



General Physics Corporation
Leading the World to Better Performance

What the Emergence of Cloud Computing Means for Your Learning Technology Deployment

White Paper

by

Billy Biggs, MBA, PMP
Director, Learning Strategies

General Physics Corporation

www.gpworldwide.com

Behind the Firewall, Externally Hosted, Software as a Service and Now Cloud Computing— Welcome to Delivery Model Madness!!

The technology-enabled business services industry is obsessed with buzz words. In reality, this is nothing new for organizations that provide or consume these services, as it is considered part of daily operations. If you're involved in technology in the Learning and Development (L&D) space, terms like Web 2.0, Social Media, Informal Learning, Portal, Personalization, Virtualization and Software as a Service (SaaS) have now become part of your strategic L&D roadmap discussions. You could say that as soon as organizations align technology with business objectives, the market presents alternative mechanisms to accomplish these objectives better and more efficiently. Case in point—cloud computing.

As is the case with other technology-related jargon, it may be difficult to come up with a clear, precise definition for cloud computing. Actually applying the description to real-world examples may be even more challenging.

“Cloud computing is a style of computing in which dynamically scalable and often virtualized resources are provided as a service over the Internet.”ⁱ

Does that help? No, then how about...

“The concept, quite simply, is that vast computing resources will reside somewhere out there in the ether (rather than in your computer room), and we'll connect to them and use them as needed.”ⁱⁱ

Or...

“Cloud computing basically lets you access software applications, hardware, data and computer processing power over the web. You don't have to purchase and run software in your computer anymore. The cloud, the abstract physical description of what the Internet can do (personal computer, processor, storage), is there for you anytime, anywhere.”ⁱⁱⁱ

Perhaps, most simply:

...common business applications, platforms or infrastructure that are hosted on the Internet (that is, in the cloud) and are accessed locally from a web browser, while the software and data are stored remotely on servers.^{iv}

What the Emergence of Cloud Computing Means for Your Learning Technology Deployment

Before we begin to take a deeper dive into cloud computing and how it relates to learning technology system deployments, let us take a look at two of the more traditional deployment models for comparison purposes: “behind the firewall” (BFW) and externally hosted by a third-party hosting provider.

Behind the Firewall

The BFW option is one of the earlier models that most of you are probably aware of. Essentially, an organization makes a capital investment in enterprise software and takes full responsibility for the installation, configuration and hosting of the application, utilizing internal resources, equipment and infrastructure. The initial investment is usually very high and the organization is required to pay annual “maintenance” that covers software fixes, access to new functionality (upgrades) and technical support.

The main advantage is the ability to customize the application as necessary through code modifications or other application integrations to achieve business objectives. With the ability to customize, organizations find little reason to modify or alter business processes to work with the application’s functionality constraints. However, customizing enterprise software like a Learning Management System (LMS) always involves risk. It also involves additional costs when upgrading, as those customizations must be applied to newer versions of the application. An often overlooked negative side effect of customization is that the more different the buyer’s system is from the base provider’s platform, the farther outside of the comfort zone and expertise of both the provider and its general customer community.^v One of the major disadvantages in the BFW approach is the obvious need for internal IT support and infrastructure, as well as the need to manage other priorities those resources have to balance (that is, other enterprise systems).

Externally Hosted – Third Party

This deployment method is a lot like BFW from a software-licensing standpoint. However, instead of relying on internal IT resources for the deployment of the application, organizations choose a third-party vendor that provides an outsourced solution. The provider is completely responsible for the installation, configuration and hosting of the application, utilizing its own resources, equipment and infrastructure. Costs are generally lower than a BFW deployment, as most outsourcing solutions are, yet organizations still maintain the flexibility to customize the application, if needed. However, just like in BFW, there are fewer concerns about customizations, security and privacy than in the newer deployment models.

Imagine you are in the market for a new car. A traditional deployment model of BFW or hosted by a third party is like buying a premium sedan at a fixed price. It has all the bells and whistles you desire, including maintenance support and an extended warranty. It likely also has additional features that you do not plan on using, such as satellite radio, IPOD auxiliary inputs or real-time traffic navigation. However, you will still be paying a premium price for a vehicle that has these features. On the other side, you have the ability to customize it by tinting the windows, installing an after-market stereo system or making other modifications. The car is yours until you sell it or trade it in for a new one.

Now, consider a corporate LMS. With an enterprise deployment, externally hosted or BFW, the LMS software is yours. You own it and automatically gain access to updated versions of the application since you pay the annual software maintenance. However, you will still need to pay for the professional services labor of the upgrade, whether the cost is realized by internal IT resources or whether the third-party hosting vendor is billing for such services.

Clearly, there are both advantages and disadvantages to conventional hosting options. But an overriding reason to migrate from these deployment models to SaaS and/or cloud computing is simple—cost.

Software as a Service

A few years ago, when applications like LMSs and Learning Content Management Systems took on a key role in Human Capital Management initiatives, BFW and third-party hosting were considered the only delivery options. While these models represented most of the market share, Application Server Providers (ASPs) became the next “hot” deployment model as vendors looked towards innovation. Ultimately, ASPs were considered a failure, but the lessons learned provided the advent for a new delivery option—SaaS.

SaaS has quickly been defined as a competitive business strategy for organizations looking to eliminate unnecessary costs while choosing a system that meets their most important enterprise business requirements. Unlike BFW or third-party deployments, you do not own the application software, only the data that resides in the application. SaaS providers make their applications available “on demand” on a subscription basis. Implementation costs and timelines are usually a fraction of a traditional implementation since these applications are deployed with a standard baseline of code. Your organization usually pays a per-user fee for access to the application, which generally includes quarterly “releases.” In the SaaS world, releases are often additional functionality and/or upgrades. Previously, the term *release* has been used to try to differentiate it from the word *upgrade* because the main premise behind the SaaS business model is that the software is essentially “version-less.” However, this often creates more confusion for customers, so more SaaS providers have reverted to using the phrase *latest functionality upgrades*, as shown in Figure 1.

Latest functionality with Seamless Opt-in Upgrades. The upgrade opt-in window offers you the flexibility to decide when you want to receive the upgrades and implement the newest features and functionality. By taking advantage of the opt-in window, you now have the time to plan for any potential change management issues associated with new releases.

Figure 1. Marketing Jargon Among SaaS Providers Creates Confusion Regarding Services

While this feature of SaaS is very attractive for some industries, it can be a game-stopper for others. The federal government, for example, will have an extremely challenging journey to a true SaaS business model, regardless of the application, due to its complex governance models, user acceptance and regression testing requirements and highly regulated security infrastructure.

A true SaaS model, however, operates on the premise that there are no customizations to the application. That usually means organizations must undergo a significant training process reengineering effort in order to align their business processes with out-of-the box functionality of the application. To complicate matters even further, it is important to understand that, like most technologies, each SaaS supplier is likely operating at a certain level of software maturity.

Not surprisingly, not all SaaS providers are offering the same type of deployment from a maturity standpoint. Think of it as going to a steak dinner. While you may be able to get decent cut of beef at Cactus Willie’s Steak Buffet, there is clearly a big difference in the quality of beef you would get at the Capital Grille. The most widely adopted maturity model for SaaS was published by the IT research firm IDC in 2005^{vi}.

- **Level 1: Ad-Hoc/Custom** – Each customer has its own customized version of the hosted application and runs its own instance of the application on the host’s servers.
- **Level 2: Configurable** – Many customers use separate, configurable instances of the same application code.
- **Level 3: Configurable, Multitenant-Efficient** – The third maturity level adds multitenancy to the second level so that a single program instance serves all customers.
- **Level 4: Scalable, Configurable, Multitenant-Efficient** – Technologies (such as load balancing, multitiered architectures and virtualization) increase the scalability of the platform.

Referring to the earlier example of purchasing a car, think of SaaS as leasing that new car instead of purchasing it outright. In the end, you get a nice vehicle but you can’t make any significant changes, as it does not belong to you. You are required to make a monthly lease payment to guarantee a certain level of service. At the end of the lease period, you can continue the lease, lease another new car or buy a car. With a SaaS solution, you will always be faced with a monthly line item on the vendor’s invoice. While the flat line investment model is attractive for some organizations, on paper it does not make financial sense after a few years.

Need help determining your supplier’s SaaS maturity? Bersin and Associates’ “Adaptive Talent and Learning Platforms” report provides some great insight to the maturity of current Learning and Talent Management SaaS solutions. Figure 2 shows levels of maturity around delivery only for SaaS providers. However, it’s important to understand that both internal IT departments and third-party hosting vendors may maintain all the technical expertise to provide all of the services listed by leveraging virtualization and other technologies to drastically reduce costs and manage utilization fluctuations.

	Level 1: Hosted	Level 2: Configurable	Level 3: Adaptive System	Level 4: Adaptive Platform
Delivery	24/7 hosted access to application	All clients share a common “immutable” code-base; Built-in relationships with content distribution networks (e.g., Akamai).	Multiple clients share an instance of the application (multitenant); Application is broken into tiers to better balance hardware to traffic.	Architecture is “efficient multitenant”; One or more of the following may be used to increase efficiency and scalability: virtualization, globally distributed infrastructure or “cloud” computing.

Source: Bersin & Associates, 2009

Figure 2. SaaS Providers and Levels of Maturity

Cloud Computing

If you have attended any relevant IT business conferences in the last two years, you have certainly heard, discussed and questioned the value proposition of cloud computing. Quite simply, organizations are now looking for the convenience and simplicity of SaaS but with the flexibility of traditional deployment models. In theory, the cloud is a computing service that only charges you based on the amount of resources you use (that is, pay as you go). Pricing for computing services can vary by provider, but imagine that you (in the Talent Management world) are only being charged based on some utilization factor. This could be the number of course completions recorded in your LMS, the amount of bandwidth used or the number of performance reviews completed in the system during a specified timeframe. This pricing feature sets it apart from traditional SaaS by correlating your costs with your actual system and computing usage. Cloud computing has been discussed as the convergence of three major trends:

1. Virtualization – Applications are separated from infrastructure.
2. Utility Computing – Server capacity is accessed across a grid as a variably priced service.
3. SaaS – Applications are available on demand on a subscription basis.

The cloud is like having a taxi at your disposal, whenever you need it. You only pay for it when you use it, without being responsible for any maintenance fees. Some months, you may need the service of the taxi several times a day; others, you may not need it at all. As a result, cloud computing is especially attractive for organizations that leverage applications like the LMS as a business model. Revenue is more likely to increase and decrease with cost, leaving operating margins relatively flat.

Cloud Offerings

Cloud computing models can be divided into three categories: SaaS (or application), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). Michael Sheehan from the *Cloud Computing Journal* illustrates the cloud offerings in a pyramid format, as shown in Figure 3. When considering the “cloud pyramid,” it may be easier to conceptualize the major components of cloud services available to consumers.

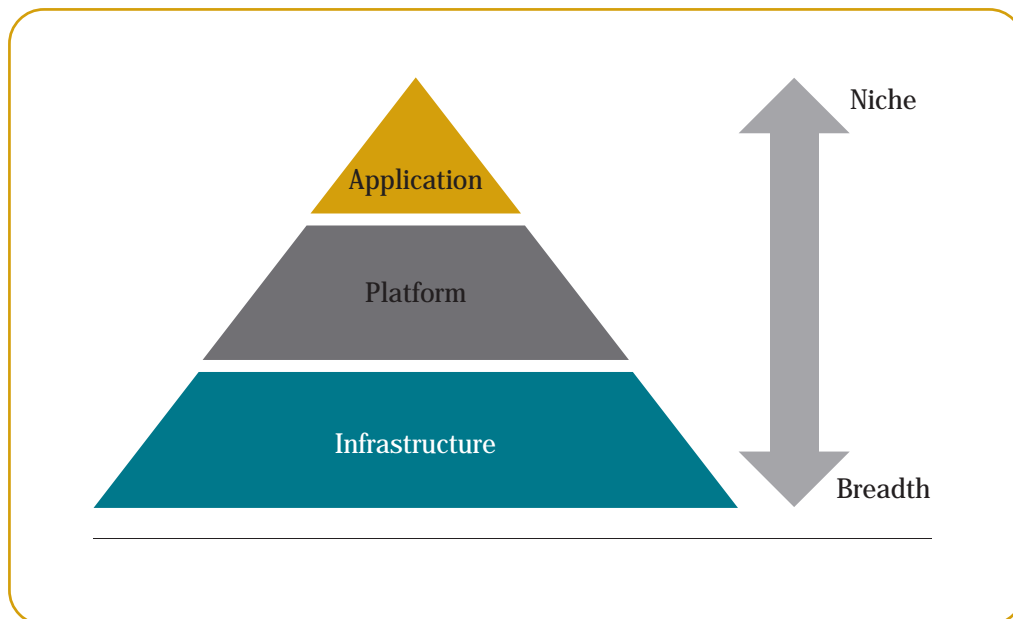


Figure 3. Cloud Pyramid

Application/SaaS

With the SaaS model, the capability provided to the consumer is the use of the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (for example, web-based email). The consumer does not manage or control the underlying cloud infrastructure such as the network, servers, operating systems, storage or even individual application capabilities (with the possible exception of limited user-specific application configuration settings).

Within this most narrow part of the pyramid, users are truly restricted to only what the application is and can do. Some of the notable companies providing this type of offering service are the public email providers such as Gmail, Hotmail, and Quicken Online. Almost any SaaS provider can be lumped into this group. Most retail consumers use the services within this cloud category, which allows for predefined functionality built into the product without much further adaptability. These applications are designed for ease of use and GTD (getting things done). For example, Salesforce, a huge cloud application/SaaS provider that has led the way for hosted software, falls into this category as well. However, their force.com product does not. Even some of the online banking offerings could fall in this group. In addition, all Talent Management solution providers currently fall into this category. The SaaS Talent Management applications are made available to users who can add and manage content within the application, but with no ability to modify or customize the product.

*Application/SaaS examples:
Mapquest, Gmail, Salesforce.com, Zoho*

Platform/PaaS

The capability provided to the consumer with a PaaS model is to deploy consumer-created or acquired applications developed (using programming languages and tools) and supported by the provider onto the provider's cloud infrastructure. The consumer does not manage or control the underlying cloud infrastructure, such as network, servers, operating systems or storage, but has control over the deployed applications and possibly application hosting environment configurations.

As you move further down the pyramid, you gain increased flexibility and control over the application's configurations and functionality, but you are still fairly restricted to what you can and cannot do with the overarching architecture and environment. Products and companies like Google App Engine, Heroku, Mosso, Engine Yard, Joyent or force.com (SalesForce platform) fall into this segment. This category is becoming more congested with competitors, many of whom are trying to leverage the cloud infrastructure. As David Mallon of Bersin & Associates suggests, "PaaS is not likely to directly affect the learning or HR application markets anytime soon, but you will see its influence as providers innovate to bring additional flexibility and extensibility to their products."

*Platform/PaaS examples: Facebook,
Google Apps, Windows Azure, Appexchange*

Infrastructure/IaaS

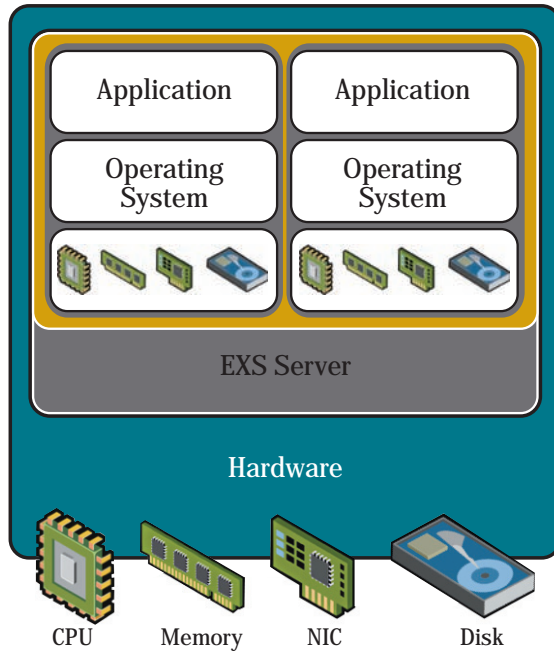
With an IaaS offering, the consumer has the ability to provision processing, storage, networks and other fundamental computing resources that allow him/her to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications and, possibly, limited control of select networking components (for example, host firewalls).

Examples of infrastructure providers that fall in this bottom sector of the pyramid are Amazon's EC2, Google, GoGrid, RightScale and Linode. Companies providing infrastructure enable cloud platforms and applications. Most companies within this segment operate their own infrastructure, allowing them to provide more features, services and control than others that are not part of this cloud service. At this foundation level, GoGrid is able to offer infrastructure in the form of both Linux and Windows, load-balancing and storage. Some Infrastructure providers may leverage other vendors within the space in order to provide competitive viability as well. For example, if vendors use the full capacity of their private clouds, they have the ability to tap into computing resources in public clouds for additional capacity. This example of technology is not considered mainstream...yet. Note that some of the industry-leading third-party hosting providers fall into this category.

*Infrastructure/IaaS examples:
Amazon Web Services, GoGrid, Rackspace*

How Virtualization Has Paved the Way

So what's the first step to cloud computing? Virtualization. Virtualization is the engine that powers the cloud computing model by adapting applications to run as virtualized images or machines. Cloud computing offerings would not be possible in today's operating environments without the concept of virtualization. Suppliers can operate in the SaaS framework without virtualization.



However, without virtualization, they cannot adapt to heavy increases in system usage such as mandatory training periods. For instance, in the federal government training space, it is not uncommon to operate at 3 percent user concurrency from November to August on an annual basis. However, when mandatory training is assigned every fiscal year, concurrency in the months of October (those who wish to take training at the beginning of the year to get it “out of the way”) and September (those who have procrastinated all year long and wait until the last minute) is sometimes as high as 25 percent.

In a virtualized environment, applications and infrastructure are independent from one another. Servers can now share many applications, and applications can run virtually anywhere. Virtualizing the application becomes an important first step in this new breed of technology. Essentially, this means “packaging” the application with everything it needs to run (databases, middleware, operating systems) to become a self-contained unit. As a result, it can run anywhere, including “the cloud.” An example of this concept can be illustrated with a common problem faced by Mac owners a few years ago. There were obvious programs that could not be run on a Mac (for example, MS Office before Apple branded a version), but were necessary to its owners. What allowed these users to work around this problem were programs like VMWare Fusion, which enable users to install a virtual machine (VM) of Windows on a Mac while still running the Mac OS X.

This type of virtualization ultimately allowed users to leverage the digital lifestyle of their Mac while having one or two Windows programs that they couldn't live without.

Now imagine an enterprise Learning or Talent Management System. Virtualization software, such as VMware ESX Server, can be used to virtualize server storage and networking, allowing multiple applications to run on virtual machines on the same physical server. No longer must the software vendor build another instance of the application when requested or necessary due to increased utilization. As demand for system resources increase, a new virtual machine can be provisioned as needed without an up-front hardware purchase. Because of the ease of relocation, virtual machines can be used in disaster recovery scenarios as well. As a result to these benefits, virtualization is one possible solution to managing the data centers of the future.

Where's the Support?

There is no question that both SaaS and cloud solutions become lower-cost options for Talent Management leaders who are faced with diminishing budgets. Yet, there seems to be another key difference when comparing these solutions to the more traditional delivery models where the application is hosted internally or outsourced. That difference is support.

In an SaaS implementation, most companies can find that the initial implementation costs are usually much lower than in traditional implementations of the past. However, Talent Management leaders should recognize that SaaS providers have opted for a more “self-service” model when it comes to the configuration and management of the application. Essentially, in many SaaS scenarios, the provider grants the client access to the application and leaves the rest in the consumer's hands. This is a drastic difference from the more traditional implementations, where the provider/integrator helps you flush out your business processes, assess them against the application's functionality and helps you configure and manage the system in a way that is most optimal to meet your organization's training needs. **These current SaaS implementations can leave you, as the system owner, scrambling to find appropriate support to work out the intricacies of how your business model should integrate with your Talent Management system, as well as to receive the much needed consulting for ongoing operations.** This, in turn, can dramatically increase your overall costs that were unforeseen—and therefore unplanned for—when you initially implement the SaaS deployment model.

Summary

In the uncertainty of accelerated technology, there are many questions to ask solution providers regarding SaaS and Cloud solutions. The greatest impacts revolve around the configurability, or lack thereof, of the application; the visual interface; data integration; delivery; and support. Security requirements will often dictate what an organization can and cannot do from a delivery model standpoint. Be certain you're not losing functionality, the ability to configure the system or overall flexibility by choosing to go to a SaaS/cloud delivery model.

In addition, remember the role that virtualization plays in this puzzle. Just because a vendor offers a SaaS solution does not mean they have the ability or expertise to run your learning applications on virtual machines that would provide added flexibility and cost savings. Also, it should be noted that there are plenty of traditional third-party hosting vendors that have the ability to run your existing application in a virtualized environment, giving you the same amount of flexibility that you've always had, yet reducing costs considerably. Lastly, remember to be careful about selecting solutions only due to lower cost, especially if what you have now meets all or most of your needs. Quality of service, expertise and a successful track record that you may have experienced firsthand should be weighed heavily in your considerations. We've all heard "you get what you pay for," and certainly technology has a reputation for being no exception.

References:

- i — http://en.wikipedia.org/wiki/Cloud_computing
- ii — http://technology.timesonline.co.uk/tol/news/tech_and_web/article3874599.ece
- iii — <http://www.smartschools.ph/SmartSchools/SmartTools/UpsDownsCloudComputing.htm>
- iv — Michael Bromley, Head of Online Services at Telstra Business
- v — "Adaptive Talent and Learning Platforms: How SaaS is Changing Talent and Learning Systems," Bersin & Associates, February 2010.
- vi — Erin Traudt and Amy Konary, *2005 Software as a Service Taxonomy and Research Guide*, IDC, June 2005.

About the Author

Billy Biggs has been in the technology-enabled business services space since 1996. Mr. Biggs serves as Director of Learning Strategies at General Physics Corporation, a global training and development firm specializing in providing human capital management, strategic technology consulting, managed services, hosting support and integrated project management solutions. He works with company executives providing strategic growth direction and business development opportunities. Mr. Biggs has extensive experience in the Talent Management industry, providing expertise to the federal government space and dozens of Fortune 100 clients. Previously, he held numerous program and project management positions focused on enterprise technology supporting Fortune 100 clients, federal civilian agencies, the Defense Information Systems Agency (DISA), and the Department of Defense. He has an MBA, a BS in Information Systems Management and is also certified by the Project Management Institute as a Project Management Professional (PMP). He presents at several conferences a year on topics such as Integration of Enterprise Systems, Integrated Project Management, the Impact of Social Applications on Learning and eLearning consulting best practices. ■



Training | Consulting | Engineering

Author Contact Information

For more information, please contact
Billy Biggs, MBA, PMP
Director, Learning Strategies
General Physics Corporation
bbiggs@gpworldwide.com

General Physics Corporation
6095 Marshalee Drive, Suite 300, Elkridge, MD 21075
888.843.4784
www.gpworldwide.com