

# Applying Cloud Computing to IT Management

Prepared by Michael Coté ([cote@redmonk.com](mailto:cote@redmonk.com))

Sep, 2010

## Executive Summary

"Cloud computing" is a concept at the height of its popularity. In addition to driving frenzied desires at the infrastructure layer, it's driving changes to how applications are developed and delivered, how IT services are provided, including automation.

This brief paper provides a working definition of "cloud computing" and then discusses three areas that IT organizations must target for change to gain the full benefits of cloud computing:

- **IT management at large scale, without IT intervention** - the efficiency of a cloud rests on operating at large scale and especially with as little human intervention as possible.
- **Configurations Management & Automation** - The configuration and provisioning needs of the cloud must be highly automated and often assume less than perfect process and diagnostics.
- **Delivering Frequent Functionality** - Taking advantage of the cloud depends on delivering applications in a rapid, Agile fashion, which depends on automation being an enabler instead of yet another moving part to wrestle with.

## What is "Cloud"?

The "cloud" has many definitions - sometimes more than one per vendor and provider! I like to start with a simple definition to have something to start with. First, what we know of cloud computing comes from recent innovations at high-scale, public web companies like Google, Facebook, Amazon, and others. These companies needed new ways to handle massive amounts of traffic and computation, and also a way to deliver new features and applications rapidly, in some cases, weekly, if not daily. Traditional IT methods and offerings weren't cutting it. These companies evolved a way of using commodity hardware and software to create pools of compute resources that were used on demand, often virtualized and always fully automated.

Corporate IT has been learning from these early cloud innovators, trying to turn the practices and technologies for public clouds into "private clouds" they can run in their own datacenters, as well as simply using public clouds like Amazon and Rackspace.

## Three Layers of Cloud

I like to divide the idea of a cloud into three slices - like any good three tiered description of an IT service delivery platform:

- At the bottom, is **Infrastructure as a Service, or IaaS**, that pools together and provides raw server, storage, and network resources. Amazon and Rackspace provide IaaS.
- In the middle, is the **Platform as a Service layer, or PaaS**, which is intended to provide a collection of middleware and frameworks on-top of an IaaS that a developer can write applications in. Think of it as an application server for the cloud. Google App Engine fits here, as well as Salesforce's Force.com and Intuit's Partner Platform - each providing the raw "stuff" developers need to write and deliver applications, beyond just servers.
- On-top is the oldest part of the cloud, **Software-as-a-Service, or SaaS**, which we've actually had long before "cloud computing" was a popular phrase. This is the hosted application a user interacts with, usually through the web but increasingly as mobile apps as well, and often multi-tenant. Think Google Apps, Salesforce, and Microsoft's hosted email.

## Changing IT Management to Benefit from Cloud

With that context, let's look at how cloud computing is driving needs in automation.

### IT management at large scale, without IT intervention

Operating at cloud scale requires using pools of homogeneous resources that can be swapped between different workloads as needed, automatically. Using dedicated resources - servers - slows things down. Each time something special is required, you get out-of-band of the standardized way you'd like to operate a cloud. More importantly, handling actual hardware - moving wires around, installing operating systems, and setting network configuration manually when servers are coming up and down all the time - eliminates the efficiencies of cloud computing.

Your IT Management, and especially automation, software needs to match these requirements, allowing you deal with pools of resources and swap them around as needed.

### Configurations Management & Automation at Scale

As you can imagine, in this kind of environment, provisioning and configuration are happening all the time. Instead of being a once-in-a-release activity, resources in your cloud are constantly being shifted around for different uses, updated even as you begin to deliver new applications and features more frequently.

Some of the efficiency benefits of cloud computing come from not diagnosing every problem, and instead just "blowing it away" and starting fresh: you assume failure, but always have the magic reset button as a fix. If the best fix for any problem is to simply re-provision a fresh node, why spend time diagnosing anything? Clearly, that's an exaggeration and there will be many exceptions, but the principal a starting fresh should top the troubleshooting and recovery list.

Your automation software needs to operate in an environment where it becomes part of the daily running, diagnosing of your, and repairing your IT services.

## Delivering Frequent Functionality

The most radical need that cloud computing is driving in IT Management is the need to collaborate with developers. In addition to driving down costs, the ultimate goal of cloud computing is to allow IT to deliver more services more frequently to The Business, customers, or whoever IT's customers and end-users are.

Instead of taking 18 months to push out a new feature, leaders in cloud computing push out new features weekly, if not daily. Note the use of the term "feature" vs. "release": orienting your automation around delivering small chunks of functionality is critical.

While developers must re-orient their application architectures and project management to deliver frequent functionality, operations must re-orient as well. Operations has to start thinking of the pools of resources in their cloud(s) as something they "program" as part of the applications and services being delivered, not just infrastructure that's manually tuned and configured.

There isn't time for IT to give loving care to each resource or even orchestrate everything manually. Instead, you need your automation system to act like a compiler and runtime engine to do your infrastructure bidding at scale.

## Guiding Principals

This has been a very brief treatment of the topic, but you should get a sense of how the "needs" of the cloud are driving requirements for automation software. To pull out some general principals from the above, reaching the full benefits of cloud computing depends on IT changing the way it operates and thinks in several ways:

- **Failure is a Feature** - assume failures are going to happen all the time and build your automation system as one of the ways to recover by re-provisioning fresh, hopefully error free resources.
- **Embrace Simplicity** - try to keep your pools of resources as homogeneous as possible instead of requiring extra attention for specialized IT. Anytime you need to specialize, ask why and try to avoid it.
- **Automate Everything** - otherwise you'll have little time to do anything. Operating at cloud scale requires that humans intervene at only the most critical moment, and not be constant hand-holders for cloud resources.

Most of all, remember: if you're can't actually benefit from cloud computing - or any new technology - don't waste your time with it, stick with what works until something better comes along.

## About RedMonk

[RedMonk](#) is the first and only "maker" focused industry analyst firm. We believe that developers, operations staff, and those who are on the front lines of implementing and using IT are the most important constituency in technology.

We focus on how new and old technologies are being applied by these makers to run businesses and help achieve the goals of their organizations. RedMonk advises both buyers and sellers of technology, providing [all of our research for free at RedMonk.com](#) in the form of blogs, podcasts, videos, presentations, and other mediums.

While it's impossible given the breadth to simply distill our coverage and views, the core thesis that guides much of our work is that technology adoption is increasingly a bottom up proposition. The supporting evidence abounds; think Linux, Apache, MySQL, PHP, Firefox, Cloud Computing, Eclipse, and the consumerization of IT. All of these are successful because they've built from the ground floor, often in grassroots fashion.

So the question we pose to you is this: you may have analysts that help you understand top down. [Who do you have that does bottom up?](#)