



The new building blocks for IT: OpenStack, continuous delivery, and devops

Jo Maitland

May 28, 2013

This report was underwritten by Mirantis.



Cloud

TABLE OF CONTENTS

Executive summary.....	3
A new approach.....	4
Highly flexible infrastructure.....	4
Why OpenStack?.....	5
OpenStack community matures.....	6
Highly flexible software delivery lifecycle.....	7
MercadoLibre’s approach to continuous delivery.....	8
Highly flexible IT organization.....	8
The future of OpenStack and continuous delivery.....	10
Key takeaways.....	11
About Jo Maitland.....	12
About GigaOM Research.....	12

Executive summary

The key building blocks of a modern IT organization include a highly flexible infrastructure, an automated software delivery life cycle, and a devops-driven IT organization. These capabilities are essential to ensuring that IT remains relevant in an era of continuous change to customer-facing services and automated business operations.

This report focuses in particular on OpenStack as an underlying cloud platform to support this new type of organization. In addition, it examines a handful of ecommerce and Software-as-a-Service (SaaS) providers and the challenges these companies seek to solve by adopting a more agile software development and deployment process using OpenStack.

Key issues this report addresses are:

- Defining the continuous delivery (CD) approach
- Deploying the cloud as the CD infrastructure: public to private
- Choosing OpenStack as an infrastructure platform on which to build a private cloud
- Addressing the benefits and disadvantages of the OpenStack community
- Moving toward a continuous delivery model to reap the full benefits of a cloud platform
- Defining the role of devops

A new approach

In a high-growth environment, the goal is to get great ideas from a businessperson's head into a new product as fast as possible. Unfortunately, legacy IT systems and processes are no longer up to this task. Companies like Netflix, Amazon, and Facebook did not become global powerhouses by running on Oracle, Microsoft, BMC, CA, or any other outdated client-based or server-based architecture. They built their businesses on a flexible cloud infrastructure, and they use continuous delivery (CD) techniques to deploy new features to their products as fast as possible, with the highest quality and reliability.

GigaOM Research analyst and devops consultant Paul Duvall describes [continuous delivery](#) as “the automated implementation of the build, deploy, test, and release process. CD provides quick, automated feedback on the production readiness of a release candidate, with every change committed to the version control system. CD embraces the devops philosophy and defines the steps and activities for regularly delivering features to users in stable environments.”

Successful SaaS providers such as Cisco WebEx and eBay are working toward a devops model, restructuring their IT departments in such a way that every member of the team has a stake in the success of the whole operation, not just in the success of an individual system or silo within it.

Highly flexible infrastructure

Deploying the right infrastructure is the first step. The traditional way of doing this — racking and stacking servers and provisioning every layer of the stack each time a company needs to test some new code — doesn't work when it comes to delivering the software quickly and reliably. Netflix and others have found that they need infrastructure that could change quickly. Virtualization was the first step along this path. Cloud computing is its continuation, evolving the IT infrastructure to be highly flexible, programmable, and efficient.

Typically enterprises take a two-pronged approach when they move to the cloud. First, they test prototypes on Amazon Web Services (AWS); then they pull those prototypes back in-house to build out a private cloud. The reasons for building a cloud in-house versus running a public cloud are numerous, but usually they are to ensure security, reliability, and control. Most SaaS and ecommerce providers say that when they have problems with their services, they need to own that problem from end to end in order to fix it. They can't afford to point fingers at a public cloud provider, nor can they afford downtime.

Additionally, to improve performance, they want to place their applications close to where their customers might be, which isn't always where public cloud data centers are located. MercadoLibre, a subsidiary of eBay in Latin America, is a good example. Its cloud was in Latin America long before AWS opened up shop in the region.

Why OpenStack?

Next comes the choice of an infrastructure platform to build a private cloud. Enterprises often start with VMware because that's what they are familiar with for virtualization. Or, if they are more open-source-oriented, they begin with platforms such as OpenStack or CloudStack. This paper will focus on OpenStack, which has the most momentum and backing in the industry today.

According to its mission statement, OpenStack is an Infrastructure-as-a-Service (IaaS) cloud computing project, run by the [OpenStack Foundation](#) to

“Produce a ubiquitous open source cloud computing platform that will meet the needs of public and private clouds regardless of size, by being simple to implement and massively scalable.”

Companies that have built private clouds based on OpenStack include MercadoLibre, eBay, Disney, Cisco (WebEx), and Argonne National Laboratory, to name a few. Service providers running public clouds on OpenStack include Deutsche Telekom, Rackspace, and HP Cloud, with the list expanding daily.

One of the major attractions of OpenStack, according to its users, is its open-source model. This allows users to customize the platform to suit their requirements. MercadoLibre, for example, developed a custom application programming interface (API) for its OpenStack cloud that gave it a central point to schedule jobs across its two clusters running in separate data centers. The open-source model also means users can perform bug fixes themselves, quickly if necessarily. In a closed, proprietary system, they would have to wait for the vendor to make the fix, which could take weeks or months if the issue is not a priority to the majority of customers.

In the latter half of 2011, Cisco's WebEx team chose OpenStack as its cloud-infrastructure platform over VMware, CloudStack, and Eucalyptus, principally because it was fully open-source. According to architects at the company, by going that route, if Cisco needed specialized schedulers or disk storage, it could simply extend the platform to suit its needs. The abundant availability of third-party commercial support was also an important factor in its decision.

Another benefit OpenStack users often cite is that the open-source model gives them more control over the cost of their cloud environments, versus the perpetual license fee approach of VMware, Oracle, Microsoft, and other vendors selling proprietary cloud platforms.

The OpenStack community matures

Following pressure from users to ensure OpenStack would stay true to its open-source roots and benefit the community as a whole, its founders at Rackspace turned OpenStack over to the newly created OpenStack Foundation in September 2012. The foundation is now the independent home for OpenStack. It has around 7,000 individual members and 850 different organizations participating in it. This ever-expanding group of participants is both a strength and potential weakness of the project. While OpenStack has tremendous support and momentum from IT giants (IBM, Red Hat, Cisco, HP, Intel, and others), it also has many cooks in the kitchen, each with an opinion on how to bake the product. Some argue this is an advantage — the more ideas, the better the final product — but it also has the potential to slow things down.

The large group of participants also means that the available talent capable of building OpenStack clouds is hard to come by. OpenStack developers are quickly swallowed up by the participating IT giants or by cloud builders and systems integrators. Consequently, acquisition of skills and knowledge to successfully stand up, automate, and operate the OpenStack cloud is rapidly becoming a differentiator.

Companies like MercadoLibre in Latin America and Nokta, a video service similar to YouTube in Turkey, have an even harder time finding engineers that know OpenStack, as the available talent in those regions is limited. Nokta's chief software developer, Erdem Ağaoğlu, had a positive spin on the challenge. He said his team of three system administrators running Nokta's OpenStack cloud believes it has a huge opportunity to build the first major OpenStack cloud in Turkey. MercadoLibre has four systems administrators managing its OpenStack cloud. It searched for developers comfortable with Python and open source as well as Linux administrators, and they learned OpenStack on the job.

“It's like 1997 all over again, except instead of trying to find good Linux talent, now it's OpenStack,” said Damon Edwards, the co-founder of DTO Solutions, a devops training company. In the long run, perhaps three to five years from now, the OpenStack market will be more mature and OpenStack engineers will be more abundant.

Filling the gap in the meantime are OpenStack consultancies and systems integrators (SIs), who are doing a roaring trade helping enterprises build OpenStack clouds. But SIs are not in the operations outsourcing business. The IT organization must still do the hard work of retooling its operations team to develop, manage, and drive its OpenStack cloud over the long term.

This is where a change in mindset around IT operations and software development best practices often comes into play.

Highly flexible software-delivery life cycle

Adding a cloud wrapper, such as OpenStack, around legacy IT processes just creates a more complex hosting environment. The business often get disillusioned at this point, as it is not getting the fast turnaround of new services that was promised by moving to the cloud.

Eventually a light bulb goes off in the heads of the IT team as someone realizes that the team must also rethink its software-delivery processes and tools and then move toward a more continuous delivery model to reap the full benefits of a cloud platform. For companies that are already familiar with the agile methodology, which uses incremental development techniques, continuous delivery isn't too big of a leap. But for those still in the world of traditional software development, it can be a daunting challenge.

In the old method of software development, the train (i.e., new release) comes around once or twice per year and the passengers (i.e., developers) must get everything ready to board the train. Massive teams muster everything they've got to pour into the big summer release. Application developers work for months in stand-alone development and test environments to build their applications. Often releases are done in the middle of the night with all hands on deck, forklifting the new release into the highly optimized production environment, accompanied by much nail-biting in the hopes that the new code doesn't screw everything up. And things do go wrong; because the code was built for months on a platform that may have subtle or even gross differences from the configurations on which it will finally run, mismatches can bring the entire operation to a halt.

In the world of continuous integration and continuous delivery, these epic releases go away, replaced by much more frequent updates. There's no rolling back to former versions of code: Updates only move forward and patch along the way. What's more, with the fluid nature of cloud operations such as Infrastructure as a Service (IaaS), configuration parameters can be captured programmatically at

development time. And these same configuration standards can be applied and enforced during development, preventing surprises downstream at deployment. This means no more nail-biting all-nighters and massive coordination efforts. It also means that the IT team is no longer slowing down the business but instead is powering it forward.

MercadoLibre's approach to continuous delivery

With 62 million registered users and 2 million requests on its auction site every minute, MercadoLibre needs to be able to deliver new features to its platform fast and reliably if it is to stay competitive. If its auction site is down and users can't make bids, MercadoLibre doesn't have a business.

The company runs its auction site on OpenStack across two data centers for redundancy, housing 6,000 to 7,000 virtual machines. It uses Opscode Chef tools to handle all the configuration and management of its OpenStack cloud. Alejandro Comisario, the cloud architect and infrastructure technical lead at MercadoLibre, noted that when a company can provision new environments and launch code in literally seconds, it must have a different set of rules and processes for using that infrastructure. Using OpenStack with Chef helped him enforce controls over the new environment. His team no longer needs to spend as much time deploying and managing physical servers. Previously, deployments could take hours, even days. Using OpenStack plus Chef, deployments are measured in minutes and seconds.

Furthermore, MercadoLibre's development team can directly deploy APIs and apps themselves, without the intermediary steps involving the IT team. Thus they can rapidly speed up time to market for new features and services. Previously those deployments were restricted to once every 24 hours. Now they happen two to three times per day.

The ability to write to an API and deploy applications in a self-service way is central to the practice of continuous delivery. It stands to reason that the underlying cloud platform needs a solid API that can be controlled remotely, is simple to use, and typically adheres to a [RESTful](#) design. According to the users interviewed for this report, the OpenStack API meets these requirements, and it was in many cases simpler to use than the alternatives.

Highly flexible IT organization

Freeing up the development team to worry only about code and not the infrastructure required to run it is a relatively new phenomenon. Most IT organizations still operate with a mass-production mentality and

from functional silos or separated teams with separate budgets and agendas. Typically, two organizations are responsible for delivering software to users — development and operations — but they have different incentives.

Development teams are measured on the features they deliver, while operations teams are measured by the stability of the system after it launches. These are often competing interests in organizations, because the development team is not incentivized to ensure the stability of a system once it delivers it to the operations team. Likewise, the operations team is not afforded the appropriate incentives to care about the frequency of releasing new features to users, because its primary responsibility is to ensure the uptime of the system. Bringing development and operations teams together, referred to as “devops,” is not easy, given these two groups’ historically opposing goals.

Devops is a fundamental change in the structure and organization of an IT department. Instead of an IT organization consisting of development versus operations or database administrators versus QA, there is only one team: a delivery team. That team must see the software system as the sum of its parts; it’s not only the application source code but also the infrastructure, data, and configuration. When a configuration change has an effect on the infrastructure, every team member should care. While not every member will work to fix a specific problem, all team members should realize that any change affecting the viability of a release candidate affects each team member.

A major online payments company, which expects to have 10,000 virtual machines in production on its OpenStack cloud by year-end, has 15 people on its OpenStack engineering team, but many more remain in its traditional IT organization and development teams. The company said the transition must be organic rather than forcing people into a devops mindset. When other groups saw how fast the OpenStack engineering team was able to launch new services, they wanted a piece of the action, they found. Depending on the application, the OpenStack team at this company is releasing new functionality into production every day.

The company acknowledged that shifting to a devops culture takes time, but as more teams learn the capabilities of the new OpenStack architecture and continuous delivery model, they will join this group.

Cisco’s WebEx division doesn’t use the title “devops” to describe what it’s doing, but this definition nevertheless accurately explains how it is tightly integrating the collaboration between its engineering and operations teams to reengineer the way it deploys and manages its infrastructure. “We’re always trying to match agile operations with agile software development,” said Reinhardt Quelle, the operations architect at Cisco.

The future of OpenStack and continuous delivery

When developers needed a new machine in the world of IT over the past two decades, they had to file a ticket and then wait days or maybe weeks for a response from IT. Now they simply sign in to the infrastructure self-service portal; a policy recognizes and authenticates their request, and the resources they need are available immediately.

By lifting the burden on the operations side, companies deploying flexible cloud-infrastructure platforms like OpenStack aim to be able to focus more on agile software development techniques and improving their processes, with the ultimate goal of matching their IT operational speed with the speed of their business.

Here are the trends to watch as this market unfolds:

- CI/CD is definitely the future, but these practices will take some time to propagate to the traditional, slow-moving enterprise.
- The first wave of adoption of these practices is being driven by technology-intensive organizations such as eBay, Cisco WebEx, and MercadoLibre.
- Because of the scale at which the above early adopters of CI/CD generally operate and the fact that they are technologically proficient, they are very opinionated about their infrastructure and need it customized to their process, road map, and physical environment.
- OpenStack is an ideal platform for these companies that are building their own flavor of the cloud, because it is a toolkit that can be molded to suit their specific needs. It allows these organizations to stay in control of their product and infrastructure road map while also containing costs.
- The question many are evaluating is the cost of prebuilt distributions versus the cost of staffing a growing team to build out their OpenStack cloud themselves. As the size of OpenStack deployments grows, cloud service providers have an opportunity to take on the overspill. This is an example of where the much-vaunted hybrid cloud computing model might eventually get some traction. Instead of an IT organization using its operations team to run development servers, these could run in a public OpenStack cloud, provisioned on-demand and dynamically from its OpenStack private cloud. OpenStack does not deliver this kind of capability yet, but it will get there.

Key takeaways

A modern IT organization requires three elements:

- A highly flexible cloud infrastructure built on a platform such as OpenStack
- A highly flexible software delivery life cycle known as continuous delivery
- A highly flexible IT department organized around a devops approach

All three are a challenging evolution from traditional IT infrastructure and operational practices of the past. But SaaS and ecommerce providers that have made the leap, such as Cisco (WebEx) and eBay, are significantly more agile and efficient at supporting their business goals.

CI/CD is an important trend that appears to be influencing the adoption of the cloud in general and OpenStack specifically. Today CI/CD is still something that is primarily getting embraced by technology-intensive organizations with an opinionated approach to infrastructure, and OpenStack, due to its flexibility, is a good fit for these companies.

Meanwhile mainstream enterprises looking for more turnkey cloud platforms will turn to systems integrators like Mirantis or packaged OpenStack distributions from the likes of CloudScaling and Piston Cloud. Right now one can think of OpenStack as a grocery store offering the ingredients to build a cloud. If you are an exquisite cook with preferential tastes you are likely to go to the grocery store and cook your own meals. However, traditional enterprises that are not technology-centric prefer to go to a restaurant. As CI/CD propagates to that segment, expect to see more and more adoption of CloudScaling, Piston, and other packaged OpenStack distributions.

About Jo Maitland

Jo Maitland is the research director for the infrastructure and cloud channel at GigaOM Research. Maitland has been a technology journalist and analyst for over 15 years and specializes in enterprise IT trends, specifically infrastructure virtualization, storage, and cloud computing. At Forrester Research and the 451 Group, Maitland covered cloud-based storage and archiving and the challenges of long-term digital preservation. At TechTarget Maitland was the executive editor of several websites covering virtualization and cloud computing. Maitland has spoken at several major industry events including NetWorld + Interop and VMworld on virtualization and cloud computing trends. She has a B.A (Hons) in Journalism from the University of Creative Arts in the U.K.

About GigaOM Research

GigaOM Research gives you insider access to expert industry insights on emerging markets. Focused on delivering highly relevant and timely research to the people who need it most, our analysis, reports, and original research come from the most respected voices in the industry. Whether you're beginning to learn about a new market or are an industry insider, GigaOM Research addresses the need for relevant, illuminating insights into the industry's most dynamic markets.

Visit us at: pro.gigaom.com.

© 2013 Giga Omni Media, Inc. All Rights Reserved.

This publication may be used only as expressly permitted by license from GigaOM and may not be accessed, used, copied, distributed, published, sold, publicly displayed, or otherwise exploited without the express prior written permission of GigaOM. For licensing information, please [contact us](#).