NEOVISE

Build Your Own Cloud:

Personalized Cloud Servers from SoftLayer

A Neovise Vendor Perspective Report

Executive Summary

Cloud computing provides a variety of benefits including flexibility and control. However, in the case of infrastructure as a service, practically all public cloud service providers offer only a relatively small number of pre-configured virtual machine images. This is somewhat analogous to offering shoes without half-sizes or widths, and then – when the fit isn't quite right – asking the customer to try on a brand or style that is sized a bit differently. The customer may end up with the right size, or the right brand, or the right style – but not all three. In the case of IaaS, this means settling for a server which does not fit the needs of an application in terms of CPU, memory or storage.

SoftLayer has recently updated its renowned CloudLayer service with a *Build Your Own Cloud* capability. This new approach allows customers to configure the amount of CPU, memory and storage used by their cloud servers. With this enhanced capability, SoftLayer customers avoid the tradeoffs forced by other IaaS offerings. These tradeoffs typically include *over provisioning*, which leads to inefficient resource utilization and higher cost, or *under provisioning*, which leads to lower levels of performance and scalability. Rather than constraining the dimensions of your server, CloudLayer lets you customize your server to meet the specific needs of your application. After all, who knows your application best – you or your IaaS provider? Neovise believes that any organization wishing to run production applications in a public cloud needs to strongly consider SoftLayer.

Background

Two of the key benefits provided by cloud computing, particularly infrastructure as a service (IaaS), are *control* and *flexibility*. Control puts decisions and actions in the hands of the customer. This commonly happens through self-service interfaces that enable automated infrastructure provisioning, monitoring and management. It also happens through APIs which enable even greater automation by supporting integration with other tools and processes, as well as by enabling applications to self-manage. Flexibility provides customers with a broad range of capabilities and choices. As related to cloud computing, flexibility includes adjusting the mix of capital and operating expenses, gaining access to additional infrastructure resources in minutes rather than weeks, responding to changes in demand for IT services, and much more.

Together, control and flexibility allow customers to make a variety of choices around the management and consumption of IT. These choices include:

- Who to handle infrastructure management (e.g. internal staff or service provider staff)
- Where to process workloads (e.g. internal datacenter or in a public cloud)
- When to process workloads (e.g. resource driven or priority driven)
- How much to consume (e.g. policy driven or demand driven)
- How to consume IT (e.g. central administration or self-service)

The Challenge

However, even with the benefits provided through cloud computing, an important aspect of control and flexibility has been missing at the system administration level: *the ability to configure certain key attributes of underlying virtual machines*. Rather than allowing custom configuration, practically all public cloud service providers simply offer a relatively small number of pre-configured virtual machine images (VMI). The choices often start with small, medium and large VMIs, each with pre-set amounts of CPU, RAM and storage. Unfortunately, there are a couple issues with this approach.

"Rather than allowing custom configuration, practically all public cloud service providers simply offer a relatively small number of pre-configured virtual machine images."

The first issue is that many more variations of overall server size are needed than small, medium and large. While some cloud service providers indeed offer a few more choices here, the fine grained control of server sizing is still not available to the customer. More importantly, a given application is likely to end up with a server which is either too small or too large. If the server is too small, performance issues may arise. If the server is too large, then the customer is paying for more resources than necessary. Offering a small set of fixed server sizes is really an artificial concept in the world of virtualized cloud servers and is contrary to the benefits of control and flexibility normally associated with cloud computing.

The second issue is related to locking down the relative quantities of CPU, memory and storage within VMIs. The truth is that servers are not accurately measured based on a single dimension of size. Servers are more commonly measured using a combination of CPU, memory and storage capacities. This is not an academic exercise. Applications perform best when running on servers with optimized configurations. Yet every application is different. This means there is no single small set of configurations that will best meet the needs of all applications. Even so, most cloud service providers have attempted to address this issue by offering a limited set of pre-configured VMIs that vary across the more granular dimensions CPU, memory and storage size. For example, a user can purchase a medium-sized, memory intensive VMI or a large-sized, CPU intensive VMI.

Packaging VMIs like this may appear on the surface to meet the needs of customers and their greatly varying applications. After all, this approach not only includes a lever for overall server sizing, but also a lever for a few common application variations, including CPU intensive versus memory intensive versus storage intensive. While the basic idea is in place, the control and flexibility is not. Customers are not allowed to set these attributes exactly as desired, so they must choose the VMI which mostly closely matches their needs. Unfortunately, the closest pre-set server configuration is still not likely to be optimized to the specific needs of a given application. If a VMI is optimized on one dimension, such as RAM, there will likely be a tradeoff: more or less CPU and / or storage than is optimal. If the VMI is over provisioned on CPU or storage, there is additional expense for unneeded resources. If the VMI is under provisioned, the application may not meet performance or capacity requirements.

The SoftLayer Solution: Build Your Own Cloud

CloudLayer™ Computing from SoftLayer lets customers deploy on-demand computing instances running enterprise-grade infrastructure in as few as five minutes. The CloudLayer Computing Instances are available with the customer's choice of operating system. Built on SoftLayer's automated data center services expertise and Citrix® XenServer™, these instances can stand alone or be seamlessly integrated with SoftLayer dedicated servers, virtual servers, automated services, and additional CloudLayer Services, to create a unified computing environment with the highest levels of efficiency and interoperability. Further, with the new *Build Your Own Cloud* capability which is now a standard part of CloudLayer, customers can configure the amount of CPU, memory and storage used by their virtual machine images.

Build Your Own Cloud allows customers to configure 1 to 16 CPU cores and 1 to 16 GB of RAM. CloudLayer Computing Instances come with 25 GB of free SAN storage to cover operating system installation and provide some additional headroom. Up to 600 GB of SAN storage can be configured. Configuration can be done manually through the SoftLayer Customer Portal or programmatically through the SoftLayer API. Customers still get to choose whether to pay for CloudLayer on an on-demand or monthly basis. The starting price for a base configuration is $0.10 per hour or $70 per month. Monthly plans include 2,000 GB of bandwidth. As shown in the screen shot below, SoftLayer provides an online price calculator to quickly determine on-demand and monthly rates for any given configuration.

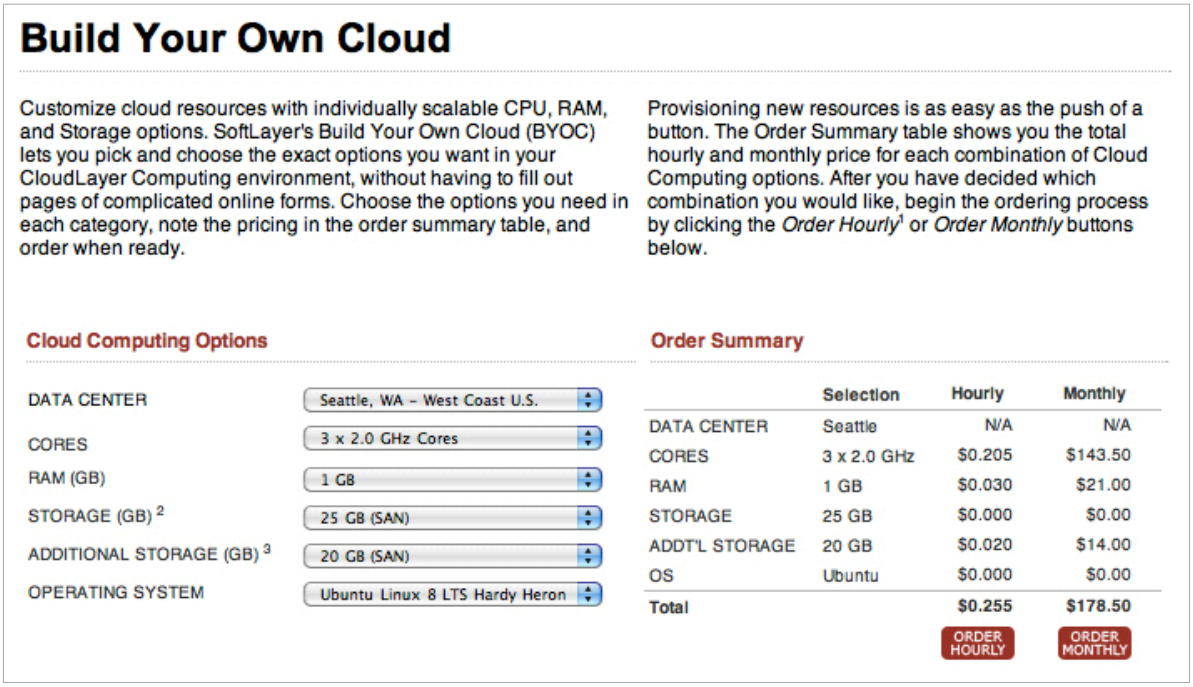

Screen Shot: CloudLayer Price Calculator (http://www.softlayer.com/byoc/)

Capturing the Benefits

Having the flexibility and control to set CPU, memory and storage parameters can be helpful for any application, not just those with extreme needs along one particular dimension. A couple specific use cases offer clarify this:

- **Typical case** – Using only pre-set VMIs inevitably leads to one or more of tradeoffs:
  - CloudLayer avoids the over provisioning tradeoff which leads to inefficient resource utilization and higher cost
  - CloudLayer also avoids the under provisioning tradeoff which leads to lower levels of performance and scalability
- **Scale up applications** – Cloud computing is outstanding for scale out applications that increase their performance by adding additional servers. However, most databases achieve improved performance by *scaling up* which means using a more powerful server rather than adding additional servers. Applications running on CloudLayer Computing Instances can easily be scaled up in whatever increments best suit demand. Production testing can be completed in advance on a variety of instance sizes, then – as demand fluctuates – applications can scale up or down in minutes.
- **Multi-tier applications** – Many enterprise applications are designed with multiple software tiers, each with different server requirements. With CloudLayer, every server used by the application can have its own configuration. If the application tier is CPU intensive, the CloudLayer Cloud Instance can be configured appropriately. Further, CloudLayer will allow one tier to scale up / down while another tier scales out / in.

Neovise Perspective

SoftLayer customers already appreciate the unprecedented levels of automation available to them through the CloudLayer API and Customer Portal. Now, with the new Build Your Own Cloud capability, they also have new levels of control and flexibility. Rather than guessing at generic application requirements – or worse, offering several hundred variations to choose among – SoftLayer has decided to hand over control to the customer. That is exactly what customers wanted.

For the cloud computing industry at large, this is a classic case of "a lot to gain and nothing to lose." There is no longer a good reason to simply settle for a limited set of fixed configurations around CPU, memory and storage. In hindsight, users of public cloud solutions will surely wonder why the industry norm for so long had been to lock these parameters together as though a few magic ratios between them could somehow meet the needs of many thousands of different applications. It may take a couple years, but Neovise predicts that the vast majority of public cloud service providers will adopt this same approach.

"For the cloud computing industry at large, this is a classic case of "a lot to gain and nothing to lose."

Interestingly, this new capability provided by SoftLayer says a lot about what the future of cloud computing can and should become. While SoftLayer has focused on delivering enterprise class capabilities for cloud computing as well as other on-demand datacenter and hosting services, a few other service providers have focused on commoditization – delivering cookie cutter capabilities at fairly low prices. Offering a commodity is not necessarily a bad approach since the market will likely support at least a couple commodity players.

However, many enterprise IT shops, government IT organizations and others need more. What they need is mass customization – the ability to get personalized capabilities like Build Your Own Cloud – without paying high custom prices. SoftLayer has set the mass customization bar high – not just with personalized cloud servers but with highly automated control and flexibility throughout its offerings. Neovise believes that any organization wishing to run production applications in a public cloud needs to strongly consider SoftLayer.

"SoftLayer has set the mass customization bar high – not just with personalized cloud servers but with highly automated control and flexibility throughout its offerings."

About Neovise

Neovise delivers essential knowledge and guidance to key members of the IT community including IT leaders, IT staff, business executives, technology vendors, systems integrators, and service providers. While each of these roles has a separate focus, they all require value added information and advice in order to formulate winning strategies and make optimal decisions.

Based on decades of industry experience and ongoing market analysis, we are your trusted, independent source of *Business Driven IT Insight*™.

For more information, visit www.neovise.com.