

CLOUD U

DOES DATA WANT TO BE FREE? Exploring Issues in the Open Cloud from Vendor Lock-in to Open Standards

Through this year-long series of whitepapers and webinars, independent analyst Ben Kepes is creating a Cloud Computing curriculum designed for technologists and non-technical users alike. The mission is to build widespread knowledge about the Cloud revolution and encourage discussion about the Cloud's benefits for businesses of all sizes. Read more CloudU whitepapers and register for upcoming webinars at www.rackspace.com/cloud/cloudU

Table of Contents

Executive Summary	1
Portability—Why Open Standards are Necessary	2
A Historical Perspective— The Case for Open Source and Open Standards	3
Open Standards and the Open Cloud Initiative	4
Open Source Built the Web	5
Using a Robot to Fight a Fruit, Android Battles Apple	6
Open Source is Crucial for the Cloud, Too	7
Innovation is Critical, No Company Can Deliver it Alone	7
Vendor Lock In Should be Avoided at All Costs	7
Major Open Initiatives in Cloud Computing	8
APIs, the Backbone of Open Clouds	8
Hypervisors, Building Blocks of the Cloud	8
Open Compute—Flexibility Beyond the Building Blocks	9
Driving Efficiency—The Open Compute Project	10
Conclusion	11
About Diversity Analysis	12
About Rackspace	13

Executive Summary

Open Source software was fundamental in building the web into what it is today. In the same way that the “LAMP stack”—Linux, Apache, MySQL and PHP- was the foundation upon which the bulk of the internet was built, we believe that Open Source solutions and Open Standards will drive uptake, innovation and choice for Cloud Computing users.

With a plethora of different Open Standards initiatives, and a number of Open Source products available upon which to build Clouds, it is important for users to understand the landscape and the different approaches towards openness for the Cloud. As with other areas of technology, it could be argued that interest groups are using “open” as little more than a marketing ploy. However, it is our contention that the use of Open Source products, built with a view to adopting and embracing Open Standards, will drive flexibility, innovation and choice for consumers.

Portability—Why Open Standards are Necessary

As we've detailed in a previous CloudU report¹, a major driver of Cloud Computing is the migration of computing from on-premise owned infrastructure to subscription-based utility computing. While this move does much to commoditize computing, ultimate commoditization is hampered by the fact that Cloud Computing services are generally not fungible, that is they are generally not readily substitutable without exporting data from one vendor and importing it into another.

The process involved in this import and export either requires an intermediary service to move the data into a format recognized by the receiving provider or it requires the adoption of Open Standards, ways of creating schema, formats and data structures that enable the transfer of data from one provider to another.

A Historical Perspective– The Case for Open Source and Open Standards

Wikipedia defines² Open Source as practices in production and development that promote access to the end product's source materials. To achieve that aim, Open Source takes advantage of Open Standards which are defined by the Open Source Initiative³ as standards which do not prohibit conforming implementations in Open Source software. In practice this means standards that are reusable, royalty-free, have no discriminatory copyright issues and have multiple implementations.

It is worthwhile reflecting upon how Open Source and Open Standards came about in a product world and contrasting that with a service world. With software products, there was a continuum from *free* to *proprietary*. At the free end, source code was often made available under a copyright license that granted users some freedoms. These freedoms allowed them to run, configure and modify the software to their needs and to redistribute the amended product at will. At the other end of the spectrum, proprietary software didn't give users any access to source code whatsoever.

The rationale for Open Source was threefold;

- To allow and encourage a community to form around a software product and to contribute to the creation and extension of that software product. Innovation as driven by tapping into the collective skills of multiple contributors.
- By encouraging the use of Open Standards, users gained certainty that their data could be moved between applications thus avoiding lock-in with any one vendor. Users were given choice and flexibility.
- The spread of technology was democratized and accelerated as Open Source and Open Standards both reduced end users concerns about adopting technology and reduced the economic barriers to entry to adopting products.

Open Standards and the Open Cloud Initiative

The term “Open Standards” has multiple meanings and interpretations. The most applicable to Cloud Computing is that adopted by the non-profit Open Cloud Initiative (OCI)⁴, a non-profit advocate of Open Cloud Computing that was launched at OSCON in 2011⁵.

The OCI contends that for a Cloud to meet the requirements of being what it calls an Open Cloud, it must meet the following requirements;

- Open Formats: All user data and metadata must be represented in Open Standard formats
- Open Interfaces: All functionality must be exposed by way of Open Standard interfaces

The OCI goes further and says that Open Standards must meet the following requirements;

- Copyrights: The standard must be documented in all its details, published and both accessible and reusable free of charge.
- Patents: Any patents possibly present on parts of the standard must be irrevocably made available on a royalty-free basis.
- Trademarks: Any trademarks possibly present on identifier(s) must be used for non-discriminatory enforcement of compliance only.
- Implementations: There must be multiple full, faithful and interoperable implementations (for both client and server where applicable) and at least one such implementation must be licensed under an Open Source Initiative (OSI) approved license or placed into the public domain.

The Open Cloud Initiative however isn't simply an exercise unrelated to reality, as we will see in the next section, Open Source code, based on Open Standards, built the web and enabled the web to grow as meteorically as it did.

Open Source Built the Web

One of the reasons that the web grew so fast is that vast portions of it were built on top of Open Source software, specifically the LAMP stack. LAMP is an acronym for a complete solution of Open Source software that goes together to build a general purpose web server, these individual components are;

- The Linux operating system
- The Apache HTTP server
- The MySQL database
- The languages sitting atop the other components, generally Perl, PHP or Python

To see just how important the LAMP stack has been to the development of the web, it is worth noting that since May 1996, Apache has had the biggest market share in terms of servers across all domains.⁶ In July 2011 Apache accounted for over 65% of global web servers. By way of comparison, Microsoft IIS, a proprietary web server, accounts for around 16% of global servers.

We contend that the rapid growth of the web was, in part, due to the widespread adoption of Open Source technologies – these technologies encouraged innovation, experimentation and the democratization of the web.

There are modern parallels where Open Source products have encouraged innovation and it is to one example of this, the Android mobile operating system, which we shall now turn.

Using a Robot to Fight a Fruit, Android Battles Apple

Android is a software stack created specifically for mobile devices. It was developed by Google and the Open Handset Alliance, a consortium of businesses involved in the creation of mobile devices. Android was initially developed by a private company, Android Inc, that Google acquired in 2005. In 2007 Google released the bulk of the Android code under an Open Source license and created the Android Open Source Project to further develop the Android operating system.

Android was, in part, a response to the meteoric rise to prominence of the Apple iPhone and was arguably an experiment to see whether an Open Source and hence more flexible operating system would encourage innovation and drive adoption.

This experiment would seem to have proved the hypothesis. From an estimated base of 2.8% global share in Q2 2009⁷, Android rose to 33% market share by Q4 2010. Figures from August 2011 estimate Android at 48% of smartphone market share.⁸

We contend that one of the main reasons for this rapid growth is the fact that Android is an Open Source product with the result that device manufacturers can customize the operating system for their particular device – this leads to a high degree of innovation and products that can be specifically tailored for particular markets and customer types.

While Cloud Computing is very different from mobile phones, we believe that an open approach towards Cloud Computing will deliver benefits in the same way as it has done for the web itself and smartphones.

Open Source is Crucial for the Cloud, Too

We contend that in the same way that Open Source has been critical to the development of past technologies, so too is it important for the development of Cloud Computing. There are a number of reasons that this is the case.

Innovation is Critical, No Company Can Deliver it Alone

As we've seen from previous examples of technological disruption, it is rare for one company alone to be able to deliver complete innovation. We further contend that, innovation is more effective, quicker and of a higher quality when it is the amalgam of multiple organizations working on it in isolation, but united by a common goal.

Cloud Computing is a paradigm change and, requires software manufactures and service providers to solve hard problems. As such, we believe Cloud Computing that is the result of many organizations working on individual but connected problems under a common framework of standards, will drive the best and most efficient result for all concerned.

Vendor Lock In Should be Avoided at All Costs

One of the key reasons that Open Source grew in popularity was because of concerns around proprietary standards and formats. These standards and formats locked customers into a particular platform, frequently adding excessive amounts of time and cost to their modernization or transformation projects.

We contend that vendor lock in is an even greater risk under Cloud Computing as the vendor gets to pick not only the data formats but also the physical location of the data. This has the effect of causing a double lock in – firstly the user has difficulty extracting their data from the provider; and secondly if and when they are able to extract their data, it is in a format that is difficult for them to use with another provider.

We believe this is especially the case with the host of different services and providers in the Cloud space, Open Standards provides choice on many different levels.

With an understanding of why Open Source is important, it is now time to look at different areas that Open has been applied in the Cloud space

Major Open Initiatives in Cloud Computing

There are a number of different initiatives that broadly fall under the term *Open*. While some of these initiatives may not be Open Source, they all contribute to a more portable, open and flexible Cloud Computing environment and hence are worth reflecting upon.

APIs, the Backbone of Open Clouds

Application Programming Interfaces (APIs) are a particular set of rules and specifications that software programs can follow to communicate with each other.⁹ To put it simplistically, APIs are the translation services that allow different services to communicate and work with each other.

APIs are especially important for Cloud Computing as they make the Cloud both easier to work with, and more flexible to develop upon. As noted analyst Michael Cote put it¹⁰;

“ Instead of manually dealing with a user interface (or “console”) you can programmatically interact with the Cloud to do your Cloud operations. To pick one of the benefits of having a Cloud API: developers hope to better automate the role of the sys admin, leaving developers more time to focus on the application rather than running the application. ”

Given this important position that APIs hold in the very creation of Clouds, it is fair to say they are an important backbone for the open Cloud.

Hypervisors, Building Blocks of the Cloud

In a previous paper¹¹ we explained how virtualization was the key driver to increasing computing efficiency. Virtualization involves the use of hypervisors, virtual machines that allow multiple operating systems to exist on one piece of physical hardware.

There are a multitude of hypervisors available, both Open Source and proprietary. While the statistics for hypervisor use are not as skewed in favor of Open Source alternatives as they are for web servers, more organizations wishing to have a higher degree of flexibility around their virtualization strategy are preferring Open Source hypervisors such as Xen¹² for the flexibility and freedom it gives them.

Open Compute – Flexibility Beyond the Building Blocks

As we explained in a previous report,¹³ compute is the most fundamental aspect of Cloud infrastructure; it is the ability to make use of physical servers lying in a data center somewhere, on demand.

When Cloud Computing was first developed as a concept, most attention was given to building the scalability and rapid provisioning aspects of the Cloud, little or no thought was given to empowering users to have flexibility with their infrastructure. Since then there have been some notable initiatives and products introduced to make compute more open;

- OpenStack, originally sponsored by Rackspace and NASA, is a free Open Source project that allows organizations to build Public and Private Clouds. OpenStack has more than 100 companies as members or contributors.
- Eucalyptus is an alternative platform for the implementation of Private Clouds. It is sponsored by Eucalyptus Systems. Eucalyptus works with most Linux distributions and multiple hypervisor products.
- OpenNebula is an Open Source toolkit for managing distributed data center infrastructures. It is being used in a number of primarily European based projects. OpenNebula provides multiple choices in both interfaces and hypervisors.
- Nimbus is an Open Source toolkit that provides a Cloud infrastructure via industry standard APIs. Nimbus supports the Xen Open Source hypervisor
- Cloud.com, now owned by Citrix, has created an Open Source Cloud orchestration product with multiple hypervisor support.

With all of these compute level products in the Open Source realm, it is now time to look at Open Source approaches towards data center design.

Driving Efficiency–The Open Compute Project

Launched in April 2011, the Open Compute Project is an initiative sponsored by Facebook that seeks to share best practice approaches towards data center design in an effort to increase overall efficiency of infrastructure. As we detailed in our special Earth Day CloudU report,¹⁴ research suggests¹⁵ that the environmental footprint from data centers is set to more than triple between 2002 and 2020. This result would see data centers become the fastest-growing contributor to the total carbon footprint of the technology sector.

While the Open Compute Project doesn't relate purely to Cloud Computing infrastructure, it is an example of how open collaboration can drive efficiencies and innovations for the Cloud.

Conclusion

Cloud Computing is now a proven technology paradigm, as such there is widespread innovation and experimentation occurring all of which builds opportunities for new and exciting products and services for end users. In order to maximize these innovations, and give users security, flexibility and the greatest degree of freedom, we believe the development and adoption of Open Source solutions, built upon Open Standards is key.

While many different solutions exist in the marketplace offering Open Source Cloud solutions, we believe that in a short space of time the industry will settle on a single Open Source standard. Organizations such as the non-profit Open Cloud Initiative are ideal places to drive this transition and ensure a good result for all – vendors and customers alike.

About Diversity Analysis

Diversity Analysis is a broad spectrum consultancy specialising in SaaS, Cloud Computing and business strategy. Our research focuses on the trends in these areas with greater emphasis on technology, business strategies, mergers and acquisitions. The extensive experience of our analysts in the field and our closer interactions with both vendors and users of these technologies puts us in a unique position to understand their perspectives perfectly and, also, to offer our analysis to match their needs. Our Analysts take a deep dive into the latest technological developments in the above mentioned areas. This, in turn, helps our clients stay ahead of the competition by taking advantage of these newer technologies and, also, by understanding any pitfalls they have to avoid.

Our Offerings: We offer both analysis and consultancy in the areas related to SaaS and Cloud Computing. Our focus is on technology, business strategy, mergers and acquisitions. Our methodology is structured as follows:

- Research Alerts
- Research Briefings
- Whitepapers
- Case Studies

We also participate in various conferences and are available for vendor briefings through Telephone and/or Voice Over IP.



About Rackspace

Rackspace Hosting is the world's leading specialist in hosting and Cloud Computing. The San Antonio-based company provides Fanatical Support® to its customers, across a portfolio of IT services, including Managed Hosting and Cloud Computing. Rackspace is also the founder of OpenStack™, an open source cloud platform with broad industry support, designed to offer cloud consumers greater choice. For more information, visit <http://www.rackspace.com>.



About the Author

Ben Kepes

Ben is the founder and managing director of Diversity Limited, a consultancy specializing in Cloud Computing/SaaS, Collaboration, Business strategy and user-centric design. More information on Ben and Diversity Limited can be found at <http://diversity.net.nz>

Endnotes

- [1] http://broadcast.rackspace.com/hosting_knowledge/whitepapers/Revolution_Not_Evolution-Whitepaper.pdf
- [2] http://en.wikipedia.org/wiki/Open_source
- [3] http://en.wikipedia.org/wiki/Open_Source_Initiative
- [4] <http://www.openCloudinitiative.org>
- [5] <http://www.openCloudinitiative.org/press/releases/2011/launch>
- [6] <http://news.netcraft.com/archives/category/web-server-survey/>
- [7] http://www.appleinsider.com/articles/09/08/21/canalys_iphone_outsold_all_windows_mobile_phones_in_q2_2009.html
- [8] <http://www.canalys.com/newsroom/android-takes-almost-50-share-worldwide-smart-phone-market>
- [9] http://en.wikipedia.org/wiki/Application_programming_interface
- [10] <http://www.redmonk.com/cote/2009/07/14/rackspace-Cloud-api-open-Clouds-brief-note/>
- [11] http://broadcast.rackspace.com/hosting_knowledge/whitepapers/Revolution_Not_Evolution-Whitepaper.pdf
- [12] <http://en.wikipedia.org/wiki/Xen>
- [13] http://broadcast.rackspace.com/hosting_knowledge/whitepapers/SayGoodbyetoDIYDataCenters.pdf
- [14] http://broadcast.rackspace.com/hosting_knowledge/whitepapers/sustainability-and-the-Cloud.pdf
- [15] Cloud Computing and Sustainability: The Environmental Benefits of Moving to the Cloud, WSP, 2010