



David Jackson
David Handy
Kershaw Mehta

Maximize IBM System z ROI with Open Cloud Management and Policy-based Optimization

IBM® System z® servers are valued for their resiliency, and companies and government agencies worldwide place their trust in them with good reason. However, as enterprises establish heterogeneous private cloud environments, a need has emerged for a middleware solution that is inclusive of System z environments. This solution must unify workload management and remove barriers to efficient use of virtualized resources.

Moab® Cloud Suite is a multi-dimensional policy engine and cloud management solution that can enable rapid implementation of heterogeneous private cloud environments. A capable, lightweight solution, Moab Cloud Suite automates the creation of standardized or customized virtual services in the System z environment. It can help users and groups in the organization tap into the high reliability and cost efficiency of running services in a System z environment. In addition, it is based on an open architecture that can take advantage of xCAT and other key components.

The Moab intelligence engine recognizes workloads, the service level agreements (SLAs) that are associated with them, and the resources that are required to optimally run them. In addition, it provides a future vision for unified cloud management, regardless of whether workloads are created and optimized on IBM System x®, Power Systems™, or System z. The Moab Cloud Suite user interface enables administrators to see, control, migrate, and terminate those workloads. Moreover, IT administrators with no System z experience, and even authorized users with varied or no technical training, can see and control workloads and can assess the health of workload environments, thus freeing the System z administrator for more strategic tasks. In addition, Moab Cloud Suite can track and report ongoing resource usage and costs to users and groups so that proper decommissioning can take place and underutilization and waste due to forgotten and endlessly running processes do not occur. This process also ensures that organizations pay for the portion of the resources that they use in the System z environment for high-cost efficiency.

This IBM Redpaper™ publication is a reference to various use cases that are available to enable and manage cloud computing by using System z and Moab Cloud Suite. By removing barriers to utilization, preventing overutilization and resource waste, and automating the process of administering virtual server services, Moab Cloud Suite enables the use of System z resources right up to the maximum of what the platform can sustain, with minimal human input. Moab Cloud Suite is the product of more than 10 years of collaboration between Adaptive Computing and IBM. This publication is intended for system administrators, IT executives, cloud architects, and other IT professionals who are looking for ways to maximize System z return on investment and to optimize the value and integration of multiple IBM environments.

Benefits of a Moab Cloud Suite and System z solution

The Moab Cloud Suite unique multi-dimensional policy engine and cloud management tools enable rapid implementation of an optimized cloud on System z environments that has the intelligence and automation to efficiently meet the needs of today's business. Combined with System z, Moab Cloud Suite can provide the following benefits:

- ▶ Reduce service provisioning and delivery costs up to 70% on System z through automated self-service virtual server and multi-VM service provisioning that minimizes manual, repetitive administrative tasks and reduces staffing requirements related to complex virtual server services.
- ▶ Improve utilization by automating the decommissioning and reuse of virtual servers that are no longer needed or in active use, and by leveraging standard catalog virtual server definitions that prevent the overprovisioning of virtual servers.
- ▶ Enable more efficient usage decisions by making virtual resource usage costs transparent to users and business groups, so that they are far more careful about the services they request and keep live.
- ▶ Improve SLA performance through intelligent policies that optimize service placement based on current logical partition (LPAR) load and demand, therefore maximizing performance.
- ▶ Reduce business risk and disruption by eliminating incidents caused by self-service under-provisioning of virtual resources and manual configuration errors.
- ▶ Speed delivery of needed IT virtual services to better support business opportunities and processes.

Moab Cloud Suite key capabilities

Moab Cloud Suite consists of the key capabilities that are described in this section.

Self-service request catalog

IT professionals everywhere are looking for ways to build efficient workload management into cloud environments. However, making that goal a reality requires reducing the workload of System z administrators. Specifically, administrators can reduce the time that they spend manually provisioning virtual servers and associated services (network, CPU, memory, I/O, and so on).

The Moab Cloud Suite portal-based self-service catalog can reduce manual virtual server administration requirements, accelerate delivery of virtual services, and reduce operating costs. Within the self-service catalog, the System z administrator establishes policies and sets parameters. Then, based on permissions, other administrators or authorized users can administer the process—making service requests, checking service status, and accomplishing virtual server lifecycle management tasks. Users can request standard virtual server services, multi-VM application services, and even custom virtual server services. Moab Cloud Suite automatically specifies the optimal resources that are needed.

Figure 1 illustrates how the Moab Cloud Suite self-service catalog portal automates the multi-VM request and creation process.

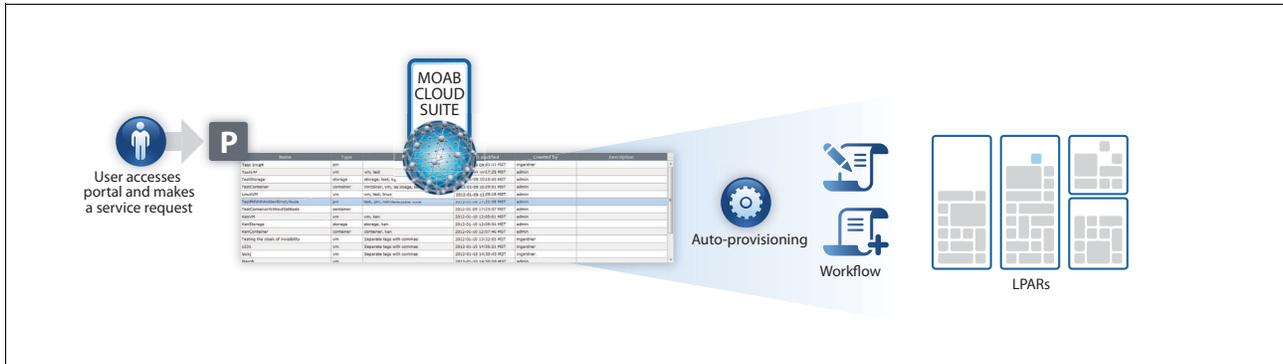


Figure 1 The Moab Cloud Suite self-service catalog portal

Aside from enabling the establishment of granular workload provisioning parameters, the Moab Cloud Suite self-service catalog provides modular, reusable service templates, which allow for services standardization. Moab Cloud Suite lets administrators create virtual server service templates that automate and enforce preferred practices, and resource setup processes for each service. This capability minimizes errors, increases compliance, speeds service delivery, and reduces management burdens and costs. In addition, administrators and authorized users can use proven Extreme Cloud Administration Toolkit (xCAT) profiles, so that they can build multi-VM profiles from standard building blocks.

The self-service catalog portal is consistent across multiple environments (that is, IBM System x, Power Systems, and System z). Therefore, there is no learning curve associated with accessing or ordering resources on disparate platforms. Moreover, with the Moab Cloud Suite self-service catalog in place, the organization can scale even with limited administrative staffing, because virtual services are started and stopped more efficiently, and LPAR utilization is optimized through easy access for more users. The result is easy, one-stop service creation and ultra-efficient use of resources.

Auto-provision optimal virtual service resources

Manually provisioning virtual machines and associated services can be time-consuming, costly, and error-prone. By automating the provisioning process, Moab Cloud Suite eliminates between 70% and 80% of misconfiguration issues. In addition, in just minutes, Moab Cloud Suite can deliver a complete multi-VM service in the specified, “right-sized,” policy-compliant configuration. For example, a web application configuration can be spun up automatically, comprised of a web server, database, and storage—all in the same request package—which can reduce the time required of an administrator.

In addition, virtual services can be standardized in templates and reused as an organization sees fit. Moab Cloud Suite can take advantage of xCAT open source, validated, and integrated provisioning infrastructure that customers already have in place. In addition, Moab Cloud Suite can reuse existing IBM z/VM® and xCAT profiles in the service catalog to make them easy for users to request as a service or part of a service. Reusing these existing profiles reduces risks and accelerates virtual service design, spin-up, and delivery. Moab Cloud Suite then automates the provisioning through xCAT to reduce total provisioning time from weeks to just minutes. Moab Cloud Suite auto-provisioning capabilities make the process inherently scalable and able to easily handle capacity utilization for multiple LPARs, as illustrated in Figure 2.

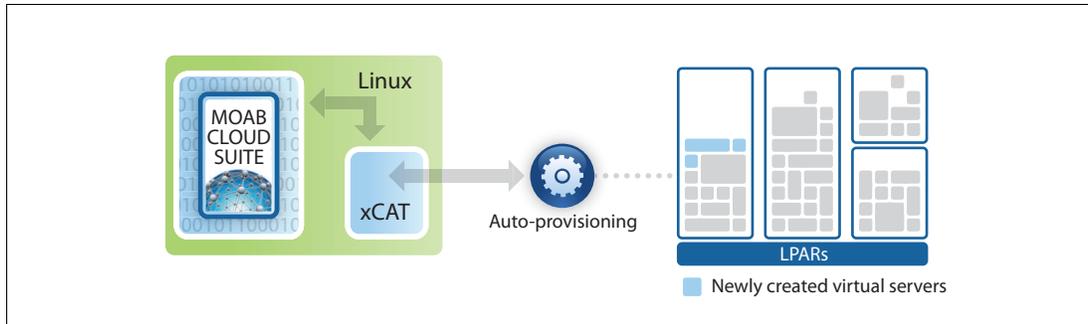


Figure 2 Moab Cloud Suite auto-provisioning capabilities

Adaptive Computing provides nearly a decade of experience and integrates with the xCAT provisioning that is used with the System z platform. Tens of thousands of servers are already under management using the Moab intelligence engine and xCAT profiles.

Optimized service placement

Better system scaling and higher administrator productivity are possible for System z with the optimized virtual server placement that Moab Cloud Suite provides. Moab Cloud Suite automates the load-balancing process between multiple LPARs based on predefined policies. It analyzes multiple load factors collected from Systems Management API (SMAPI) or CP sources in each LPAR, such as available versus in-use memory, disk space, CPU, and I/O, then determines the optimal LPAR in which to place a workload. The process is precise, efficient, and automatic. Administrators are no longer required to make initial placement load-balancing decisions. And, frankly, intelligent policies do a better job of it. Manual placement is no longer necessary either.

Figure 3 shows how optimized placement policy determines the best LPAR for a new virtual server to be created to ensure service performance SLAs will be met.

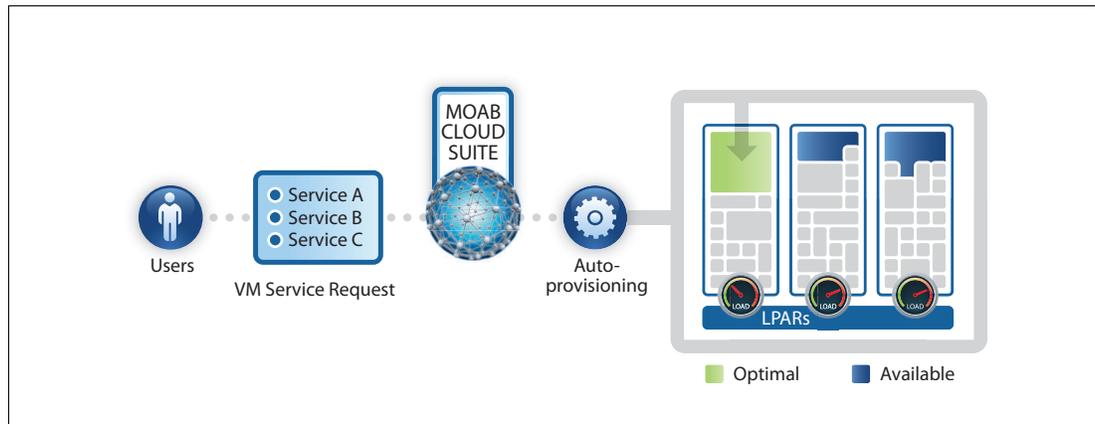


Figure 3 Moab Cloud Suite optimized placement policy

Adaptive Computing is collaborating with IBM engineers to research additional metrics that can further optimize virtual services placement such as I/O or network considerations among others. The goal is to assess multiple dimensions of resource consumption and allocation—identifying not just what has been requested but what is actually being utilized when automating the process of selecting which LPAR is a better candidate for additional loading. Research findings will lead to more granular and nuanced policy-based placement that can prevent LPAR overloading and ensure with greater certainty that service performance requirements are met.

Automated showback and chargeback

These capabilities are useful in any IT environment, but they are especially important in System z environments because of the extreme scalability options that the System z platform offers. *Automated showback* reveals how services are being used and by whom. *Automated chargeback* provides the accounting mechanism for billing resource usage against budget quotas, as well as charging different rates based on things, such as SLAs or usage metrics.

With Moab Cloud Suite showback and chargeback capabilities, costs are transparent and can be set to be based on usage. Administrators have the flexible control that enables creating custom charging models, and models can be set up on an information-only basis or for live, real-time notification and enforcement. In addition, authorized users can see what their costs are going to be before they order virtual services, as the portal can be set up to require that they place those services in a shopping cart before ordering them so that they can plan accordingly.

Figure 4 illustrates how the Moab Cloud Suite automated showback and chargeback policies enable fair use and eliminate wasteful use of System z resources.

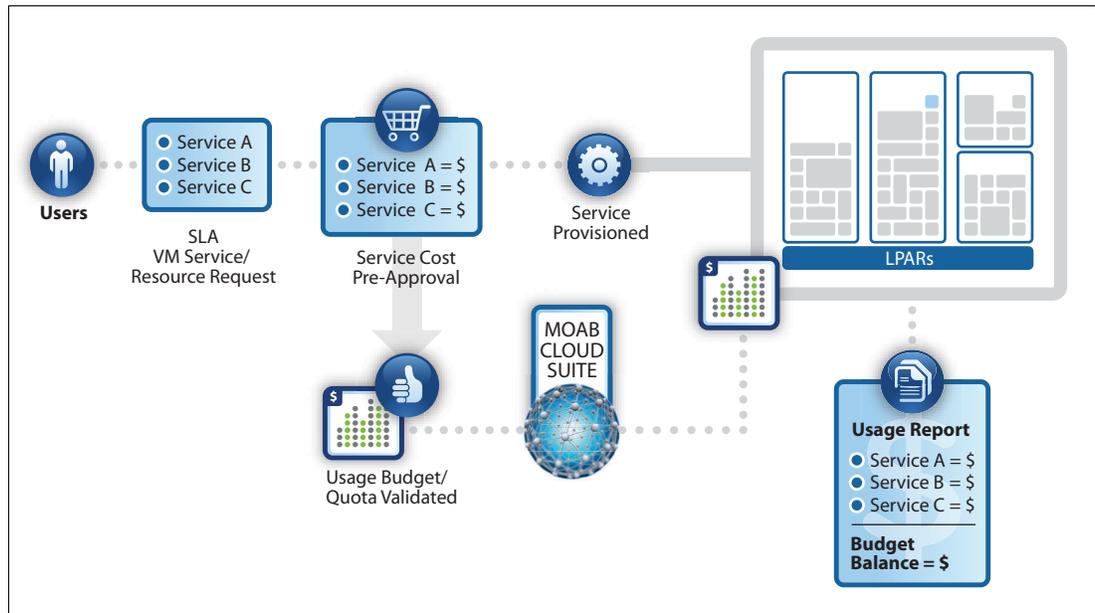


Figure 4 Moab Cloud Suite automated showback and chargeback policies

Moab Cloud Suite includes the following automated showback and chargeback capabilities:

- ▶ Total service cost calculations across all resources for preapproval and to help requesters avoid costs and waste.
- ▶ Flexible charge rates and models based on each resource type or service level. Supports showback only, setup and ongoing periodic usage tracking, budget/quota enforcement, or chargeback accounting modes. Itemized breakdowns of each charge are available for transparent pay-for-use chargeback.
- ▶ Replication of any organizational hierarchy structure so usage charging and showback reporting can be done against business units, cost centers, workgroups, and so on.
- ▶ At-a-glance views of event notifications give administrators a window into failed or successful accounting events that need to be acted upon. This minimizes management of usage budgets and quotas across multiple groups by administrators. In addition, groups using the cloud always have immediate visibility into their usage and budget availability.
- ▶ Budget or usage limits and enforcement, validating service requests against accounts on a per-user or project basis.

Self-service virtual server lifecycle management

The Moab intelligent policy engine is workload-centric rather than resource-centric. This characteristic enables awareness of services and SLAs that are associated with each virtual machine. Moab Cloud Suite ties this information together, giving administrators a dashboard view that lets them see, control, migrate, and terminate virtual server services as appropriate. In addition, Moab Cloud Suite accelerates the migration to a user-centric IT services delivery model.

With Moab Cloud Suite, authorized users can also see and control workloads in the user portal and quickly view and filter services by almost any service attribute. Users can also assess the health of their virtualized environments. This control is especially important in test or development environments where users can request resources and schedule them to be automatically deprovisioned within a planned time frame, so that start and stop times can be efficiently calendared. From a virtual server lifecycle perspective, this level of computing intelligence and transparency makes users more responsible while reducing the administrative burden of IT.

Figure 5 illustrates the Moab Cloud Suite user and administrative dashboards. These dashboards provide access to details and management actions for services and virtual servers across their lifecycle.

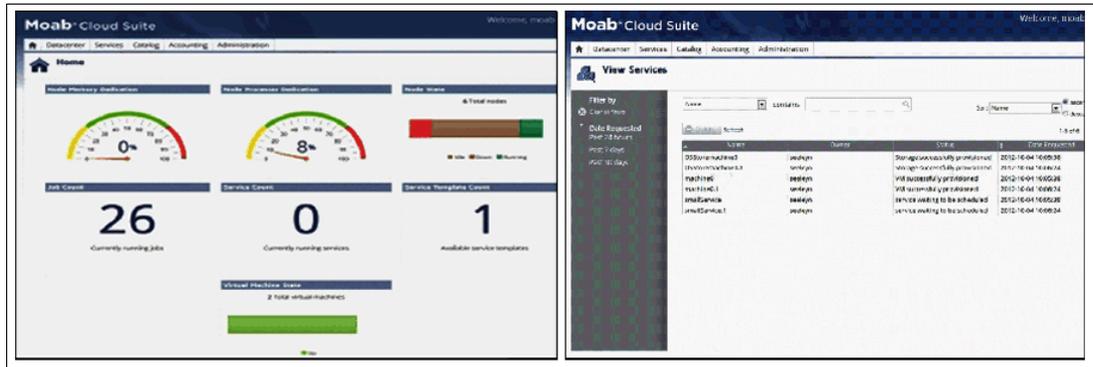


Figure 5 The Moab Cloud Suite user and administrative dashboards

Automated decommissioning of virtual services

Moab Cloud Suite enables automated decommissioning of virtual services based on the expiration of the original request or manually through a self-service portal when an administrator or authorized user initiates a deprovisioning order. Either way, resources can be reprovisioned quickly and adapted for new services. Virtual services can be set to have a predefined and finite lifecycle, with extensions only through an explicit request from an authorized user. Infinite services can also be allowed if desired for certain virtual services.

Figure 6 illustrates how Moab Cloud Suite automated decommissioning policies and capabilities reduce resource waste and maximize available capacity.

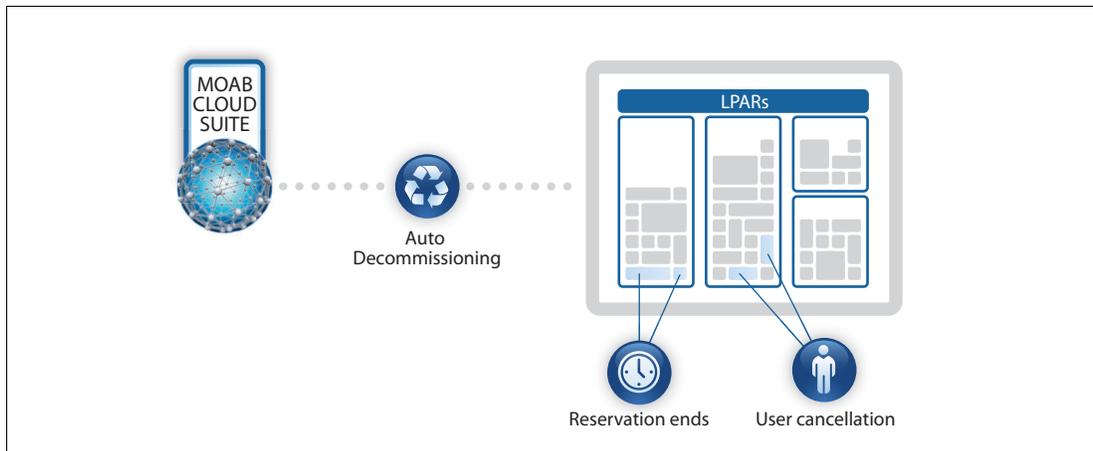


Figure 6 Moab Cloud Suite automated decommissioning policies and capabilities

Automated decommissioning replaces labor-intensive manual processes that can leave resources unused for long periods of time when they can be put to use for new business services and initiatives. In addition, automated decommissioning keeps operational and capital costs down by maximizing utilization and available capacity. Also, similar to the Moab Cloud Suite in its entirety, the automated decommissioning capability maximizes return on investment by optimizing the efficiency with which System z environments are used.

An open, policy-based cloud architecture

Adaptive Computing is providing an important option that IBM System z clients can use to establish heterogeneous cloud environments based on open source technology. Moab Cloud Suite leverages and integrates with existing IBM System z environments, including those that want to leverage an open source xCAT provisioning infrastructure. It harnesses Moab Cloud Suite and System z metrics information to power policy-based optimization around provisioning, usage accounting, and efficient load balancing. This process enables Moab Cloud Suite to provide capabilities in a light technology implementation that complements System z without adding complexity.

This open, policy-based architecture, and Adaptive Computing's continuing collaboration with IBM and IBM customers, makes it ideal to address additional challenges in the future, such as unified cloud management across multiple System z mainframes, and optimized service request and placement for System z, System x, and Power Systems environments.

Moab Cloud Suite open cloud architecture, illustrated in Figure 7 with the System z hardware control point (zHCP), uses xCAT and policies powered by data within each system with a future vision to unify management across different IBM system environments.

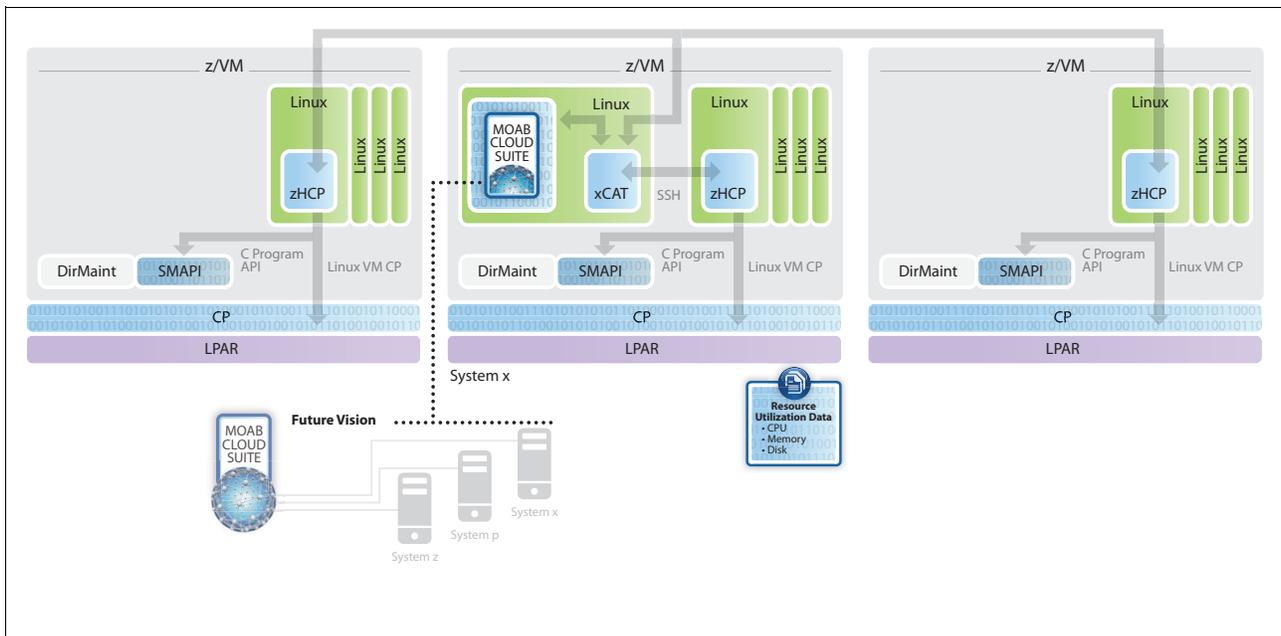


Figure 7 Moab Cloud Suite open cloud architecture

Conclusion

CIOs are looking for a single cloud solution that encompasses a common administrative domain on top of all their diverse resources and architectures. Moab Cloud Suite provides the policy-based optimization and automation with the self-service user interfaces proven in some of the most demanding cloud environments and brings them, tuned and honed, to System z environments. Moab Cloud Suite is unique as the solution that can be built on top of existing System z toolsets—no rip and replace required. In addition, it is the solution that enables organizations to maximize the efficient use of System z environments—right up to the highest levels the architecture can sustain, for maximum return on investment.

With Moab Cloud Suite, System z gains a simplified user's self-service experience, automated provisioning with no System z expertise required, usage charging or accounting to reduce waste and overprovisioning, and a highly efficient load-balancing system. That is the promise of the Moab Cloud Suite approach: A unified experience throughout the enterprise.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization.

David Jackson is an established thought leader with over 16 years of experience in policy-driven automation and workload management for adaptive enterprise data centers, cloud, and high-performance computing (HPC) environments. As a result of his early pioneering efforts, Dave has been awarded 29 patents in resource management, cloud, data center automation, and workload management and has more than 100 additional foreign and domestic patent applications pending. David has consulted on solutions for some of the world's largest and most advanced private clouds and largest HPC sites. David is the Founder and CTO of Adaptive Computing and is also the Executive Director of Adaptive Computing Labs, which leads the research and development efforts to evaluate, prototype, and advance strategic technologies.

David Handy is a software architect and a leading member of Adaptive Computing Labs, a research and development team working on advanced projects within the company. His primary area of interest is integrating Moab Cloud Suite into new platforms and environments. David has 10 years of experience developing multitier business applications. He earned a BS degree in Electrical and Computer Engineering from Brigham Young University.

Kershaw Mehta is the IBM Chief Architect for Cloud Computing on System z. He has over 20 years of experience in System z as a software engineer in areas, such as open standards, virtualization technologies, and application enablement. He has led development teams around the globe and often speaks at worldwide events on various System z topics. His current focus is enablement of the System z platform for cloud computing.

Thanks to the following people for their contributions to this project:

Debbie Willmschen
International Technical Support Organization, Raleigh Center

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks publications

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-5008-00 was created or updated on May 29, 2013.

Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

| | | |
|----------------|---|-----------|
| IBM® | Redpaper™ | System z® |
| Power Systems™ | Redbooks (logo)  ® | z/VM® |
| Redbooks® | System x® | |

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.