

Accelerating Time to Market:

Application Development and Test in the Cloud

Paul Speciale, Savvis Symphony Product Marketing


June 2010

Software Development and Time to Market

Software development has become a business critical, yet highly resource and budget intensive endeavor for corporations in a wide range of industries. For many corporations, software is a key enabler of their business processes, and the availability and stability of that software has direct impact on corporate revenues and customer satisfaction. A major transition is also occurring in the world of software development, as traditional Independent Software Vendors (ISV's) transform themselves into Software-as-a-Service (SaaS) vendors to deliver applications in the new streamlined world of cloud computing. For example, the delivery of consumer oriented financial, tax and investment software has become an extremely prevalent SaaS solution, as applications that were previously only available as host or desktop based software, are now delivered as secure online application services to thousands of concurrent users. For SaaS vendors, a major driver of their business success hinges on speeding time to market for these key applications, as any additional months and quarters of providing online services directly accelerates revenue, and therefore impacts their financial bottom line. In addition, the SaaS delivery model itself aids companies in accelerating their time-to-market, as the vehicle for application delivery is directly over the Internet, rather than traditional style packaged distribution, or downloadable modules.

Cloud Benefits for Application Development

- Agile, on-demand provisioning of all data center resources in under an hour, through self-service access
- Multiple service levels to address the wide ranging requirements of the application lifecycle
- Fast scale-up or scale-down of resources, with usage-based billing
- Support for virtual environments and multi-tier application architectures
- Migrate existing virtual server images and workloads into the cloud



SaaS Application Delivery Model

In many vertical industries, delivering in-house applications in a timely manner can also impact the business dramatically. For manufacturing companies, new software development can involve the delivery of a critical design or simulations application for new products that need to be brought to market. In media and entertainment companies, creating a new brand or event specific web site can accelerate end-user traffic, thereby driving advertising

which can have material impact on revenues. In the financial services industry, the need to provide new online capabilities for trading, portfolio and wealth management impacts customer attraction and retention. Clearly, the speed at which software applications can be developed, tested and brought into production is a critical factor for thousands of corporations worldwide. For many such corporations, any solution that accelerates the time to market for key software applications will have an immediate and measurable return on investment.

An emerging class of cloud services delivering Infrastructure-as-a-Service (IaaS) stands to become a critical enabler of rapid application development, as many SaaS vendors and in-house development organizations will choose to employ them for application development and test, as well as production deployments.

Bottlenecks in Software Development

Application developers work in highly dynamic environments, whereas IT administrators prefer to create development environments that are more static development static in nature, to provide much needed stability. These opposing requirements have been at odds for years, and provide a key insight into the lack of software development productivity.

For application developers, there is a regular need to build up new environments to design, write or debug code and compile software. This requires specific environments consisting of the proper libraries and tools. In many cases, these environments are transient, where they are utilized for several hours, and then can then later be eliminated. A key to developer productivity is access to these properly configured compute and storage resources in an on-demand manner. Users need to rapidly provision servers, with underlying storage capacity and fast network access in order to perform their tasks. Unfortunately, in many environments the speed of provisioning is an inhibitor to these tasks, as provisioning of these resources is controlled by IT, or the number of resources a developer has access to have been limited to a certain level. This forces developers to do development steps in a serial fashion, and precludes them from performing tasks in parallel to speed the development process. The inability of developers to rapidly provision, de-provision, or copy (clone) properly configured development environments are therefore several inhibitors to productivity, and hence time to market.

Development also involves iterative testing, which creates a need for additional data center resources above and beyond that which has been dedicated to development. In many cases, test environments are extremely dynamic: they are created to run a series of test cases and to review and analyze results, but then become dormant until the next test iteration. This again points to the need for a highly agile IT capability, where test environments can be provisioned on-demand (ideally by the person assigned to perform the testing).

Finally, the Quality Assurance (QA) process has its own requirements which bring agility to bear. On top of running functional test cases, and regression sequences, QA typically bears the additional burden of performance and scalability testing. These last two requirements make QA environments unique when contrasted with the developer's iterative code/test cycle. To address performance, the QA environment needs to provision data center resources that perform at production levels. This includes fast networks, with production level QoS guarantees, and fast storage systems, with the types of throughputs and latencies that will be seen in production. Servers may also be provisioned with higher levels of CPU and RAM resources in production than are used by developers. In short, performance testing requires the QA environment to largely mirror the production environment. This is often very difficult to attain in many corporations, due to the complexity of provisioning an environment that mirrors the production environment, even on a smaller scale. Even if the resources are available, the time involved in establishing a QA configuration purely for performance testing may be unwieldy, and may be skipped until software is in production.

In addition, testing at higher scale is becoming of paramount importance. With delivery of applications in a hosted, or SaaS model, the number of concurrent users accessing the system can skyrocket. This impacts network and server loads, as well as storage capacity and performance. Software that performs well for 10 users, often encounters bottlenecks or bugs when exposed to 100s or 1000s of concurrent users. Such workloads and beyond are typical for SaaS based applications. The need for QA to have these data center resources at their disposal, on-demand, can be viewed as a major accelerator for time to market, with the additional assurance that scalability and performance have been thoroughly tested.

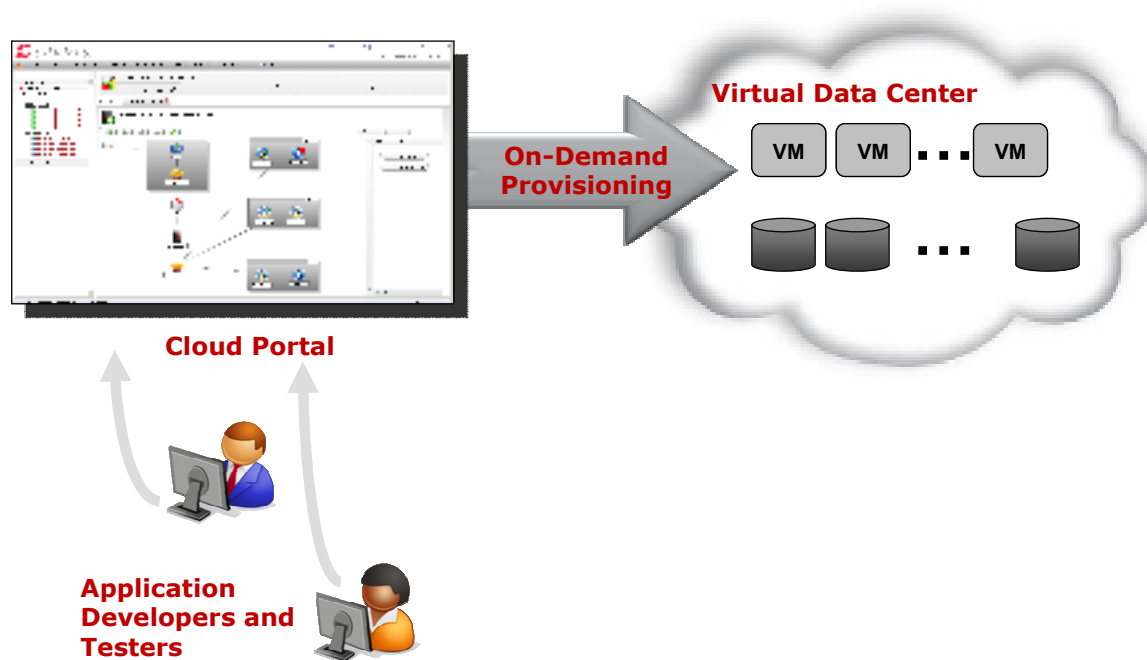
Next Generation Cloud Services: Enabling Agility for Developers

More agile IT environments are required, to match the highly dynamic, and resource intensive needs of the application development environment. Solving these problems would provide a tremendous enabler and time-to-market accelerator for applications. A new generation of cloud services is emerging, that can address these requirements for more agile IT, by providing the following powerful capabilities:

- **Agile, on-demand provisioning of all data center resources in under an hour, through self-service access**
- **Multiple service levels to address the wide ranging requirements of the application lifecycle**
- **Fast scale-up or scale-down of resources, with usage-based billing**
- **Support for virtual environments and multi-tier application architectures**
- **Migrate existing virtual server images and workloads into the cloud**

In turn, these capabilities can be leveraged by application developers to dramatically speed time-to-market for new applications.

The ability to perform on-demand provisioning of data center resources enables developers to quickly provision development server environments for any of their required tasks. This can also enable parallel streams of development, as multiple environments can quickly be provisioned for development in multiple branches, or for running multiple test runs concurrent with development.

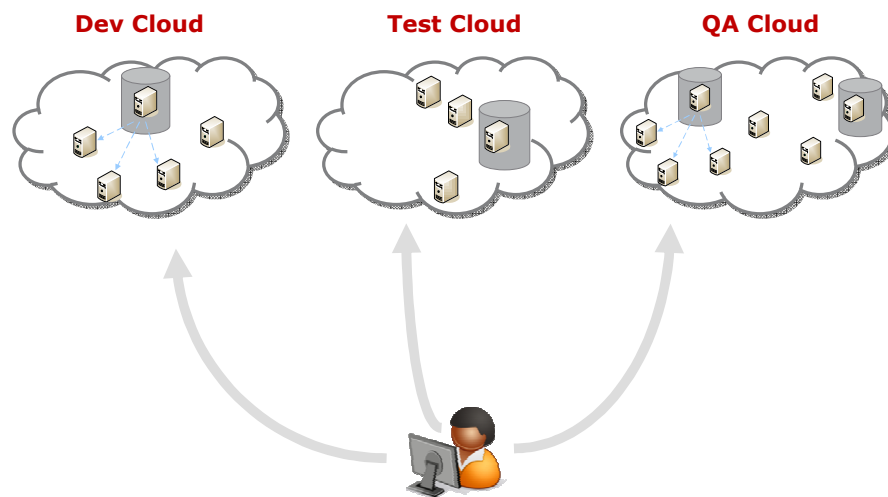


On-demand resource provisioning in the cloud

The ability for developers to perform this provisioning on a “self-service” basis is another accelerator of productivity, as it frees key IT resources from becoming bottlenecks in this process. This provides developers with the ability to simply request resources and make use of them as needed, avoiding the IT bottleneck. IT administrators will rest easy knowing that they can provide credit limits on the amount of resources consumed by a developer, or by the development organization as a whole. These credit limits can be monitored and adjusted as needed to maintain a level of governance on development expenditures.

Cloud environments that provide multiple service levels bring another dimension for application developers: the ability to optimize price/performance for each stage of the application lifecycle. For example, the ability for developers to create an environment in a cloud that is optimized for the needs of development and testing would entail lower levels of performance guarantees, but at the return of a lower hourly cost. As the application lifecycle

moves to QA, the environment needs to more closely mirror the production environment, so higher levels of performance, security, and availability need to be assigned in exchange for a higher cost of infrastructure. QA will also benefit from the on-demand up and down scalability provided by the cloud, as this finally solves the problem of testing applications at large scale, but without the need to indefinitely reserve and pay for these resources when they are unused. With usage based billing the cloud will bring closure to the need to pay for unused resources, enabling customers to finally pay for resources used, as opposed to resources deployed.

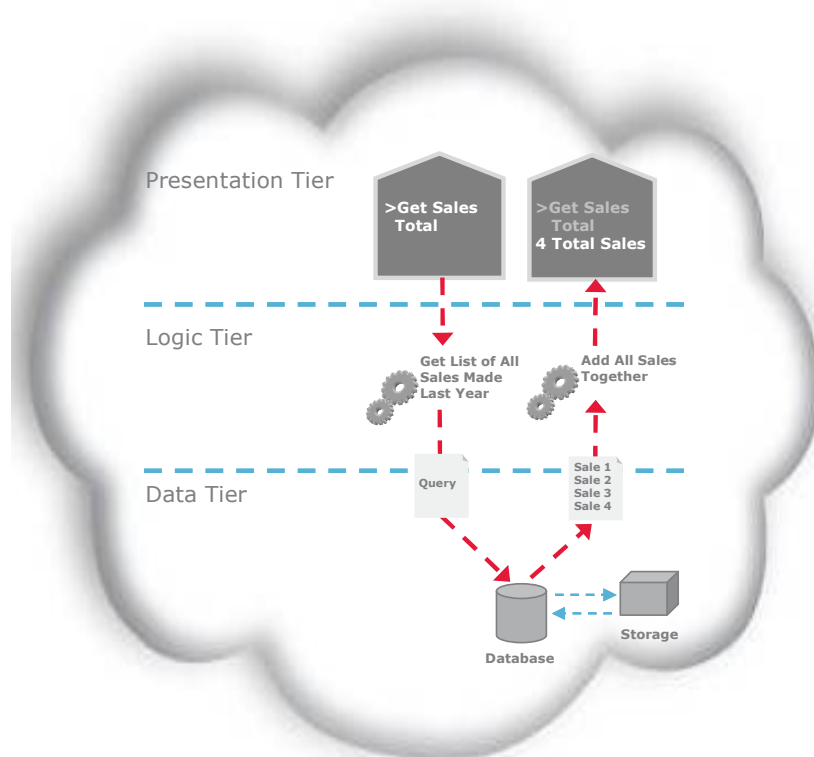


Accelerating Development through Parallel Clouds

Finally, for a cloud environment to fully support the application lifecycle, it should be provide the capability to optimize the environment for production applications with a variety of needs, ranging from web hosting, to eCommerce and finally to mission-critical applications. Cloud environments that provide such automation of multiple service levels for these “virtual data centers” are now emerging on the market to deliver on this promise.

Cloud services have offered virtualized environments for the last several years, based on a variety of commercial and open source virtualization technologies. Beginning with virtual hosting services, these clouds have provided on-demand virtual server provisioning capabilities for end-users. Now, a select set of cloud services are also beginning to support modern, multi-tier application architectures. Typically, these are 3-tier architectures consisting of a presentation (or UI) tier; a logic (or application) tier; and finally a data tier. This directly supports the current model of web application development, by providing virtual data centers that support deployment of virtual servers into each tier, and furthermore provide security separation between the tiers. Since the 3-tier model was designed to make applications highly scalable, it ensures that new applications will be able

to grow to accommodate the large workloads of web and eCommerce applications. By supporting multi-tier architectures directly, this ensures that applications can easily be moved from development and test clouds into production clouds.



Multi-Tier Application Architectures in the cloud

Advanced cloud services also provide support for uploading existing operating system and application images into the cloud, from the developer's private virtualized environment. This will enable rapid movement of new or existing applications from a development sandbox, into a larger scale test or QA environment as needed, without the need to configure the application from scratch.

Cloud Services: Return on Investment

The cloud services market is now directly addressing the needs of application developers, with an emerging class of services that provides fast, on-demand provisioning; instant scale up and down capabilities, and support for multi-tier application architectures, as well as workload migration capabilities.

A variety of cloud services supporting these services is becoming available. These services in many cases offer a "zero CapEx" model, where there is no up-front fee, and no long term

commitment required in order to use the service. Instead, these services are based on simple usage-based metering and billing models. This enables application developers to truly leverage the on-demand nature of these clouds, and pay only for actual resources consumed.

For many corporations, this type of pay-as-you go model for development infrastructure will enable tremendous cost savings, as well as earlier time to market. This enables the end goal of revenue acceleration. The return on investment is clear: developing applications in cloud services can provide near term ROI, through reduced CapEx, usage-based billing, and most importantly through accelerated time to market.