IDC EXECUTIVE INSIGHTS

Cloud Computing: The Need for Portability and Interoperability

August 2010

By Jean Bozman, Research Vice President, Enterprise Servers and Gary Chen, Research Manager, Enterprise Virtualization Software

Sponsored by Red Hat, Inc.

Introduction

Cloud computing has become a major discussion thread in the IT world, but it has been redefined in recent years to speak to the increasing number of deployment scenarios – and their practical uses for business workloads.

The business benefits of cloud computing are clear: the ability to leverage virtualized infrastructure for greater efficiency, and to be able to dial up — or dial down — a specific service, based directly on business need. Rapid provisioning is a key element in supporting this scalability and elasticity. In this era of economic constraints, the ability to "stand up" a new cloud service, and then use it only as long as it's needed, is a valuable business attribute of leveraging cloud computing within the enterprise.

Cloud services are driven by user demand, enabled by self-service catalogs of cloud workloads, and workloads are managed across the infrastructure in a transparent and self-healing way. This means that cloud services can be accessed by non-technical users; the technical details of managing the infrastructure are left to the cloud provider. As IDC defines it, cloud services are shared, standardized services that are available from a self-service catalog; able to scale in an "elastic" fashion as needed; priced based on actual usage; accessible via the Internet; and supportive of published application programming interfaces (APIs).

Public, Private, and Hybrid Clouds

Public clouds are the best understood type because there are multiple, high-profile public offerings available today, including Amazon Web Services, Google Apps, and salesforce.com. Early adoption of cloud computing was driven by application development, collaborative workloads (including email), and scale-out Web applications. Because anyone can subscribe to these services, familiarity with them is widespread and a business case can be made for their use, whether for special projects, or adoption of the public cloud for specific, well-defined workloads.

Because enterprises want to adopt cloud computing to break down the "silos" in their internal IT infrastructure, private cloud scenarios — in which cloud technology is deployed inside the firewall — have been growing over the last two years. Users of private clouds seek to apply the cloud computing model to their own IT infrastructure. Given the data-security and availability characteristics supported inside enterprise IT, these private cloud deployments are seen as supporting compliance regulations and governance standards acceptable to company auditors and government regulators.

Finally, the concept of hybrid clouds speaks to the dilemma of IT managers in industry segments that are highly regulated, such as financial services or healthcare. These industries must balance the wish to access outside, hosted cloud services with the need to retain some types of applications and data inside the security firewall of the enterprise or organization. A hybrid cloud makes it possible to tap both types of resources, public and private cloud, and to link these resources together into a cohesive pool.

Moving Between the Clouds

The ingredients for cloud enablement are well known: software stacks ready to provision, on-demand deployments, a self-service catalog of cloud services, the scalability to meet growing demand for cloud services — and the flexibility to scale down resources when they are no longer needed. Three major models for cloud deployments fall readily into place, delivering various layers of the computing stack as a service: IaaS (infrastructure as a service); PaaS (platform as a service); and SaaS (software as a service).

But the "who, what, when, where, and how" of enterprise-style cloud computing are only now coming into view. This is cloud computing that supports stateful applications, and places customers' concerns about security and multi-tenancy high on its priority list. It's the kind of cloud computing that many IT customers are waiting to see before widely adopting cloud computing for business-critical and mission-critical applications.

Portability and Its Role in Cloud Deployments

Now, if there are multiple types of cloud deployments, then it should be possible for an organization to move between them – or to "federate" cloud services on behalf of the end-users or end-customers of the enterprise.

That is why the idea of portability of cloud applications has become so compelling to enterprise customers: In the event that a given organization decides to switch cloud services, they want to know the "how" of switching and the opportunity cost of doing so. Given that cloud computing is such a new phenomenon for IT organizations, many of the blueprints for switching cloud services have yet to be drawn up – and some are in the beginning processes of deployment for public and private clouds.

While cloud providers have been working furiously to build out cloud features and to deliver enterprise-class features, IT has begun to worry about potential lock-in to specific technologies, such as software tools and programming interfaces.

Enablers for Switching Cloud Services

What would it take to move workloads from one cloud service to another? Key enablers would include the following:

- **Standardized programming interfaces.** This would allow applications written to those interfaces to be moved to a new "target" cloud. That is why well-identified programming toolsets can be leveraged by companies readying IT tasks for deployment on a cloud computing service.
- **Layers of abstraction.** Abstraction layers insulating cloud services from underlying infrastructure will ease portability between cloud resources. Reducing dependencies on underlying processors, operating systems and virtualization software will increase the speed of switching cloud services – and will reduce the cost and pain of doing so.
- **Management capabilities.** These would allow IT staff to easily manage the cloud services to which they subscribe – without introducing new layers of complexity into an already-complex

enterprise infrastructure. To be successful, these management capabilities would have to work well with existing management products, both for managing physical servers and virtual servers.

Benefits

According to IDC demand-side research, business customers see the following business benefits of incorporating cloud computing into their IT organizations (see Figure 1):

- Less IT infrastructure to buy and deploy
- Less IT management demands on existing IT staff
- Pay-as-you-go model for computing that would reduce both capital expenditures (capex) and operational expenditures (opex)

Figure 1

The Appeal of Cloud Computing

Q. Rate the benefits of the Cloud model

Scale 1 - 5; 1 = Not at all important, 5 = Very important

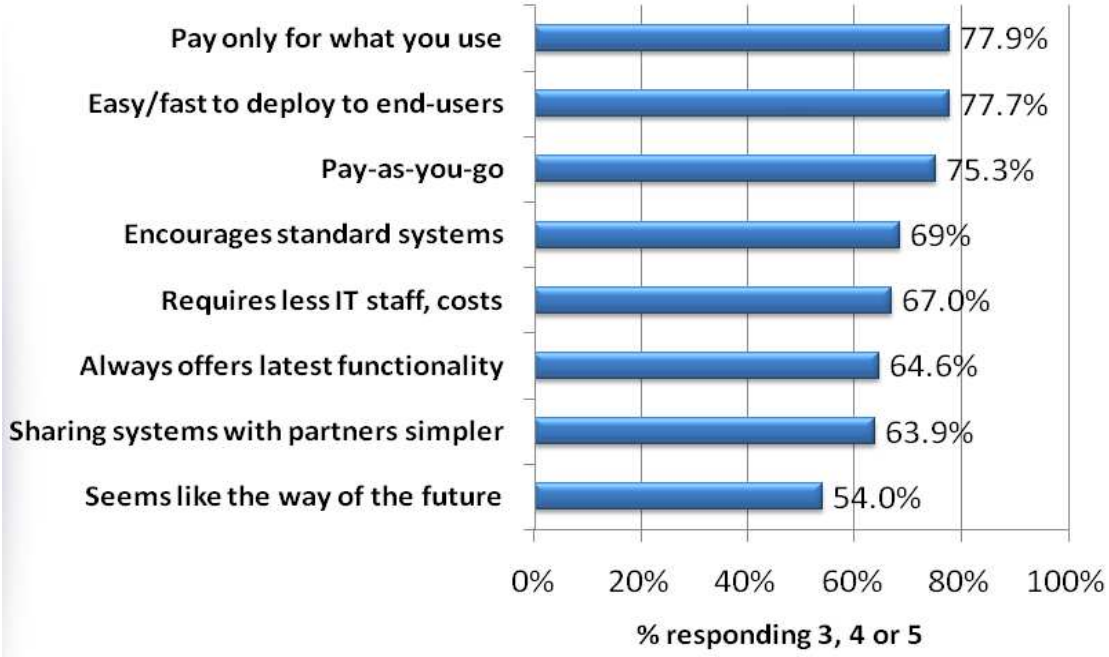

Source: IDC Enterprise Panel, 2010

Cloud computing has the potential to affect not only servers and operating systems, and the way they are deployed, but also the associated storage and services that complete the enterprise solutions. Data can be "parked" on a cloud-enabled storage service, improving high availability for production data that has been replicated (copied) to the cloud. Combined with high availability of applications, or

virtual machines containing applications, the presence of this data allows restart – and supports disaster recovery and business continuity in the event of outages.

That availability of cloud services means online business can proceed even when network outages or power outages or natural disasters occur in any given location or geographic region. When business continues, revenues continue – and customer satisfaction, via access to cloud-enabled services, is protected and retained.

Challenges

However, customers have some concerns about cloud computing, as well: Will it be enterprise-ready, with security, availability and quality of service (QoS), comparable to what is provided by traditional IT infrastructure? Will there be vendor lock-in, which prevents, or slows, applications or data moving from one cloud service to another? IDC demand-side research shows that customers are concerned about protecting their ability to change cloud providers and moving their applications and data from one cloud to another, via federation.

These concerns about cloud computing must be addressed by vendors across the board in the cloud computing environment. This means that hardware and software vendors, virtualization providers, along with services providers, must consider, and embrace, a common vision for how the individual components of cloud computing will be deployed together, either in an enterprise IT datacenter or in a cloud provider's datacenter.

No matter where they are deployed, there must be ways to glue these piece-parts together or be able to import and export the applications, virtual machines (VMs), and data that comprise an enterprise software solution. That approach will support IT flexibility and agility, and it will support the kinds of changing business environments typified by today's organizations — which may gain business units rapidly via merger and acquisitions (M&A) activity — or may divest units. Business initiatives also change, as do priorities for the business.

Solutions: Building a Portable Cloud Environment

Silos are the norm in IT today, existing on multiple levels. There are many dependencies on specific technologies, including dependencies on servers, network devices, storage devices, operating systems, middleware, or applications. For that reason, many types of data and applications are encased within traditional IT systems – and have poor mobility across the infrastructure. Encapsulating applications within virtual machines (VMs) that could be moved to other locations, for IT flexibility and HA/DR purposes, lends more IT flexibility (and business agility) to an organization.

Public cloud providers can be responsible for shielding end users from the complexity of and dependencies on the infrastructure they've put into place. But for private cloud deployments, IT organizations must be more aware of these complexities and dependencies, including those of various virtualization technologies that are the basis for the private cloud deployments. IT organizations will also have requirements for managing the life cycle of the software they are deploying across the cloud – even in instances that tap federated clouds to support an end-to-end solution. As applications age, they must be updated, and this must be done in a holistic way across all the components of that end-to-end application.

To the degree that many of these dependencies on infrastructure components are abstracted away from the end-user's viewpoint, then it follows that the datacenter, and its support for software solutions, will become more dynamic. The business result will be clouds that support change and portability at higher levels of abstraction – much closer to the applications and further away from the underlying infrastructure and its hardware specifications.

Key Interoperability and Portability Considerations

The following is a summary of key interoperability, portability, and flexibility issues across the stack that IT organizations should consider when building a portable cloud environment:

- **Support for multiple hypervisor technologies such as VMware, Hyper-V, KVM, and Xen.** The virtual machine (VM) is becoming a fundamental unit of work and encapsulation, particularly for IaaS and some PaaS clouds. Customers today are increasingly likely to have mixed virtualization environments inside their enterprises -- and to be able to pick up a VM and move it, regardless of the underlying platform is a primary first step to cloud portability.

- **Choice of operating system.** While application virtualization is emerging to break the bonds between applications and the operating system, most applications have hard dependencies on a particular version of a specific operating system. Over time, enterprises have built up a heterogeneous mix of operating systems (or operating environments) in their datacenters — which means that forcing any given operating system on customers in the cloud is highly impractical and would not be able to address a large number of workloads.

- **Programming models.** Most enterprises already have chosen their programming models a long time ago. While some of these programming tools will need to adapt to a cloud model, IT staffers will likely want to avoid a wholesale change, given their familiarity with a given programming model – and their skill-sets. However, new cloud programming models are certainly worthy of consideration. Clouds need to support the programming models that are in use today, both the workhorse enterprise models such as .NET and Java, as well as lighter weight Web models such as PHP and Ruby on Rails. Ideally, a cloud would support all these programming languages, but this isn't always going to be possible, which could result in sourcing from multiple clouds — all of which would place an even higher emphasis on hybrid cloud management.

- **APIs.** Platform as a Service clouds (programming to a given software platform) already provide new application frameworks and APIs that provide special cloud functionality, such as massive scale-out capability. However, there are many considerations when it comes to leveraging these APIs. Generally, these APIs are useful only to new applications. Existing applications may need to be modified to leverage these APIs, which is generally a difficult and expensive procedure. Using these APIs also may commit the application to a certain level of lock-in, because the API may not be accessible externally (from a new "target" cloud environment) and there may be no equivalent functionality in the new environment. One also needs to examine the complete API set to look not only at importing and migrating to the cloud, but also at moving out of the cloud later on. Cloud providers are eager to migrate customers onto their platform and readily provide tools to do so, but customers have voiced their concerns about the inconvenience of moving applications from cloud to cloud.

- **Unified management.** Every new technology seemingly has a dedicated management console, with public clouds being no exception. With the desire for tight coordination between on-premises and off-premises resources in a hybrid cloud, management must become more unified. Otherwise, a tangle of multiple management consoles will continue to co-exist, defeating the goal of simplification via cloud computing. This requires that clouds follow existing management standards — and that they also should be "open" to third-party management applications. If the APIs and management data are accessible, then that would allow unified cloud consoles, supporting plug-ins to uniformly control things such as access control, scheduling, resource management, and billing.

- **Data portability.** While considering the above issues of hypervisor and operating system choice and portable programming models, which mainly address application portability, the ability to move data also looms large as a management issue for cloud computing. The data the

application relies upon must be portable, and this reaches far beyond simply being able to copy files from location to location. Enterprise data may exist in structured (database) and unstructured (files) forms — and may be reached via various access methods. Data portability really means that when applications are moved, they will be able to find, and to access, the production data they need to operate upon. Further, simply moving data from one location to another will not solve the portability problem because applications require data that is formatted in specific ways. Customers will need to ensure that the cloud provides appropriate standards to support data export, along with data conversion from one format to another, as well as compatible or abstracted storage-access services.

- **License flexibility.** Software licensing has traditionally been static and often bound to locations and hardware. With virtualization and cloud opening up IT to a much more dynamic world, software licensing has not caught up. Customers need to be aware of the restrictions and problems that license portability may have on their use of cloud, and work with their ISVs to implement new policies that are cloud-friendly. Cloud portability, and the avoidance of lock-in regarding specific vendors' product sets, is an issue that is just starting to gain recognition. There is no simple fix to this problem; it spans the entire enterprise computing stack, which over time has become very heterogeneous. To further complicate the matter, enterprise IT has a wide web of dependencies that increases the complexity of the environment exponentially. Take a typical large multi-tier application: various parts could be virtualized or not, running on different operating systems, different security models, programming models, and management interfaces. Customers must become aware of the many areas of potential "lock-in" and take steps to mitigate their exposure. It will require due diligence at every step of the cloud migration.

Conclusion

In summary, IT has never completely solved the portability problem within the four walls of the enterprise datacenter. And now that cloud computing has arrived as a new option for IT, the evolving cloud computing ecosystem of ISVs, cloud providers, and customers must think about how they can support more portability in this new era. The open-source model, which supports collaborative development, provides a mechanism for improving the portability and interoperability aspects of cloud computing. Clearly, cloud migration will not be a "one-click" task anytime soon – and the discussion about portability and interoperability is still evolving, along with cloud computing adoption.

In coming years, the goal will be to provide reasonable options for migration and portability in cloud computing. With the wide range of cloud use that is expected to grow rapidly, expanding from $17 billion spent on IT infrastructure (servers, storage, and software) for cloud services worldwide, to more than $40 Billion in 2013, focusing on design goals that support portability of applications and data will certainly avoid expensive and time-consuming fixes in the future.