster by renting capacity, and accommodate more workloads without a commensurate build-out of the infrastructure.

Using modern workload automation approaches, you can move batch workloads to the cloud with assurance. You can achieve faster time to market, reduced costs, and greater flexibility — all the objectives you would be able to achieve if you were using an internal virtualized environment. You will maximize your hardware efficiencies and energy savings, improve business service levels, and make the data center much more dynamic.

For more information about BMC Control-M, visit www.bmc.com/solutions/msm-main/Mainframe-Job-Scheduling-Workload-Automation.html

ENDNOTES

1. IDC, "IDC's New IT Cloud Services Forecast: 2009-2013," <http://blogs.idc.com/ie/?p=543>, October 2009. Spending is across five IT categories: applications, servers, storage, infrastructure software, and application development/deployment.
2. BMC news release, www.bmc.com/news/press-releases/2009/BMC-Software-Brings-the-Power-of-Business-Service-Management-to-Enterprise-Cloud-Computing.html, July 15, 2009.

Business runs on IT. IT runs on BMC Software.

Business thrives when IT runs smarter, faster, and stronger. That's why the most demanding IT organizations in the world rely on BMC Software across both distributed and mainframe environments. Recognized as the leader in Business Service Management, BMC provides a comprehensive and unified platform that helps IT organizations cut cost, reduce risk, and drive business profit. For the four fiscal quarters ended December 31, 2009, BMC revenue was approximately \$1.90 billion.

About the Authors

Mike Moser, a product management director and program executive within BMC Software's Mainframe Service Management business unit, focuses on issues related to reducing costs while improving IT efficiency and service delivery. Moser's more than 20-year career has spanned a wide variety of engineering, IT management, technology consulting, and product management positions across both technology vendor and IT end-user organizations.

Tony Morelli, a global product manager for the BMC Control-M brand management team at BMC Software, supports product management activities and is focused on helping organizations realize business value by migrating to BMC Control-M. He also assists with BMC Software's Control-M field enablement and marketing for all of the geographic regions. Morelli has been with BMC Software since 1994 and brings with him more than 20 years of technical and management experience.

