

The Journey to Cloud Computing: from experimentation to business reality

Judith Hurwitz, President
Marcia Kaufman, COO



**HURWITZ
& ASSOCIATES**
Insight to Action

Sponsored by IBM



The Journey to Cloud Computing: from experimentation to business reality

The greatest impact of cloud computing on enterprises is its ability to more efficiently and quickly satisfy business requirements. The typical organization now relies on IT based services as the foundation for new products and services and as the primary way to focus on customer value. It is a well-known story that companies have too many underutilized IT resources in their data centers. Ironically, despite the existence of these untapped resources, IT is often unable to respond to the needs of developers and business units in a timely and cost effective manner. In order for the business to meet growth and profitability objectives, it requires higher levels of performance and agility from IT. Thus, cloud computing is becoming increasingly important because it enables organizations to make efficient and secure use of a pool of IT resources that can expand and contract based on the business needs.

In this paper, Hurwitz & Associates will provide insight into the journey to enterprise cloud computing from a quality of service perspective. Many companies have initially looked to virtualization to drive higher utilization of IT resources, lower capital expenses, and enable the scalability required to support business growth. These companies have also focused on standardizing and automating IT environments to reduce costs and improve the quality of service to customers. However, as organizations move from experimental cloud implementations to using the cloud for mission critical workloads, quality of service requirements take on a new dimension. Service issues that may have been overlooked with less critical workloads may become magnified, exposing the company to significant risk. To protect against higher levels of risk it becomes increasingly important to translate existing internal data center policies into policies that make sense in the cloud. Therefore, we will provide insights into the key elements that these organizations need to focus on to ensure that policies for security, governance, and scalability are maintained as mission critical workloads are moved to the cloud.

In addition, we will provide an overview of what capabilities IBM Power Systems™ cloud solutions offer to support customers.

While the business imperative is clear, the journey to leverage the cloud in a way that protects the company's assets while providing a platform that is secure, scalable and easy to manage, can be a challenge. In addition, enterprises needing to balance a complex mix of workloads can find it daunting to ensure that quality of service requirements are met in hybrid IT environments that include the traditional data center, private clouds, and public cloud services. Therefore, many customers have approached the journey to the cloud with caution. Management wants assurance that the platform itself will not compromise the integrity of company confidential data. In addition, management wants to know that performance can be guaranteed and that the quality of service required by customers, suppliers, and internal business units will be supported. Therefore, many companies have begun their journey by experimenting with Software as a Service (SaaS) applications in areas such as

Service issues that may have been overlooked with less critical workloads may become magnified, exposing the company to significant risk.



**HURWITZ
& ASSOCIATES**
Insight to Action

sales support. Other starting points in the journey to the cloud have included the adoption of private and public cloud services for software development and test as well as data storage and retrieval.

The next phase, however, is the most important for businesses – creating private clouds and leveraging enterprise focused commercial cloud services as a means of transforming the way technology is used to become an agent for business growth and change.

The Commercial Cloud Transition

Now that many innovative businesses understand the foundational requirements and benefits of cloud computing, they are beginning to evaluate how to put cloud computing in context with their enterprise requirements. In order to accomplish this goal, companies need to start thinking about the specific requirements for managing mission crucial workloads and what this means in a cloud-based environment. There are a number of key requirements including:

- There needs to be an overarching governance strategy that reaches across the enterprise to ensure the integrity of all types of workloads.
- Security, scalability, reliability, and performance all need to meet the highest standards when workloads are critical to business operations.
- Managers need to consider how to manage quality of service objectives without compromising on these important requirements. For example, you may want to support sophisticated service and performance management requirements in a hybrid environment without expanding staff overhead.

Other important considerations include eliminating risk of downtime while maintaining cost control and eliminating unexpected limitations when leveraging virtualization to increase performance.

Managing Mission Critical Workloads – A Customer Perspective

Today many businesses find that they have an increasingly diverse group of IT workloads—such as customer-facing web portals and email services – and more of these workloads are considered mission critical. IT workloads – defined as an independent service or collection of code that can be executed - can be a simple set of business services or a full set of applications. Examples of business workloads range from a simple sales management application used by a small sales team to enterprise-wide order, inventory and customer service applications or compute intensive analytics applications. Faced with a complex mix of workloads to manage, more companies want to take advantage of the scale up and scale down benefits of cloud computing. However, these organizations often find it confusing to try to determine what types of workloads are most appropriate for cloud computing. Many IT executives are concerned about giving up control of the security and overall management of the workloads that are

... companies need to start thinking about the specific requirements for managing mission crucial workloads and what this means in a cloud-based environment.



**HURWITZ
& ASSOCIATES**
Insight to Action

tightly integrated with business operations. Therefore, many companies follow a multi-phased approach that allows IT management to gradually incorporate critical workloads into their cloud environments. The experience of a typical company is highlighted below.

Phase 1 - Exploratory/Pilot. The company selected the process of developing and testing applications as the first phase of its cloud computing journey. The intention was to establish a best practices approach to providing developers with a self-service option while working with non-production workloads. The company established a private cloud and also supported the use of a public cloud option. The idea was to give each developer “a personalized copy in the cloud” and allow him or her to use the IT resources they needed for a limited amount of time. Development and test workloads were selected for this initial phase because there were fewer concerns related to security and information governance. One specific pilot implementation was to provide a small ecosystem for the SAP development staff to do some testing for a couple of weeks. They wanted the team to be able to select a system from a repository. One of the benefits of this exploratory phase was that developers learned first-hand how easy it is to make use of the right IT resources at the right time with a self-service approach. They also learned about image management and what it was like to move between public and private clouds. By focusing the pilot effort on development and test, the IT management team was able to understand how the developer community would react and accept the self-service cloud without exposing the production environment to any potential risk for this new approach.

Typically, a development task involves a set number of developers who need access to a server, an operating system, and a set of development tools. Therefore, the workloads are bounded and precise. The same can be said of the software testing process. Starting with development and testing is a good starting point because it helps companies understand what it means to provision resources in a self-service manner based on a policy set by management. It also makes it easier for management to keep track of resources in a more effective manner. At the end of the development cycle, those resources used for development are put back into the resource pool so that they are available for the next project. Additional workloads more common in the exploratory phase include other pre-production workloads and disaster recovery.

Phase 2 - Support Mission Critical Workloads. The workloads supported in the company’s first phase of cloud computing were development projects with a small team tackling a specific task that has a set beginning and end. Because these are development projects they don’t impact operational systems. In contrast, mission critical workloads are often either customer facing or are essential to the day-to-day operations of a business. Although the company was initially concerned about business risk, the success with early cloud deployments convinced them to expand on the types of workloads they could support in the cloud. They are currently extending their cloud infrastructure to take on mission critical workloads.

By focusing the pilot effort on development and test, the IT management team was able to understand how the developer community would react and accept the self-service cloud ...



Phase 3 - Support Range of Workloads with Hybrid Cloud Environment.

The economics of managing workload requirements is a critical issue for this company. They want to have the flexibility to be able to manage capacity on demand. In the past, it has been necessary to anticipate workload requirements years in advance. The IT team has been plagued by dramatically changing business conditions making it impossible to plan far ahead for workload requirements. For example, there have been several major acquisitions that led to unexpected changes in workloads. In addition, there have been short bursts of computing needs based on unanticipated projects or business changes.

With the growth of cloud computing, management now understands that there are more options than ever. The company expects its management of mission critical workloads to require a focus on a hybrid approach to cloud computing. The organization will need to create linkages across its traditional data centers, its private cloud environment and its public cloud services to meet its diverse computing requirements. This may include everything from a web facing customer sales channel to transaction management systems. Rather than needing to estimate IT requirements in advance and hoping to be prepared for business changes, the company can use external as well as internal cloud resources to augment the data center capacity when needed.

As this company moves to incorporate mission critical workloads into its cloud strategy, one of its key concerns is on providing the right quality of service.

The quality of service imperative – how do you establish quality of service?

The first priority for any company wanting to leverage the benefits of cloud computing for mission critical workloads is to focus on quality of service and scalability. If either of these characteristics is missing, the business will be at risk. So, what does this mean from an operational perspective? First, organizations have to think about the purpose of putting mission critical workloads into a cloud computing environment. For example, is the requirement for day-to-day management of an ERP system or a customer facing portal? In some situations, a company may want to provide a critical self-service application to their suppliers and partners. In other cases, a cloud computing service might be used as a disaster recovery platform or as a technique for being able to acquire additional computing resources during a peak period. Despite the fact that there are many different reasons to leverage cloud computing there are some common issues relevant to all customers. They include:

- Automated provisioning of resources
- Leveraging virtualization to optimize performance and ensure application availability
- Ensuring security of information
- Having the ability to monitor and control performance and service within the cloud environment

... organizations have to think about the purpose of putting mission critical workloads into a cloud computing environment.



**HURWITZ
& ASSOCIATES**
Insight to Action

- Having the ability to optimize the environment to reduce operational expenses

Automated provisioning of resources. IT resources need to be provisioned in a way that is consistent with the varied requirements of a mix of different workloads. For example, in many situations, customers need to be assured that their workloads will support self-service or on-demand provisioning. However, if a mission critical workload begins to expand because of increased demand, additional resources need to be made available as quickly as possible. In order to quickly respond to this change in demand, automated provisioning of additional resources to support the new demand must be supported. There may also be situations where workloads become less active and in order to optimize their asset utilization, an organization might also like to automatically deprovision resources to reduce the number of resources actively deployed.

Leveraging virtualization to optimize performance and ensure application availability. A well-designed virtualization environment allows IT to consolidate workloads running on different computing platforms. This enables these platform resources to be allocated dynamically to meet mission critical workload requirements. Within the virtualized environment, workloads can be transferred between servers, mission critical workloads can be balanced so that performance is optimized and application downtime is reduced. For example, when a university medical center recognized that it was quickly running out of CPU as a result of a fast rate of growth in the business, it used processor virtualization to leverage untapped capacity without requiring the company to purchase additional IT resources. The organization uses shared-processor pools to achieve close to 80% CPU utilization across its servers. A wide variety of workloads can be balanced throughout the day to achieve optimum performance.

Ensuring security of information. Whether a cloud computing service is deployed in a private environment or based on a public capability, security of mission critical workloads and their data is an imperative. Identity management needs to be addressed for users who may use a service on a continuous basis or for short durations. Virtual machine images must also be secured since they can open up access points for unauthorized users. In addition to virtualization, security reaches many levels within cloud computing including storage, data, network and applications. Organizations need to be able to secure the environment, the interactions between environments, and be able to govern their overall security.

Monitoring and controlling performance and service. Cloud computing workloads do not live in isolation. Most organizations have a hybrid environment including both traditional data centers as well as both public and private cloud services. Therefore, it is critical that the IT organization be able to monitor the entire environment based on the required service level and overall quality of service of these workloads. The more mission critical these workloads are, the more important it is for IT be able to have oversight into overall performance, availability, and scalability.

A well-designed virtualization environment allows IT to consolidate workloads running on different computing platforms.



**HURWITZ
& ASSOCIATES**
Insight to Action

Optimizing the environment to control costs. The ability to reduce overhead expenses is one of the primary reasons that businesses are beginning their journey to cloud computing. Many organizations assume that the process of acquiring cloud services will reduce capital expenses overall. However, the benefits of cloud computing can be much more far reaching. By balancing and optimizing those workloads, the IT organization can actually turn on and off servers and storage systems, automatically without IT intervention, when their capacity is not needed. This type of automated and granular control over internal and external computing resources can make existing resources more efficient, reduce IT management costs and help organizations do a more accurate job of responding to changing business demands. By turning off unused resources, IT can save on electricity and power. In addition, IT can avoid purchasing excess capacity for peak loads. Instead, IT management can leverage sophisticated cloud service providers through self-service provisioning to meet peak demands or handle high-intensity compute workloads that it might not be able to support in its internal data center.

These qualities of service concerns are important for companies as they look to transform their data centers into a more flexible, scalable, and self-service environment. They understand that hybrid cloud environments will be a reality if they expect to leverage the benefits of cloud computing for their mission critical workloads. The following section provides an overview of IBM Power Systems cloud solutions and how they are designed to help companies support mission critical workloads.

How cloud solutions from IBM Power Systems can support enterprise cloud computing

Once companies are convinced that they can safely move to the cloud for non-production workloads, they often are ready to move to use the cloud to support mission critical workloads. At this point, they need to be sure that the quality of service provided will match up to the increased requirements. Because characteristics such as virtualization, security, and energy management are built into the POWER7 hardware, it can support these quality of service requirements for workload management. For example, POWER7 systems are designed to improve performance by as much as five times compared with their POWER6 predecessor. This can be accomplished without increasing the physical size of the server. The platform's active memory sharing is intended to distribute memory across workloads so that they can be balanced automatically without requiring the application owner to make changes. Customers are also able to manage overall virtualization in the data center through a graphical interface based manager. Another key requirement for many customers moving to cloud computing is to improve energy management and expenses. The built-in energy optimization of the POWER7 will help companies be successful from an energy management perspective. The system evaluates the workloads and based on user-established requirements for performance and quality of service, it can automatically turn off unused capacity.

Because characteristics such as virtualization, security, and energy management are built into the POWER7 hardware, it can support these quality of service requirements for workload management.



For customers planning to move to cloud computing, the fact that the POWER7 platform can support the familiar IBM AIX operating system along with IBM i and Linux operating systems, has proven to be a benefit. By supporting multiple operating systems on the POWER7 systems, customers avoid costly retraining and additional software acquisitions. At the same time, the fact that virtualization is built into the firmware and expanded with IBM PowerVM means that not only can performance be optimized, but enterprise-class security can be achieved.

PowerVM is the built-in virtualization capability in the POWER7 environment. It enables customers to consolidate large and diverse groups of workloads onto fewer systems resulting in greater efficiency and lower costs in the data center. There are three key types of virtualization controlled by PowerVM: processor virtualization, memory virtualization, and I/O virtualization. The pooling concept is important across all areas of virtualization enabling resources to be called on when needed to boost performance. As many as 10 partitions are allowed for each processor core and as a result of the resource pooling, processor, memory, and I/O resources can be moved between these partitions to support changing workload requirements. With processor virtualization, a single partition in the server can act as a separate AIX, IBM i, or Linux operating environment. PowerVM provides the dynamic logical partitioning (LPAR) capabilities that make this possible. Each of these partitions can have dedicated or shared processor resources. The memory virtualization capability of PowerVM provides for allocation of memory where and when it is needed most to help increase flexibility and performance. Physical memory can be shared across the logical partitions in a single server. In addition to providing for shared memory, PowerVM allows for portions of memory to be dedicated for specific purposes. The role of I/O virtualization is handled by the Virtual I/O Server (VIOS). It is used to enable the sharing of I/O resources. Network adapters, disk adapters, disk drives, tape adapters, and tape drives can be shared across multiple partitions. Similar to the situation with memory virtualization, VIOS hosted I/O for some partitions can co-exist with dedicated I/O devices for other partitions in the same system.

Advanced virtualization with PowerVM can be used to consolidate very large numbers of database and application workloads because as many as 1,000 partitions per server are allowed. With PowerVM, customers are able to dramatically increase the number of CPUs and the memory available to the most demanding workloads. When deploying mission critical workloads with PowerVM, customers can scale to use all the resources of the host server. In addition, these IT resources can be changed as needed without requiring a reboot of the system. PowerVM Live Partition Mobility enables active partitions to be moved between servers, virtually eliminating planned downtime. Live partition mobility can also be used to upgrade workloads between POWER6 and POWER7 processor-based servers without an application outage. With POWER7, PowerVM and IBM Systems Director VMControl™ virtualization software will support up to 1,000 virtual machines on a single system.

With processor virtualization, a single partition in the server can act as a separate AIX, IBM i, or Linux operating environment.



IBM Systems Director VMControl is designed to complement PowerVM by providing automated virtualization management that minimizes time to provision virtual machine images and enables the creation and management of system pools. One of the key requirements for working with mission critical workloads is the ability to dynamically add and remove virtual resources as your needs change and VMControl provides this capability. VMControl automates the management of both physical and virtual servers providing a single virtualization management tool. VMControl is used to create and manage standardized virtual images. By automating this provisioning step customers have been able to reduce deployment times from days down to hours. This management process allows companies to optimize system pools. The host system can monitor for failures, predict failures and automate the process of recovering from the failure as well.

IBM Power Systems cloud solutions include several additional product offerings that are intended to support customers moving to cloud computing environments. These include IBM CloudBurst, IBM Service Delivery Manager, and the IBM WebSphere CloudBurst Appliance. IBM has packaged hardware components, software, storage, and services together to help companies deliver private clouds quickly and cost effectively. The next section will provide a description of each of these offerings.

IBM CloudBurst

The IBM CloudBurst on Power Systems solution is a tightly integrated cloud service delivery platform built on the IBM Power Systems platform, bringing the advantages of the latest POWER7 processor-based system capabilities to the cloud environment. It includes a comprehensive set of components required to create, and manage a private cloud within the data center. This includes the network, virtualized CPU, memory and storage as well as the platform and service management software. These components are integrated, pre-installed and pre-configured at the factory and are sold, delivered and supported as a single product making it easier for companies to get started. Implementation and support services are included in the solution and provided to ensure the solution is up and running as quickly as possible.

The IBM CloudBurst on Power Systems solution includes a self-service portal that enables users with authorization to request resources based on a service catalog and have those resources dynamically allocated. Customers can use the portal to reserve the various IT resources ahead of time via the self-service interface ensuring that the requested services are made available to users at the time they are needed, in a safe manner and in accordance with governance rules. The solution provides an easy way to access the service catalog that helps IT track and manage the codified components and services they offer to their customers. The automated provisioning capabilities provided not only provision requested resources, but also help companies control the recovery of IT resources so they can be quickly returned when no longer needed. From

VMControl automates the management of both physical and virtual servers providing a single virtualization management tool.



**HURWITZ
& ASSOCIATES**
Insight to Action

a management perspective, the IBM CloudBurst solution includes platform, virtualization, energy, storage, service, and financial management capabilities that provide support for tuning workloads for high availability, energy optimization, flexible billing for services and IT cost reduction.

The IBM CloudBurst on Power Systems solution is designed for customer private cloud implementations from initial pilot projects through to mature hybrid cloud environments incorporating mission critical workloads. Customers have found the IBM CloudBurst solution to be a cost-effective platform for a mix of workloads including software development, test and pre-production, and data intensive processing. Some examples of specific workloads are SAP, Oracle, Sybase, DB2, Business Intelligence, and Transaction processing.

IBM Service Delivery Manager

IBM Service Delivery Manager for Power Systems is a pre-installed and pre-integrated service management software stack intended to be used in private cloud implementations. It is deployed as a set of virtual images that automates IT service deployments and provides self-service provisioning, resource monitoring, cost management, and high availability of services in the cloud. One of the key features of this offering is the self-service portal designed to make it easy for customers to provision and de-provision IT resources as needed. The customer chooses the services they want from the services catalog available through the portal. IT administrators can define the services and automated workflows to ensure that cloud resources are provisioned in a consistent and standardized way throughout the enterprise. Users of cloud resources benefit because they can request services when needed for the right amount of time. With a more highly automated provisioning and de-provisioning process, administrators and users benefit because there are fewer system outages caused by human errors.

IBM Service Delivery Manager includes an integrated process for monitoring physical and virtual cloud resources. This process enables organizations to analyze usage of IT assets that may be shared across different business entities and bill each unit for the resources they actually consumed. The monitoring, usage, and accounting capabilities provide comprehensive reports of how resources are used so that IT is able to plan better for future needs and improve customer satisfaction. IBM Service Delivery Manager for Power Systems also includes platform and virtualization management, so that customers not only get a complete picture of their systems and how well they are operating, but also the tools to deploy, optimize and maintain these systems at maximum effectiveness and efficiency. This also enables IT managers to automatically migrate virtual machines across systems to maintain or improve service levels. IBM Service Delivery Manager enables customers to quickly deploy the cloud service management stack onto their existing POWER6 or POWER7 resources and integrate with their existing IT infrastructure to create a private cloud computing environment.

With a more highly automated provisioning and de-provisioning process, administrators and users benefit because there are fewer system outages caused by human errors.



IBM WebSphere CloudBurst Appliance

The IBM WebSphere CloudBurst Appliance speeds application deployment to cloud and virtualization environments. It is intended to leverage existing data center resources and implement well tested WebSphere virtual images and patterns. Although the WebSphere CloudBurst Appliance comes pre-loaded with best-practices patterns, it is designed so that customers can choose to customize these pre-configured patterns based on company experience and requirements. The preloaded patterns range in complexity from a simple single server topology to a highly available cluster topology.

The best practices focus enables a consistent and repeatable deployment of an application environment based on patterns. Because the Appliance is not designed as a black-box environment, customers can apply maintenance and federate cells. In essence, within the WebSphere CloudBurst Appliance, WebSphere patterns are implemented into a private cloud based on virtualized hardware running supported hypervisors including IBM PowerVM. It also can be used to integrate with existing infrastructure through programmable REST APIs.

Conclusion

Organizations are moving to cloud computing as a means to help create more efficient and cost effective data centers. Many of these organizations begin their journey to the cloud with self-contained workloads such as development and test. However, after seeing the benefits of providing a self-service environment for developers, many of these organizations are beginning to reach a tipping point. They recognize that as they move along a path towards incorporating mission critical workloads in to their cloud strategy, they need to re-evaluate requirements for quality of service and scalability.

As a company constructs its private cloud environment, decisions that are made at the virtualization layer will impact the capabilities that the infrastructure will be able to provide for future workloads. Mission critical workloads demand the highest levels of scalability, security, and reliability. Power Systems cloud solutions provide many options for customers to take advantage of the workload optimized systems, enterprise quality of service virtualization, automated platform management as well as the service management capabilities needed to support mission critical workloads. Companies wanting to support mission critical workloads in the cloud need a flexible, secure, and powerful virtualization environment.

As a company constructs its private cloud environment, decisions that are made at the virtualization layer will impact the capabilities that the infrastructure will be able to provide for future workloads.



About Hurwitz & Associates

Hurwitz & Associates is a consulting, market research and analyst firm that focuses on how technology solutions solve real world business problems. The firm's research concentrates on disruptive technologies, such as Cloud Computing, Service Oriented Architecture and Web 2.0, Service Management, Information Management, and Social and Collaborative Computing. We help our customers understand how these technologies are reshaping the market and how they can apply them to meet business objectives. The team provides direct customer research, competitive analysis, actionable strategic advice, and thought leadership. Additional information on Hurwitz & Associates can be found at www.hurwitz.com.



© Copyright 2010, Hurwitz & Associates

All rights reserved. No part of this publication may be reproduced or stored in a retrieval system or transmitted in any form or by any means, without the prior written permission of the copyright holder. Hurwitz & Associates is the sole copyright owner of this publication. All trademarks herein are the property of their respective owners.

175 Highland Avenue, 3rd Floor • Needham, MA 02494 • Tel: 617-597-1724
www.hurwitz.com