



HRG Insight: Cloud Computing

Keeping apprised of terminology in today's constantly changing IT landscape can be a fulltime job for IT decision-makers. Some terms lend themselves to a fairly educated guess as to their meaning. But "Cloud Computing"? It's no wonder even the most technology-savvy among us shy away from attempting to define the term directly. This HRG Insight provides guidance on what Cloud Computing means and more importantly, why you should care. It will also discuss how one vendor, Red Hat, has approached it.

What Problems are Being Solved?

One of the most vexing challenges facing IT today is how to optimize computing resources most effectively. How can you increase computing capabilities in "real-time"? When you need to suddenly increase computing power during an *unplanned* spike in business activity can you do it "on demand"? A sudden increase in complex financial transactions or a short but powerful need for increased computing performance or a sudden increase in web traffic are examples of an increase in load that may simply not be available "on premise". On the other hand, what impact does *over planning* for spike situations have on your costs for additional people, additional servers, software licenses, training, or networking capabilities?

The twin business challenges of having enough "flexibility" and "freedom" to optimize your enterprise's business may not always be met by a purely "in-house", on-premise solution. Suppose you could:

- Have the flexibility to just "plug-in" when you need to and receive the computing capabilities you need at the moment, not planned during a window between "11 PM and 1 AM".

- Have the flexibility to NOT use computing power during business slow-downs.

- Have the freedom from costly software operating system and application upgrades/revisions, IT management redundancies, and database administration headaches with resulting unplanned downtime.

- Have the freedom from costly infrastructure start-up costs, particularly if you're a start-up yourself.

- Have the flexibility and freedom to be able to scale up or down a computing infrastructure when you need computing power most.

Cloud Computing is another 21st century technology term that today's CIO and IT decision-makers are struggling to understand. The benefits of being able to utilize resources when you want them, on your terms while optimizing IT infrastructure and application lifecycle costs are too good to dismiss lightly. Red Hat, a pioneer in open source solutions, partnering with Amazon, a pioneer in deploying Cloud Computing have developed an approach which makes the journey smoother for IT management.

You'd be able to sleep better at night if these challenges had a more economical solution than adding more resources, managing and administrating more vendors products while still not optimizing resources.

That's where Cloud Computing comes in.

What is Cloud Computing?

Cloud Computing can mean many things to many people in IT. There has been a significant reticence in business and IT environments to define it explicitly with some IT technologists simply smiling when asked to define it. Is it simply another name for "On-demand" computing? How about "Utility Computing", akin to plugging into an infrastructure and getting a service instead of a stream of electrons? Is it another name for SaaS (Software as a Service)? How about Web Services? Some technologists classify Utility Computing, SaaS, On-demand and Web Services as forms of Cloud Computing. In some respects, Cloud Computing encompasses all of these terms – and more.

Cloud computing can be deployed within a single corporation or as a service offered by a 3rd party to their customers. Cloud computing provides compute resources to internal customers and can be used to dynamically reallocate compute resources to meet demand on an as needed / when needed basis from existing resources. Primarily a technical architecture, cloud computing lets organizations more effectively leverage capital expenses.

Figure 1 depicts the general architectural concepts of Cloud Computing. Local servers are owned and managed by their respective companies (company A and company B). As more computing resources are needed, both company A and company B can take advantage of the servers in the cloud, which is owned and maintained by company C. Administration functions control when and what applications are executed by the cloud servers.

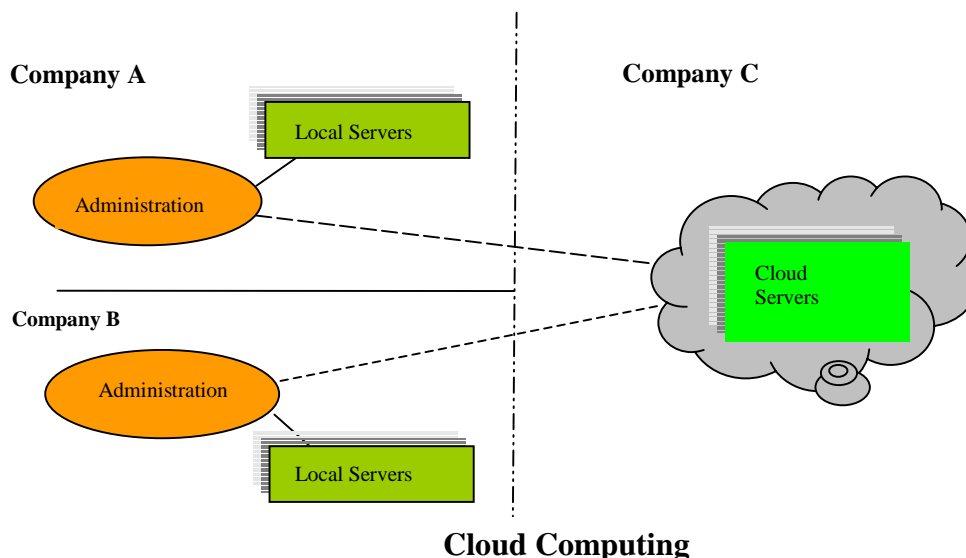


Figure 1

When offered by a services provider, Cloud Computing generally means getting computing resources and services from a remote location over the Web when you want them, for as long as you need them, sometimes for a single or unique application such as a CRM or ERP application. In this regard, Cloud Computing is similar to Web Hosting or outsourcing specific applications or functions. However, there are some important differences, at least regarding the current deployments of Cloud Computing. Most of these differences focus on extending scalability, availability of resources, or scope of applications that can be deployed.

- Size and scalability – perhaps the most obvious element that sets Cloud Computing apart is the potential to utilize far greater computer resources than a simple Web hosting company can provide. Recently companies such as Amazon, Google and Yahoo have begun offering Cloud Computing services in part because this strategy is a logical extension of their own data center infrastructures. These companies already have enormous computing power in order to conduct their web-based search and transaction businesses. Why not extend these resources to other enterprises who would like to “tap into” a cloud? Cloud Computing is set apart from general Web services because it can offer Web services that need significant computing power as well as data-intensive applications such as BI (Business Intelligence) and Data Warehousing.
- Simplified business arrangements – Cloud Computing also changes the business models companies use to pay for IT. Traditionally, companies pay acquisition license fees, fees for upgrades, fees for new versions, fees for support. Often there is little correlation with actual computing usage. Some budgets are in effect left unused. Cloud Computing posits a business model that allows a user to essentially pay for the resources and capacity actually used with more flexibility to scale capacity up or down depending on business need.
- Use of virtualization – Cloud Computing often uses virtualization (see HRG Insight Paper “Virtualization”) of servers as the key component of the service allowing the Cloud Computing provider to offer more flexibility and optimized use of the Cloud Computing infrastructure. Virtualization allows running multiple images of an operating system and application infrastructure on the same physical server, providing a more optimal use of overall computing resources.
- Utility Computing parallels – some have dubbed Cloud Computing as “utility computing that works” or utility computing 2.0. Like initial deployments of utility computing, Cloud Computing offers the benefits of real-time, “on-demand” access to resources. However, due to virtualization, in particular, Cloud Computing overcomes many of Utility Computing’s shortcomings. Resource provisioning and overall management of resources and time are made much simpler.
- Flexibility in applications – unlike some SaaS solutions, Cloud Computing has the potential to offer a range of solutions from one, compute-intensive CRM application, for example, to running several general mundane, less mission-critical applications, thereby allowing an enterprise to run only those applications it chooses in-house.

Cloud Computing therefore encompasses such terms as Utility Computing, On-demand, SaaS, and Web Services. In some respects these terms are predecessor architectures for Cloud Computing. Generally, SaaS solutions provide a single application through Web browsers to many customers, such as “SalesForce.com”. Utility Computing has been used in the industry for some time but never on a scale that Cloud Computing can offer. On-demand has been used as a marketing term as well as a technology trend.

There is no question that Cloud Computing is not just the trend “du jour”. It offers a direction to enterprises that may wish to explore alternatives to acquiring or leasing their own “on-premise” computing resources (hardware, software, etc.) or committing to a hosting contract.

Cloud Computing challenges - Some enterprises may not want to “give up” control of resources, particularly for mission-critical applications or for applications that use proprietary or trademarked data. As with any hosting arrangement, some organizations feel that a loss of direct control may also mean a loss of data or knowledge security. Security and access policies of cloud computing providers must be well understood to address these concerns. Some organizations may move more slowly in adopting new trends for mainstream applications. This is why Cloud Computing has initially caught on primarily with smaller enterprises, particularly startups who want to avoid the potentially significant IT infrastructure costs.

In any event, Cloud Computing will only continue as a trend based on how well it is built on open standards, such as the use of Web Services standards that enable APIs to integrate with application services. Further, how well the “clouds” within these mega-datacenters can be securely integrated when needed is important to ensure the future success of Cloud Computing.

One company that has built its reputation and loyalty on open software and service is Red Hat. How have they risen to the challenge of Cloud Computing and what does it mean for IT decision-makers?

Red Hat and Cloud Computing

Red Hat is aggressively looking to the future of Cloud Computing by building on its foundation of open source software and offering an integrated Cloud Computing solution directly aimed at the commercial enterprise. It is deploying this solution by partnering with Amazon, one of the pioneers in Cloud Computing. Red Hat Enterprise Linux (RHEL) is the first commercially supported operating system available on Amazon’s Cloud Computing service dubbed “EC2” for “Amazon Elastic Compute Cloud”.

Figure 2 depicts a cloud computing environment using a Red Hat Enterprise Linux and Amazon Elastic Compute Cloud solution. Company A still has its local servers, but in addition it has access to a cloud of servers that it does not own or maintain. However, Company A can control where their applications run (locally or in the cloud) and control what resources are required. Web service interfaces are used to requisition machines in the cloud and load them with the appropriate application environment.

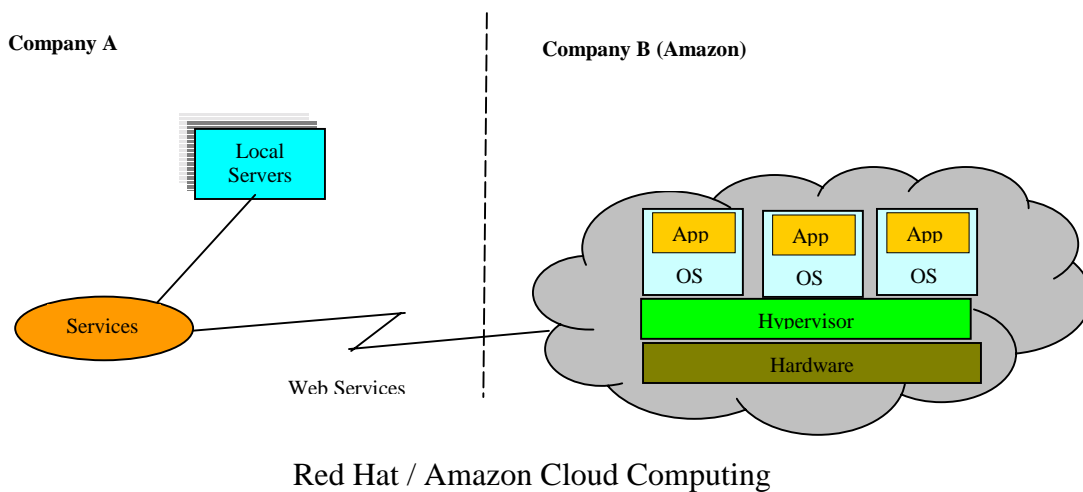


Figure 2

The overall service provides an enterprise with an open-standards, web-based, computing environment hosted within Amazon’s data center infrastructure.

- Benefits - the solution allows a user to take advantage of Cloud Computing’s benefits of using compute capacity, network bandwidth, storage and application availability “anytime, anywhere”, when the enterprise needs these resources.
- Foundation – The foundation of the service is built on RHEL, the leading open source operating system. By having a identical development and deployment environment both on-site and in the cloud, customers are assured that their applications will perform consistently regardless of their deployment.

- Virtualization - The solution also uses virtualization within the Amazon data centers assuring optimized platform and server use of resources and higher overall application availability. When a new resource is allocated within the cloud, the resulting “instance” is an official release of Red Hat Enterprise Linux and users can expect the same capabilities as they receive in their on-premise deployments of Enterprise Linux.
- Automated Linux - Red Hat’s solution provides a common set of management and automation tools integrated with on-premises clients and the Amazon Cloud Computing environment. Enterprises can utilize these tools to start or stop servers, perform software services, data transfers, management and configuration tasks as necessary.
- Application Optimization – These common management tools permit an enterprise to optimize the performance, availability and delivery of applications. They allow enterprises an easier solution for developing, deploying and managing applications. Over 3000 applications have been certified to run on RHEL.
- Control – Overall, this solution allows an enterprise more dynamic control of computing resources rather than contracting with a traditional, off-site hosting provider.
- Flexibility - True to the spirit of flexibility, the Cloud Computing resources can be used on a hourly basis allowing enterprises in order to minimize expenses and pay only for what they consume.

Beyond this partnership with Amazon, Red Hat’s approach to Cloud Computing, as is its approach to “virtualization”, is based on its overarching strategy called “Linux Automation”. Regardless of the computing technology or the terminology used, enterprises today are struggling to optimize overall IT investments, consolidate resources and increasingly control application lifecycle costs. Whether an enterprise chooses to deploy virtual servers “on-premises” or use the “on-demand” advantages of Cloud Computing, there needs to be an application foundation that allows an IT manager to “certify once, deploy many times” in an open software environment. Red Hat’s approach to Linux Automation provides a comprehensive set of operating system capabilities, application, and management tooling to help businesses easily implement and administer complex environments across varying compute infrastructure architectures.

Just as virtualization is a concept to “hide” the physical elements of the technology and “abstract” them into a virtual image, Linux Automation is Red Hat’s strategy to “isolate” the complexities of these technologies and extend an enterprise’s open environment investments. Linux Automation solves the key challenges of reducing overall application lifecycle costs. It ensures that if Cloud Computing is chosen, there will be an integrated foundation of common software and tools which can be used “on-premises” as well.

Summary

Cloud Computing is more than just the technology term du jour. It encompasses the basic concepts of On-demand, SaaS, and Utility Computing. However defined, CIOs and IT managers should pay attention and leverage benefits of being able to utilize resources when and where needed. Red Hat, a pioneer in open source solutions, partnering with Amazon, a pioneer in deploying Cloud Computing, is an approach which can only help make the journey much smoother for IT management in the future.

Harvard Research Group is an information technology market research and consulting company. The company provides highly focused market research and consulting services to vendors and users of computer hardware, software, and services. For more information please contact Harvard Research Group:

Harvard Research Group™

PO Box 297
Harvard, MA 01451 USA

Tel. (978) 456-3939

Tel. (978) 263-3399

e-mail: hrg@hrgresearch.com

<http://www.hrgresearch.com>