The Evolution of the Cloud
The Work, Progress and Outlook of Cloud Infrastructure

by

Ari Liberman García

B.S. Electronic Systems Engineering
Instituto Tecnológico y de Estudios Superiores de Monterrey, 2006

SUBMITTED TO THE SYSTEM DESIGN AND MANAGEMENT PROGRAM IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE

OF

MASTER OF SCIENCE IN ENGINEERING AND MANAGEMENT

AT THE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
FEBRUARY 2015

Signature of Author_____
Ari Liberman García
System Design and Management Program
January 2015

Certified by_____
Michael A. Cusumano
Thesis Supervisor
Sloan Management Review Distinguished Professor of Management
MIT Sloan School of Management

Accepted by_____
Patrick Hale
Director, System Design and Management Program
Massachusetts Institute of Technology

This page has been intentionally left blank

The Evolution of the Cloud
The Work, Progress and Outlook of Cloud Infrastructure

by

Ari Liberman García

SUBMITTED TO THE SYSTEM DESIGN AND MANAGEMENT PROGRAM IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN ENGINEERING AND MANAGEMENT

AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Abstract

Cloud computing has enabled the deployment of systems at scale without requiring deep expertise in infrastructure management or highly specialized personnel. In just a few years, cloud computing has become one of the fastest growing technology segments in the Information Technology industry and it has transformed how applications are created and how companies they manage their growth. The cloud market has quickly become one of the most competitive industries with companies committing their efforts to the creation of cloud platforms and aggressive pricing strategies in an attempt to gain market dominance.

This work shows the origins of the Infrastructure-as-a-Service industry and an analysis of the market dynamics by looking at the portfolios and strategies of the top competitors in this space. Also, this report shows what are the developments that will drive the innovation in the cloud industry years to come.

Thesis Supervisor: Michael A. Cusumano
Title: Sloan Management Review Distinguished Professor of Management

Acknowledgements

First of all, I would like to dedicate this thesis to my wife Alejandra and our daughter Eva for their love and continuous encouragement throughout my studies at MIT. I thank my parents and my family for always supporting me in my education and professional endeavors.

I also thank Professor Michael A. Cusumano, for providing his mentorship and advice while writing this thesis. Finally, I would like to thank the staff and my friends in the SDM program, especially to Pat Hale, Bill Foley and Amal Elalam, for making this a unique experience that I will cherish all of my life.

Table of Contents

List of Figures

Introduction

"Computing may someday be organized as a public utility just as the telephone system is a public utility. Each subscriber needs to pay only for the capacity he actually uses, but he has access to all programming languages characteristic of a very large system … Certain subscribers might offer service to other subscribers … The computer utility could become the basis of a new and important industry"

-Professor John McCarthy, 1961

Cloud computing has transformed how systems and applications are created, distributed and scaled and is arguably one of the most important technological developments in the last decade. This technology has allowed for individuals and companies to share a pool of resources of computing, storage and networking by allowing control of the underlying hardware that is required to operate much of the Internet while reducing the operational costs and complexity. While the concept isn't entirely new the underlying virtualization technology that powers cloud computing can be considered one of the most important breakthroughs in recent times.

In this thesis, I will give an overview of the Infrastructure-as-a-Service industry, starting from the mid 2000s until the present day by providing the market direction, growth strategy and the competitive landscape.

The first chapter is an overview of the cloud industry including a brief history, the classification of the different types of clouds and a high-level overview of the basic building blocks of this technology.

The second chapter of this thesis will focus on proprietary clouds and how they have risen given the wide adoption of cloud computing, the dynamics and drivers of this divide and a deep dive into the product catalog for the large players in this space.

The third chapter on this thesis will focus on Open Source clouds such as OpenStack, CloudStack and will give a deep dive into Rackspace, a leader in the Open Source cloud space.

Finally, the last chapter of this thesis will focus on containers, a new disruptive model that is changing how Cloud computing is delivered and I'll present a case study for one of the leading startups in this space, Docker.

This document should serve as a guide for entrepreneurs and IT managers that are new or intermediate users of the cloud and the main objective of this thesis to help them navigate the different options available and make the best decision based on their needs and requirements. Also, in this thesis we will look at some of the new developments in the industry that are shaping the cloud infrastructure market such as the creation of new platforms that will ease the deployment process and reduce inefficiencies in the cloud that arise from virtualization technology.

Motivation

Prior to joining MIT in the System and Design Management program, I had been working for several Internet companies including Amazon and Yahoo. While I was working on developing the user interface for applications such as the Kindle Cloud Reader and Yahoo! Mail, I had always dealt indirectly with the underlying infrastructure that powered these applications but never managed it directly.

The first time that I was exposed to the concept of the cloud was in early 2007 when I was working as a Software Engineer for Appian, a company that produced enterprise software for Business Process Management. The company was considering offering the software suite in some sort of subscription by using Amazon's EC2. Later, when I joined Yahoo to work on Mail, I found myself working on application software that would run on the company's clusters around the world, but found that there were immense inefficiencies in the usage of this hardware.

It wasn't until 2012 when I joined Amazon, one of the pioneers in modern cloud computing, that the potential for this technology became obvious to me. Any company, small or large has the problem of requiring different types of infrastructure, and these needs vary with time considerably, so the best option to scale with these changes is to use the cloud. I became more interested in this technology for developing new applications and creating new startups particularly for developing markets such as Mexico and eventually started a role as a Product Manager for one of the largest cloud providers in the world, Rackspace. I wanted to use this opportunity to learn and expand my knowledge of cloud computing particularly into its origins, the different product

offerings and strategies followed by the top providers in the space, and finally, evaluate some of the developments that are changing the space.

1 Overview of the Cloud

The precise origin of the term *cloud computing* is not clear, however, it was born through the practice of drawing network systems using clouds diagrams to denote computing, storage and networking systems. Even though the term only became popular around 2008, the theory has been in practice since the first mainframe systems in the 1960s and 1970s with time-sharing computing(1).

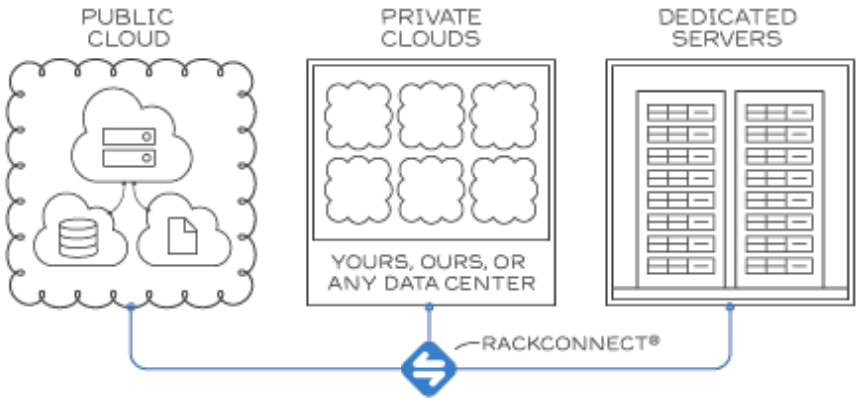

Figure 1: Rackspace Cloud Computing Diagram(2)

The dotcom boom in the 1990s paved the way for advances in telecommunication systems, increased bandwidth, easy access to personal computing devices and the proliferation of high-speed networks. Even with these advances, it wasn't until a decade later that the cloud computing industry gained notoriety.

Cloud computing is used to describe many different services, most commonly through Software-as-a-Service models but also through utility pricing models known as Infrastructure-as-a-Service. Some of the early offerings of modern cloud computing

include Salesforce.com and Amazon Web Services as the lead adopters, but more recently it has become one of the most popular business models for many of the large and established technology companies such as Apple, Microsoft, IBM, HP, Google and Rackspace.

Cloud providers can either charge for their services through utility-pricing plans or also commonly through subscription models. Given that this model usually offers products at lower prices than traditional models in the software industry, it is changing how the value is delivered to the customer and accelerating the rate of creation of new products and startups around the globe. Another benefit of cloud computing services is that it decreases the investment required to run data centers and retain technical staff to maintain and support infrastructure, therefore refocusing these efforts to the creation of value to their customers.

One of the main drivers of cloud computing is the underlying virtualization technology. This technology allows computers to run multiple operating systems and applications with the same hardware simultaneously allowing for more efficient use of the resources(3). Infrastructure-as-a-Service is built on this technology by allowing “multiple tenants” to live in the same server by spreading the costs with other customers. With cloud computing, organizations can use shared resources instead of building and maintaining their own infrastructure. Finally, cloud computing allows for flexible and elastic delivery of IT capabilities using Internet technologies.

1.1 A Brief History of the Cloud

The concept of cloud computing dates back to the 1960s with ideas such as the “intergalactic computer network”, introduced by J.C.R. Licklider while presenting ARPANET through a memo to the scientific community(4). The concept of a cloud didn’t exist at that time, but it was the birth of the Internet and the abstraction of the networking layer that gave birth to the concept of an abstract path between different computer systems. Not much later, John McCarthy suggested that “Computing may someday be organized as a public utility just as the telephone system is a public utility” while speaking at MIT’s centennial celebration in 1961. He also described the potential for this technology as a system where “Each subscriber needs to pay only for the capacity he actually uses, but he has access to all programming languages characteristic of a very large system … Certain subscribers might offer service to other subscribers … The computer utility could become the bases of a new and important industry”(5). These words prophetically describe what we’re experiencing today in modern cloud computing.

In practice, it wasn’t until much after when cloud computing became mainstream when companies such as Salesforce.com pioneered the delivery of enterprise systems. Salesforce delivered their software by using a new approach, instead of charging customers to buy a software license upfront, they would charge through a monthly subscription model. By using this model, better described as Software-as-a-Service (Saas), it became possible to access software hosted remotely through a simple website that could scale depending on the usage and would be billed in similar fashion

to a utility such as an electric or water grid. These companies became proxies of cloud computing to their customers, but didn't offer actual infrastructure that could be used for a generic purpose.

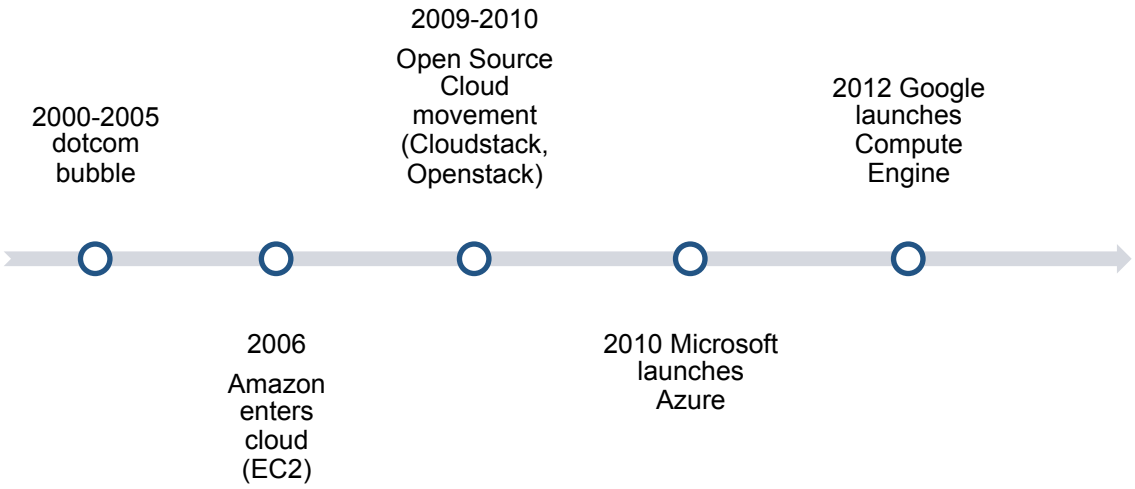

Figure 2: Cloud Computing Milestones

The most important milestone for modern Cloud computing was the launch of EC2, the Elastic Compute Cloud developed by Amazon Web Services, and what later became the blueprint for Infrastructure-as-a-Service (IaaS). IaaS allowed individuals and small companies to rent computers on which they would run their applications by paying only for the resources that they used by the hour with the promise that they would be able to scale in size at any given point with virtually unlimited capacity (6). This capability was already supported internally through the use of virtualization, but traditionally incurred in capital expenditures with large initial investments and lower maintenance fees.

Amazon had an early start in the IaaS market, and they were virtually alone in that market for some time, but it wasn't much after that other companies such as Rackspace entered with similar product offerings, and eventually Microsoft and Google caught up in 2010 and 2012 respectively. The next section of this chapter will focus on the components of the cloud, the history and positioning of the major players in the cloud IaaS market.

1.2 Building Blocks of Cloud Computing

Cloud computing is used to describe a broad range of services and products; however, it is typically composed of three core building blocks, *compute, networking and storage.*

- **Compute:**

Compute is the heart of the cloud offering. It is offered as processing capabilities powered by virtual machines that run on a physical host server. There are several variables that are part of a compute offering: **RAM**, **vCPUs**, **Disk** and **Bandwidth**. Typically, these four variables are combined to produce different types of servers that address different computation workloads. While most of today's compute is achieved via Virtual Machines, there are some offerings available on bare hardware such as Rackspace's OnMetal(7) offering, however, they are still managed through the same methods and the pricing model remains similar. Another form of Compute is available through PaaS offerings that run code without requiring the user to configure or maintain the underlying infrastructure. The latest example of this is with the release of AWS Lambda, a Compute product that is intended to run event-based functions that are triggered under certain conditions and that are able to spinoff the underlying Virtual Machines without requiring the user input. This allows the developer to focus on their business logic, rather than running services. Finally, the third type of computing product is containers. Containers provide the ability to reduce inefficiencies of repetitive underlying operating systems in a virtual machine so that only the core libraries and the application are running on top of a shared layer of the operating system, therefore reducing the overhead necessary to run full images. The last chapter of this thesis will give an overview on Docker, the leading container framework.

- **Storage:**

Cloud providers typically offer storage through different methods block storage, Content Delivery Networks and object storage. Storage can be attached directly to the physical server as in traditional computer architectures, but it can also be attached to

the Virtual machine through Block Storage or Volumes. There are also several tiers of Storage: Edge servers, for fast and reliable static content geographically near to the consumer, Object Storage solutions to allow for simple object storage solutions, and even backup solutions for low-use data that doesn't need to be fetched immediately.

- **Networking:**

Networking encompasses a large number of products including domain name systems (DNS), subnet creation capabilities, sharing of Internet Protocol addresses (IPs), Virtual Land Area Networks (VLANs) and the bandwidth necessary to connect the different pieces of the infrastructure, specifically virtual machines and storage solutions.

1.3 Cloud Classification

There are three types of cloud depending on the needs of the customer: Public, Private and Hybrid. These three types offer different levels of management, security and pricing.

A **public** cloud general consists of resources that are located off-site over the Internet and that are shared with other customers. In a public cloud, the infrastructure is typically multi-tenant, which means that multiple users can share the same underlying hardware or server, running on standard server and resources such as networking, storage, power, cooling and computing that are all shared. The customer usually has no visibility into where this infrastructure is hosted except for choosing the geographic region. A public cloud is typically priced per units of time per Gigabyte of memory and storage.

In a **private** cloud, the infrastructure is effectively dedicated to a particular customer. This type of clouds can be hosted on-premise or at a remote data-center and are typically more expensive to operate as they require dedicated hardware to be secured for a single tenant. Customers of the private cloud have more control over their infrastructure and therefore could achieve their compliance and security requirements.

Finally, a **hybrid** cloud is composed of public and private infrastructure. Typical aspects of applications running on a hybrid cloud is that they require more control over some services while needing bursting capabilities of the public cloud when needed. The public infrastructure is used to manage peak load requirements while the private cloud could be used to meet certification requirements such as those implemented by payment gateways.

1.4 Cloud Application Services

Cloud infrastructure providers are also offering application services in the form of Platform-as-a-Service, Analytics, Mobile services, deployment and management services and even marketplaces for third-parties to offer products and services on top of their infrastructure. They're able to offer these products by hosting them using the core building blocks of infrastructure and creating tools for developers to have easy access to common-usage applications, For this reason, Cloud computing is not only used to describe the core building blocks, but it actually includes a very broad set of services that are built on top of the infrastructure and that can range from platform services to user-facing applications.

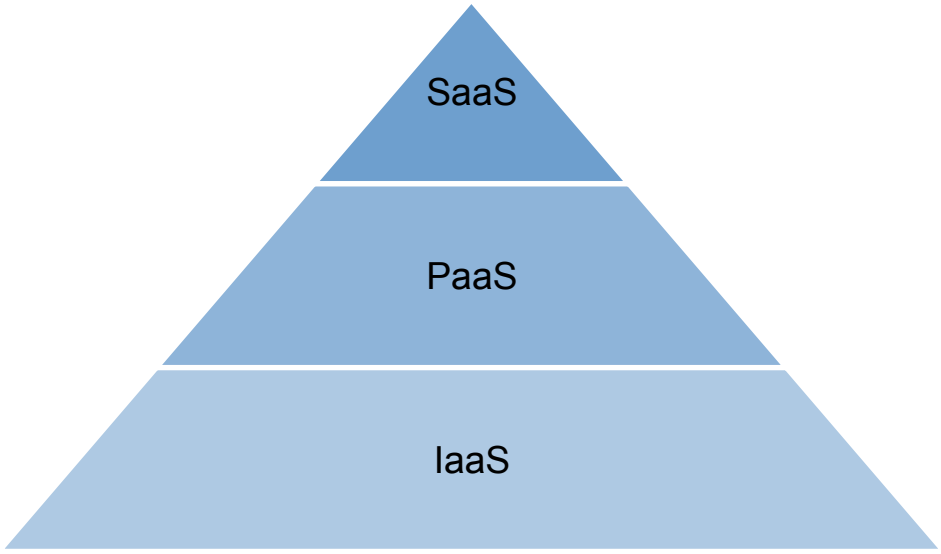

Figure 3: Cloud Computing Stack Pyramid

The leaders in Cloud computing, such as Amazon, Microsoft, Google and Rackspace started by offering a combination of raw cloud infrastructure with the core building blocks (compute, storage and networking), but quickly shifted their focus towards higher levels of the cloud pyramid. The main drivers for this shift are the ability to lock-in customers to their platforms and generate larger profit margins by doing so. This can be clearly observed by the dynamics of the large competitors in the space by lowering continuously the price of the core building blocks in an effort to drive growth in a rapidly expanding market.

The wide adoption of as-a-Service models has given birth to a broad range of services that have adopted the utility-based approach of Infrastructure-as-a-Service models and that have borrowed two main aspects: the ability to scale gracefully without requiring managing the underlying infrastructure and flexibility on pricing modes by the use of subscriptions or pay-per-use models.

Some examples of these types of services include:

- Backend-as-a-Service
- Database-as-a-Service
- Security-as-a-Service
- Delivery-as-a-Service (Fulfillment-as-a-Service)
- Mail-as-a-Service
- Analytics-as-a-Service

1.5 Market Analysis

In 2014, the cloud computing market is composed by some of the largest and best-known companies in the technology industry. Since the introduction of Amazon Web Services Elastic Computing Cloud (EC2) in 2006, the market has become more competitive by the entry of large technology companies such as Google and Microsoft that didn't offer similar services before. Some companies had roots in managed hosting infrastructure such as Rackspace, but most others had roots in the hardware industry such as HP and IBM. In the figure from Gartner(8) shown below, this space can be clearly represented in the Magic Quadrant for Infrastructure as a Service.

Figure 4: Gartner Magic Quadrant for Cloud Infrastructure as a Service

In Figure 4, the Magic Quadrant from Gartner depicts the IaaS market that hasn't diverged too much from the original introduction of AWS. Amazon has been consistently shown as the leader in this space, but more recently, it can be seen closely followed by Microsoft in the upper right quadrant. Below Amazon, a large number of companies in the challenger quadrant diagram depict the rest of the market.

Looking at the similar picture from the perspective of estimated revenue will corroborate the report from Gartner. AWS leads the space by a large factor as they alone get more revenue than its top 3 competitors, Google, Microsoft, followed by

Rackspace. AWS by itself is larger than the rest of its competitors in terms of revenue and market share. By AWS' own estimates, they are adding enough capacity to power Amazon.com from the year 2005, on a daily basis.

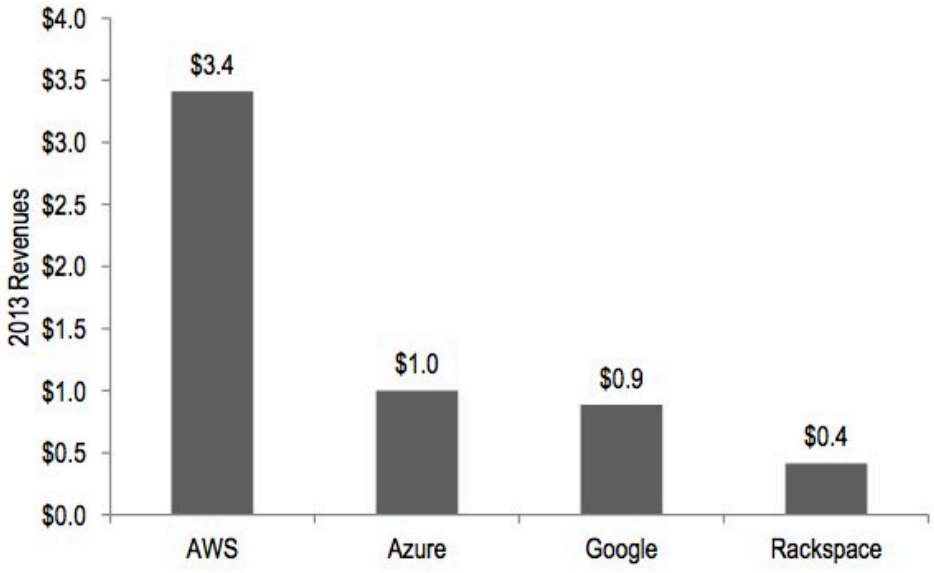

Figure 5: Estimated 2013 Revenue (9)

The cloud Total Addressable Market is expanding rapidly due to a mass migration from traditional data center solutions to flexible systems that the cloud is best suited to offer. Also, IaaS companies have started offering products well beyond commoditized infrastructure in the form of Platform-as-a-Service and other cloud offerings higher in the cloud pyramid where they can benefit from larger margins and greater system lock-in.

The cloud market has also started to diversify into two large segments, managed and unmanaged cloud. A managed cloud is one where the provider will support the underlying infrastructure by offering monitoring, troubleshooting and around the clock

customer service. In an unmanaged cloud, the infrastructure is self-service, unreliable and in any case of a failure, it is the responsibility of the customer to have mechanisms in place to restore their operations.

Figure 6: Gartner Magic Quadrant for Cloud-Enabled Managed Hosting

The Gartner Magic Quadrant for Managed IaaS was introduced in mid-2014, and it shows a different picture of the Cloud market. While the industry analysts, customers and the press used to focus only on product offering and price, in this new type of analysis, emphasis is given to how to effectively support these services. Companies like

Rackspace only offer Managed Cloud infrastructure while Amazon Web Services only offers support through tiered pricing models on top of the base infrastructure. This split was created due to the fact that as the market expands, some customers will prefer to have some level of support, especially in cases where these customers don't have a dedicated staff that is well trained in cloud technologies.

Since 2012, a price war has unfolded in the market where AWS has driven competitors to a race to zero. In the Figure below, this race can be observed with most competitors in the space dropping prices in similar intervals. However, it's very difficult to compare accurately the pricing as each provider's model varies dramatically depending on the type of product. While this race has unfolded over the past two years the trend has started to slow down in late 2014. Therefore we can assume that costs of infrastructure have reached commodity levels and that price differentiation will be less relevant with time. While the market is increasingly behaving like a commodity, each of these platforms is very different from the others and offers little interoperability, and for this reason, this market will not become completely commoditized in the near future.

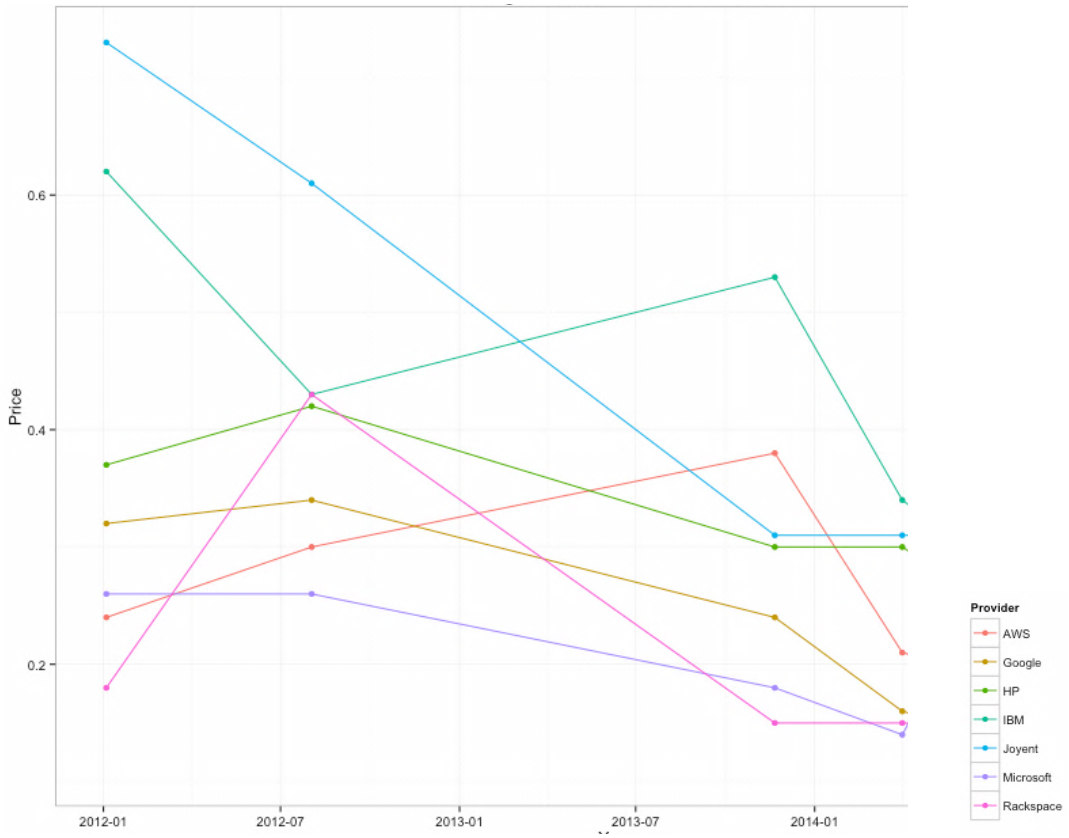

Figure 7: Average Base IaaS Price Over Time(10)

Figure 7 shows the average base price of compute instances over time. As companies have added new offerings to the portfolio this sort of analysis becomes more complex, however, it's a relative simple way of comparing how the compute offerings have dropped prices at similar intervals while this trend has stabilized and slowed down in 2014, we still expect to see price reductions over the next few years, especially as more typically IT-driven workloads are moved to cloud environments.

1.5.1 Survey

In an electronic survey targeted towards the Technical Sales Staff for a Cloud IaaS provider, the participants were asked to rank the most important criteria they thought influenced customers for choosing a cloud vendor. From 27 responses, the survey showed that price is the most important factor for their customers when considering an IaaS provider.

For the question, “When your customers look for new cloud services, what are top things they generally consider? (Rank 1 through 5)”, survey respondents selected price as the most important attribute, while product portfolio was the second highest ranked answer. This shows that the strategy followed by the leaders in this space is resonating with their customers and that the price war that unfolded in 2013 is indicative of the effort from these companies to gather market share.

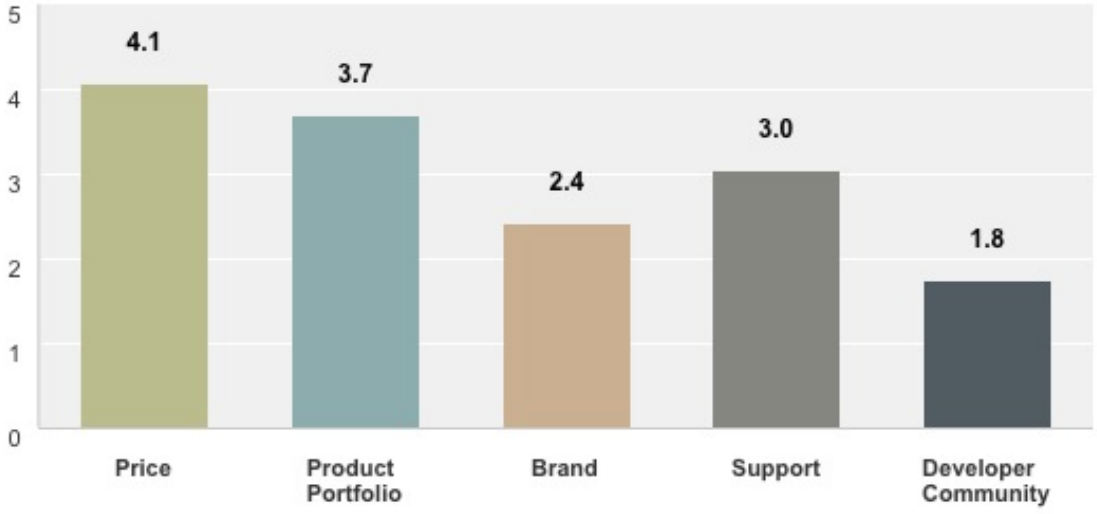

Figure 8: Survey for important attributes that new customers consider when selecting a cloud provider

From the survey results, it's clear that customers are looking for the best product offering at low prices, however, the commoditization of the core building blocks of the cloud have provided difficult to generate profit from these services, so for this reason, the product portfolio has been growing rapidly into PaaS offerings in an effort to deliver added value and generate lock-in, while driving larger margins. In the survey, Support came as the third most important criteria, followed by brand and developer community. In the next section, I compared the top IaaS providers by price per use-case instead of raw infrastructure.

1.5.2 Use-case Cost Comparison

In an analysis conducted by RBC Capital Markets to compare the costs of different providers based on actual use-cases that cover compute, storage, bandwidth and support, RBC showed the overall cost of running a 4GB server and the associated costs(11). While the costs of standalone infrastructure are very straightforward to calculate, as they are equivalent across all providers, the service costs pose a different problem.

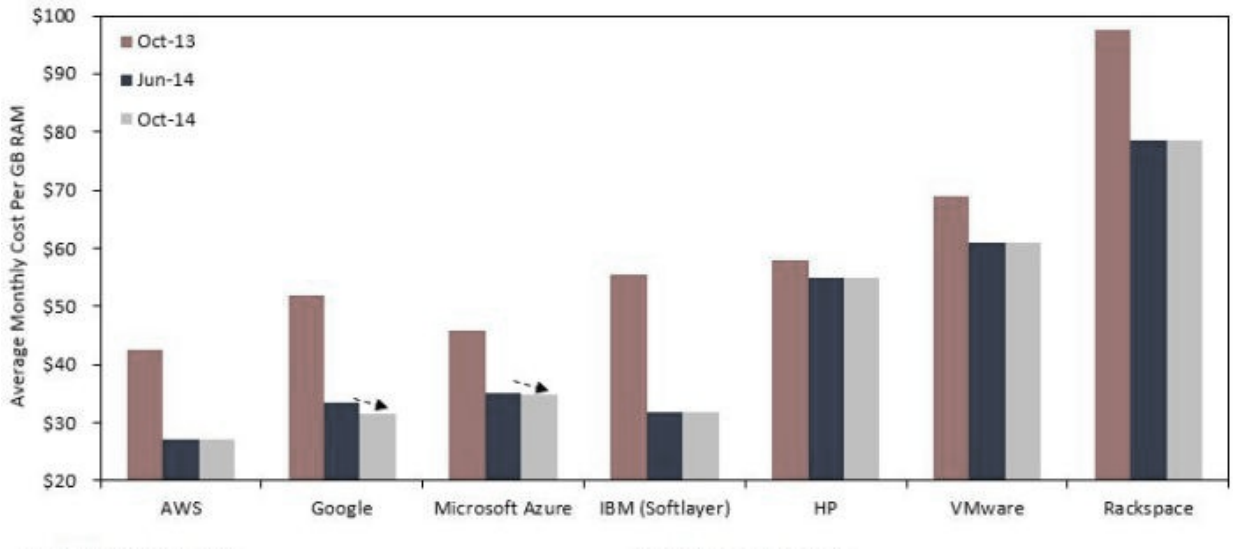

Figure 9: Cloud Use-Case Pricing per vendor(11)

This analysis posed several issues as it didn't represent the actual support and infrastructure costs effectively as every vendor has very different offerings. For the case of Rackspace, the actual costs in my calculations resulted in $59 for the infrastructure and support. This shows that comparing two providers at the surface can be misleading as there are different types of offerings besides the commoditized resources offered underneath.

2 Proprietary Clouds

The Infrastructure-as-a-Service market can be split in two large categories, proprietary and open source. The latter was created as a response to the market dominance by the leader in this space, Amazon Web Services, and with the goal of spreading out the development effort across different companies that individually don't have the resources or the expertise to compete. This chapter will cover the product portfolios for several proprietary cloud platforms and the next chapter will give an overview of popular open source platforms.

Cloud providers that have closed ecosystem are considered proprietary for the purpose of this thesis. Some of these platforms, like AWS, were introduced before there were any open-source platforms, so the distinction didn't exist until 2012 when OpenStack and CloudStack gained popularity. This section will cover the product portfolio and strategy for the three major proprietary clouds, Amazon Web Services, Google Compute Engine and Microsoft Azure.

2.1.1 Amazon Web Services

Amazon Web Services was officially launched in 2006 at the MIT Emerging Technologies conference when Jeff Bezos announced the launch of EC2 as a publicly available service. The forefather of EC2 at Amazon is Chris Pinkman, an engineer in charge of Amazon's global infrastructure in the early 2000s came up with the idea of an "infrastructure service for the world"(6). Since 2006, EC2 has grown to become the

cornerstone of Amazon's Web Services product offering. Since the launch of EC2, Amazon has experienced exponential growth and has launched a vast lineup of product offerings around the Web Services division.

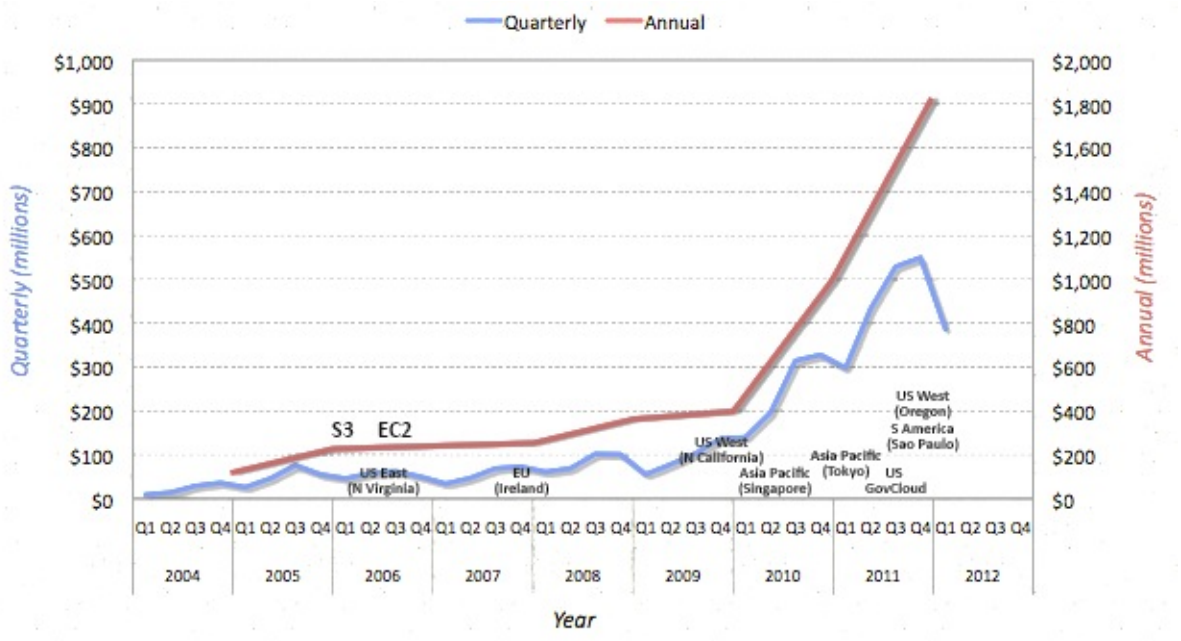

Figure 10: Amazon Fixed Capital Assets 2004-2011

Unfortunately, Amazon doesn't report their financial data and revenue breakdown so there is no clear indication of what the exact revenue numbers are for each of their divisions, however, by looking at their capital assets, total revenue and an estimation of the market, we can get an idea of what the magnitude of AWS is and the extent of their growth.

Largely, Amazon has been able to control the market through network effects, aggressive price drops and by targeting the developer community with a limited free usage tier. According to Cusumano and Suarez, the "Strengh of network effects often tilt

the balance in favor of platforms that can build their installed base faster than competitors”(12). Also, by heavily subsidizing pricing, the company is able to obtain the desired effect in adoption in a more crowded competitive space. As of January 2014, Amazon had reduced the price of their offerings over 40 times.

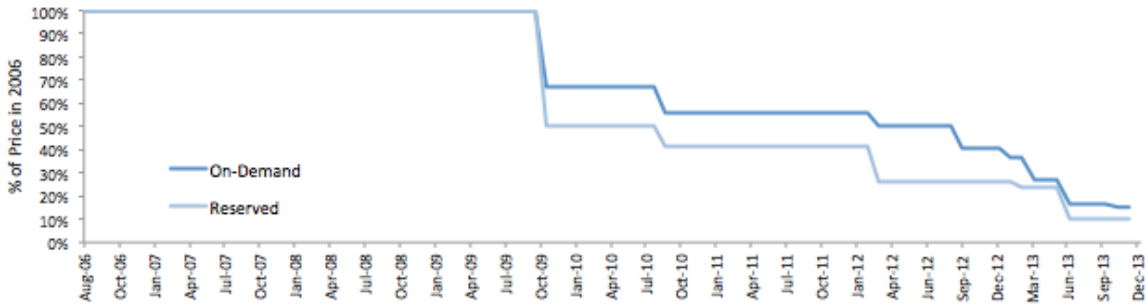

Figure 11: EC2 Price Drop History since 2006

This aggressive strategy has forced competitors to reduce prices at the same time with seemingly coordinated reductions in which the competitors would react quickly and match or lower prices even more. To some extent, the market is following a commoditized product, so Amazon has moved up to different tiers of the cloud stack (PaaS and SaaS) in search of higher margins and lock-in.

Amazon has the largest and most varied product offering of all Cloud IaaS providers. This offering is diverse and quickly expanding to new product categories. In late 2014, AWS announced several products targeted towards developers, specifically around tooling, source code management, and Platform services. This shift shows that the company is looking for new areas to expand their growth while driving higher margin

revenue sources and ensuring lock-in to the platform. Some of the strengths of AWS are the rapid delivery of features and services to their customers, the network effects given their scale, the brand perception and high-profile customers such as Netflix and startups that promote their services.

Some of the weaknesses of Amazon Web Services are that they don't offer support to customers unless they pay considerable fees. Having a closed proprietary ecosystem and their focus on mostly self-service systems and their lack of dedicated hosting services.

2.1.1.1 Amazon Web Services Product Offering

Amazon's product portfolio is organized in 7 categories(13) shown below. AWS started with a small product catalog with basic services around Compute, Storage and Networking, but has expanded their portfolio rapidly into the PaaS and SaaS segment of the cloud-computing pyramid.

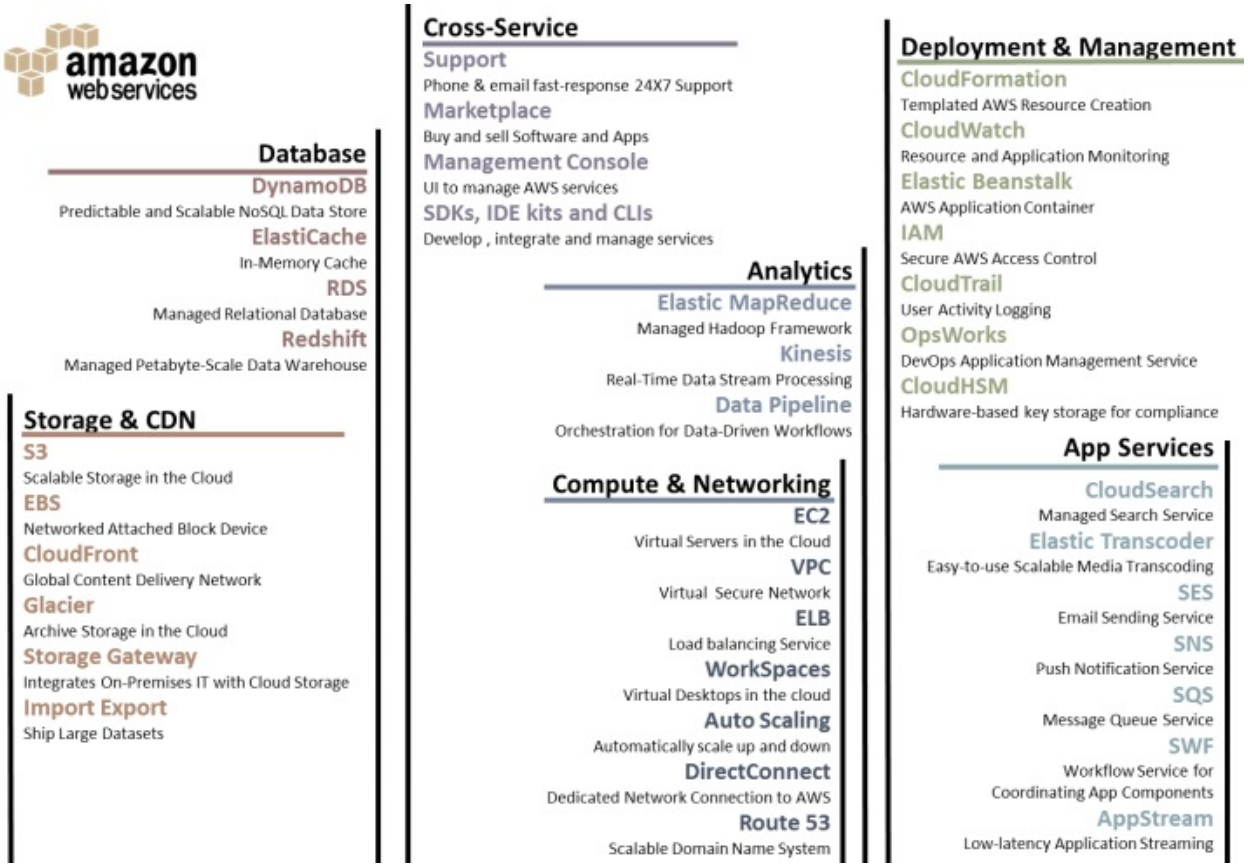

Figure 12: Full AWS Product Catalog (2014)

Compute & Networking:

EC2: Amazon Elastic Compute Cloud (EC2) is a service that provides compute capacity. It is designed to scale with the organization with minimal effort and provides full control of the computing resources. This service is the cornerstone of the Amazon offering as it was the first service that was publicly launched in 2006. Customers are billed only for the capacity that they use. EC2 is offered through a variety of instance types with different specifications.

Lambda: is a compute service that runs code in response to events launched in late 2014. With Lambda, you can execute code (functions), within milliseconds of an event

and it allows developers to operate a back-end service without having to maintain any infrastructure. Initially, these functions will be written in Javascript, with the possibility to run in more languages in the future. Billing for Lambda is metered per request in increments of 100 milliseconds and price also depends on memory allocated per function. Each function also has access to 512MB ephemeral disk space.

Mobile applications, web browsers and connected devices rely on events to run logic and deliver value to the customer only when they require it. Lambda allows the developer to only care about responding to these events instead of the infrastructure around it. The central point of this service is to offer compute capabilities without requiring customers to manage their own architecture and infrastructure to scale with their application. Other advantages include the ability to develop in similar manner on the browser or application and on the web and even share code among them.

Storage & CDN:

Amazon Web Services provides pay-as-you-go data storage services. S3 (Simple Storage Service) is used for content and object storage. A parallel product, Glacier, offers lower cost alternative to be used primarily for backup. Elastic Block Storage (EBS) provides volume storage that is directly attached to virtual machines. Finally, Amazon offers other services such as a Content Delivery Network (CDN), called CloudFront, for edge serving capabilities.

Database:

Amazon offers different types of databases through DynamoDB and RDS (Relational Database Service). DynamoDB offers NoSQL capabilities while RDS is a service used to host typical database solutions such as Oracle, MySQL and SQL Server.

2.1.2 Google Cloud Platform

In 2008 Paul McDonald, the product manager for the Google App Engine announced a PaaS offering that allowed developers to run web application's on Google's infrastructure. This was Google's first foray into cloud computing and PaaS(14). Since the release of App Engine, they have quickly expanded the service's capabilities to add more features around identity, storage, search and other services.

Google didn't enter the IaaS market until much later. In 2013 Google released the Compute Engine to the general public with enterprise Service Level Agreements. Since the launch of these services, Google has followed a similar approach to Amazon and Microsoft, by reducing the overall price of the core infrastructure while releasing services on top of their platform such as their Big Data offering.

One of Google's main attractions is the close-knit integration with Google API's for popular services such as Books, Maps, Places (15). This is the main benefit for adopters of the App Engine since developers can leverage ready-to-use services and integration points.

The Google Cloud Platform is divided in four categories, similar to AWS, but more simplified:

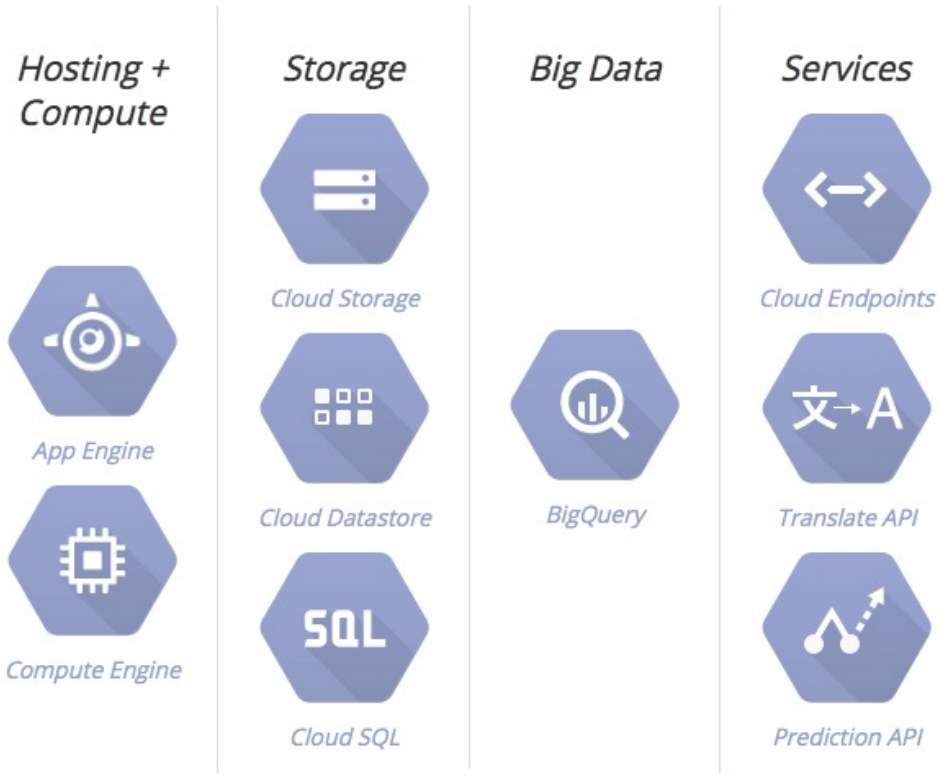

Figure 13: Google Cloud Platform Offering(16)

Hosting and Compute:

Google's Compute offering consists of three main products: Google App Engine, which is a PaaS offering that has gathered some success through it's wide range of APIs. This offering abstracts the underlying infrastructure while only exposing development frameworks. The second compute product for Google is the Compute Engine, which competes directly with Amazon's EC2 with little differentiation. Finally, Google launched in 2014, an alpha preview of the container engine, which targets docker deployments(17).

Storage:

Google offers Databases, block storage, and a Datastore. The datastore offers managed, NoSQL database for storing non-relational data. In terms of product offering, here are little differences between Google's offering and Amazon's in respect to storage.

Big Data and Services:

Google has focused on Big Data analytics services, but also interestingly, translation services, prediction APIs and EndPoints as three scattered tools.

With Google being spread-out across different fronts with Android, Search, Apps, Maps and many more services, it seems as if they're not a strong contender in the cloud IaaS arena, however, given their large scale and sources of revenue, they're able to subsidize this business as internally they would also benefit from having these tools developed.

The company has extensive experience running large and complex applications that run at very large scale such as their search and mail services. Using this experience they're able to bring a solid product offering that is able to compete through continuous price cuts. Similar to AWS, their focus is on self-service applications and they don't offer support with their basic tiers. Finally, the company doesn't have a dedicated hosting offering as they focus only on the public cloud.

2.1.3 Microsoft Azure

Microsoft launched Azure, an Infrastructure-as-a-Service and product platform on the cloud in February of 2010 in 21 countries(18). To date, the company provides one of the largest and most comprehensive product portfolios in the industry, and according to Gartner, Microsoft is a leader in the space while still following Amazon Web Services. The platform was announced in 2008, but it wasn't until 2 years later that it was available to the general public.

Microsoft initially launched with only support for Windows OS and the company poured considerable resources into the development of their platform, however, in 2012 the company started offering a few distributions of Linux, therefore competing effectively with Amazon, Google and Rackspace as the top cloud providers.

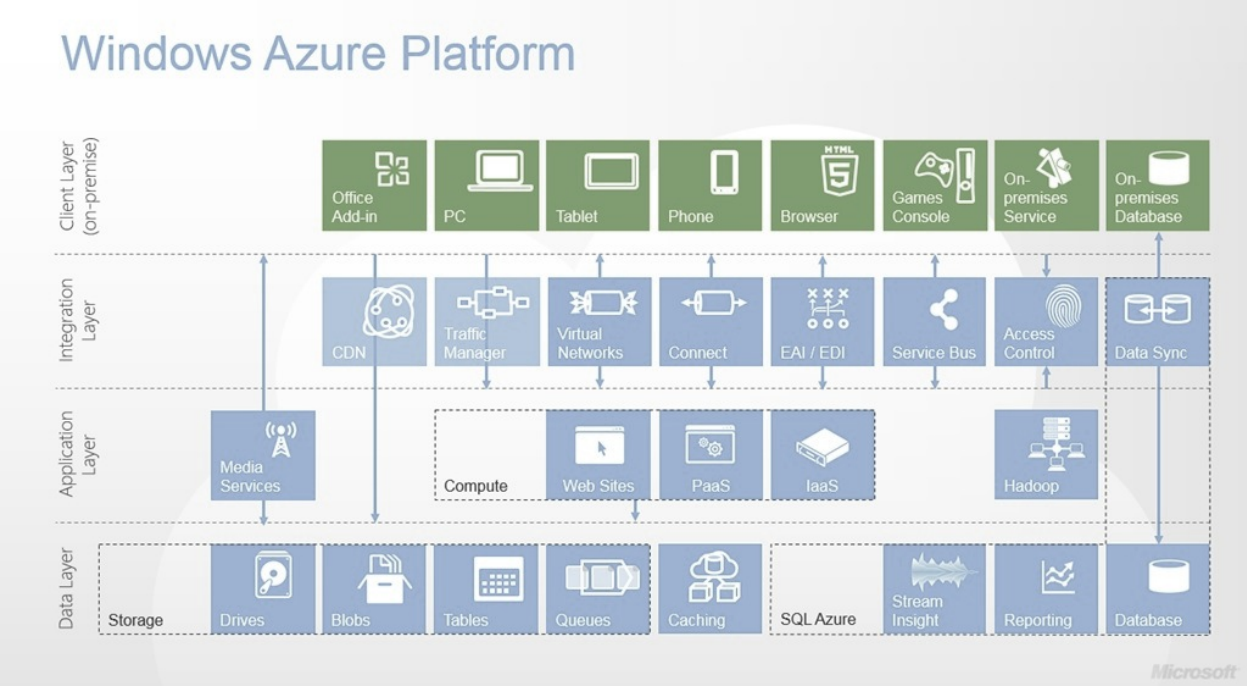

Figure 14: Azure Product Portfolio(19)

Microsoft has made considerable investments in their cloud, and their commitment to cloud computing has grown considerably as shown with the promotion of the leader of their cloud division to the role of CEO, Satya Nadella. Some of the strengths of Microsoft are that they appeal to windows customers given their focus around their platform. Also, they have a solid offering on Active Directory and SQL that has made it easier for typical enterprise companies to adopt their cloud. Finally, they have also participated in the aggressive price cutting strategy that AWS and Google Cloud have led. Some of their weaknesses are that the Windows platform to not be best suited for cloud-centric applications and they've encountered considerable outages that have diminished their reputation.

3 Open Source Clouds

Open sourced alternatives to the proprietary clouds exist today in several forms and provide benefits beyond proprietary clouds. The main benefit of deploying on an open-source cloud is the promise of portability across clouds and consistency in the implementation. This promise has not come to fruition yet, but as OpenStack matures it will become easier to port implementations or even have dual-vendor solutions. The most popular open-source projects are Openstack, Cloudstack, Eucalyptus, OpenNebula, however, these implementations lag considerably in adoption behind proprietary clouds. This section will cover the history, product offering and strategy behind Openstack and Cloudstack.

3.1.1 Openstack

Openstack was created in 2010 in a joint venture between Rackspace Hosting and NASA. The idea behind this project was to share the source code and contribute to the development of two products, compute (Nova), and storage (Swift). As of 2014, over 17,000 developers and 435 companies were involved in over 145 countries quickly becoming the largest open cloud development effort and overtaking competition such as CloudStack, OpenNebula and Eucalyptus. Companies such as IBM, Dell, Cisco, Hewlett-Packard and AT&T are some of the contributors to this effort.

Openstack is divided into several projects that mimic the product offering from proprietary clouds such as AWS and Azure. These software projects are:

- **Compute** (Nova)
- **Storage**: Object Storage (Swift), Block Storage (Cinder) and Database (Trove).
- **Networking** (Neutron).
- **Shared Services**: Identity (Keystone), Dashboard (Horizon), Orchestration (Heat), Data Processing (Sahara), Telemetry (Ceilometer) and Image Service (Glance).
- **Under Development**: Bare Metal (Ironic), Queue Service (Zaqar), Shared file system (Manila), DNS (Designate) and Key Management (Barbican).

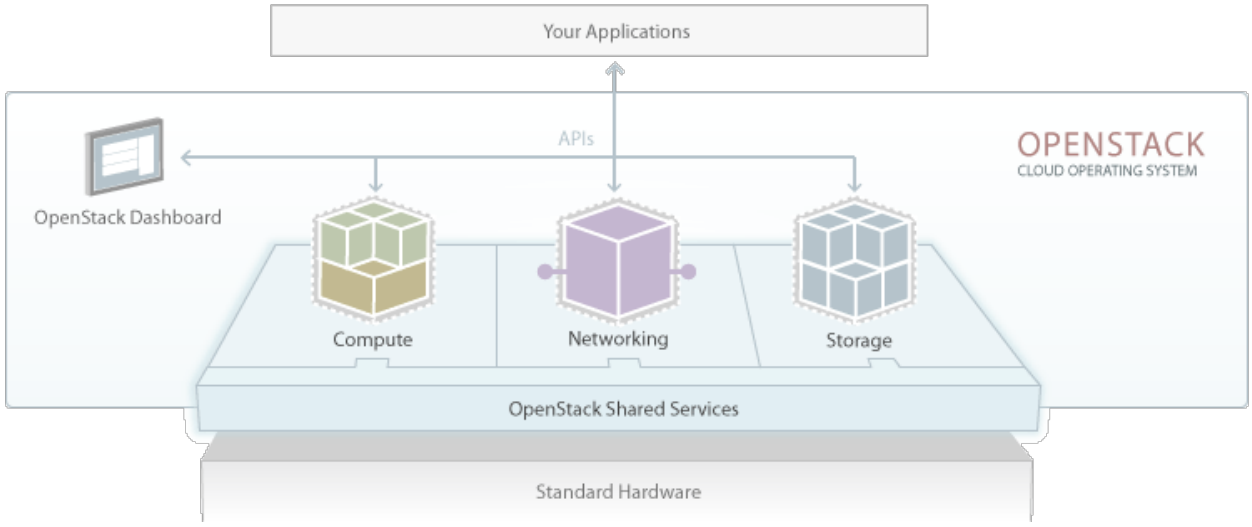

Figure 15: OpenStack Software Services(20)

Since OpenStack is an open-source project with a very distributed development effort across many companies, the community holds several conferences around the world where developers and also where companies show their progress either contributing to OpenStack or their implementation details.

In 2013 the OpenStack market size was estimated around $600 Million dollars, According to the analyst firm 451 Research, the market for OpenStack will surpass $1.7 Billion dollars by 2016.

One of the main promoters of OpenStack is Rackspace, it's founding member together with NASA. Rackspace's cloud offering is based on OpenStack and the company has committed considerable development resources to contribute and operate their cloud infrastructure in open-source.

Hundreds of companies have contributed to OpenStack and it's many projects. As of late 2014, over 7800 contributions were accounted for in the GitHub project.

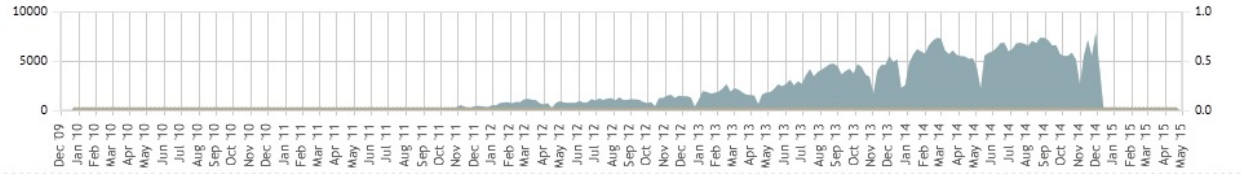

Figure 16: OpenStack Contribution Count Timeline(21)

The top overall contributors for OpenStack are Rackspace, RedHat, HP, Nebula, and some others such as IBM, Mirantis as seen in Figure 16, however, in the past release, HP has taken the stop spot as the largest contributor for the Kilo release (late 2014).

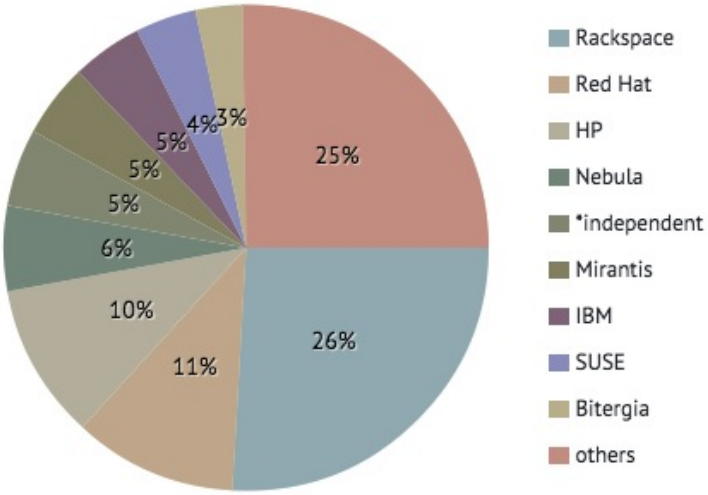

Figure 17: OpenStackTop Contributors(21)

3.1.1.1 Rackspace

Rackspace is considered one of the early entrants in the cloud space. The hosting company was founded in 1998 by Richard Yoo, Pat Condon and Dirk Elmendorf and has grown from being a small, dedicated hosting service to a large multi-national corporation with a strong product portfolio.

From 1998 until 2008, Rackspace was mostly dedicated to providing dedicated hosting services with support services, marketed as Fanatical Support. The company found that many enterprises wanted to have somebody manage and support the infrastructure for them while providing high-availability, redundancy and around-the-clock management. In 2008, Rackspace acquired the startup Slicehost, a cloud entrant that provided a multi-tenant virtual machine based cloud and invested heavily into deploying the first cloud offering for Rackspace. Slicehost provided compute power on

demand in slices of 256MB of RAM, hence the company name. In 2009, Rackspace launched the first generation cloud leveraging the Slicehost acquisition and was one of the early entrants in this market together with AWS. However, due to the entries of Microsoft, Google, HP, IBM and many other technology companies with large budgets to invest in the cloud, Rackspace looked for ways to be able to compete in this market by leveraging a larger community of developers and open source economics. For this reason, developers from Rackspace co-founded OpenStack in 2010 by contributing a Block Storage project code-named Swift, while NASA contributed their compute offering, code-named Nova. Overall, Rackspace continues to be the largest contributor to OpenStack even though in the latest releases, and given the increased popularity of the platform, other cloud providers such as HP have taken the lead as the top contributors.

Rackspace launched a new generation of their cloud in August of 2012 using primarily OpenStack. The strategy of the company was to create an open platform for cloud providers and enabling customers to migrate workloads from different clouds without having to learn new APIs and computing terminology.

Since the first generation release for Rackspace, the company has launched products in Networking, Storage and expanded their Compute offering while providing different tiers of support to customers that still value this service. When the price war unfolded, Rackspace changed strategy and decided to focus on the **managed cloud** market segment in 2014 and stopped decreasing prices at the same time as competitors were doing the same. The goal of this strategy is to discourage independent developers and early startups from signing up, as they tend to have small

deployments, therefore minimum revenue for the company while incurring in high support costs.

As part of the Fanatical Support, Rackspace offers two tiers: Managed Infrastructure and Managed Operations. In managed infrastructure, the support team will help with architecture assistance, code development guidance and launch management at a fee of $50 per month as the minimum. In Managed Operations, Rackspace support manages the cloud for their customers including patching, monitoring, application maintenance. There is a third tier that is know as DevOps, which is a service level where Rackspace support staff will work with the developers using the cloud to provide automation for deployment at every level of the development process, development, testing and deployment to production.

Recently Gartner released a new magic quadrant (seen in Figure 6), in which Rackspace is considered the leader in this market segment. Given this shift on focus, the company is now targeting customers that appreciate reliability, support and architectural guidance on top of pricing. Also, Rackspace considers itself a leader in Hybrid cloud given it’s roots in managed dedicated hosting and the company has pledged a commitment to open standards.

In the Managed Cloud segment, Rackspace is focusing on businesses instead of selling directly to consumers. They are able to target small businesses to large enterprises by leveraging their large enterprise sales team and helping organizations explore, grow and optimize their resources. However, one of the areas that the company struggles to address is its strategy with the developer and student community. While programs like Developer+ exist, the company is not incentivized to grow the

customer base in this manner since the program is not actively marketed at developer events or startup conferences. One of the disadvantages of this program is that developers don't get access to the support organization as regular customers do. This effort has resulted in a decrease in user-acquisition rate for customers, while the revenue per user has increased.

3.1.2 CloudStack

CloudStack is an Open Source Apache project designed to manage cloud infrastructure resources, mainly through management of virtual machines. This project is focused mostly around computing and this is one of the differentiators with OpenStack. The project began in 2008 as a startup called VMOps and it has been available as open source since mid-2010, however, the project has been incubating with the Apache foundation since 2012 after the Cloud.com acquisition by Citrix.

CloudStack is a software platform used for creating clouds on-demand. It provides capabilities for creating virtual machines, storage volumes and networks, in a similar way that OpenStack allows. This platform also allows for the use of underlying hypervisors such as XenServer, vSphere and KVM and it is deployed into a set of servers that manages this infrastructure. The management interfaces that are offered are available through the use of a user interface or a set of APIs(22).

This OpenSource project is also available in Git and offered with an Apache License and it is an actively contributed project with many companies committing resources to the development of the platform. Given the recent popularity of OpenStack,

much of the effort was diverted from this project. For example, OpenStack has been given over 450 stars, while CloudStack had around 183. Finally, CloudStack has 168 contributors to date, while OpenStack has over 400 (23)(24).

With the popularity of OpenStack and the recent reorganization Citrix has taken CloudStack into a new direction. With several key executives leaving the company, the open source project will suffer and likely fall into irrelevance.

4 New Developments in Cloud Technology

Traditionally, the cloud has been powered by virtualization technology, which hasn't changed much since the early 2000s. Incremental improvements such as increased efficiencies of the cloud, live migration, lightweight operating systems and the propagation of IaaS APIs have made it easier to operate a cloud at scale to the point where the cloud has become a sort of a commodity. However, since 2013 a new deployment technology, based on lightweight containers called Docker, has led to a considerable shift towards developer operations and a reduction of the overhead between the application and the underlying hardware. Almost all large cloud providers have pledged that they would incorporate Docker and container offerings sometime in 2014 and 2015. This Chapter will cover an overview of the virtualization technology, how it differs containers and mainly focusing on how Docker implementations have changed how developers work and deploy applications into production environments while achieving even greater efficiencies.

4.1 Virtualization

The idea of virtualization has been in use since the 1960s and 1970s through the use mainframes and time-sharing solutions. Time-sharing allowed for multiple applications or users to share the same underlying hardware in time-slots. This development led to the lowering in the cost of computing as large organizations could

now make better use of the hardware and users would be able to perform computational tasks without needing a dedicated machine(25).

Today, virtualization technology allows sharing resources through pools of CPU, memory, networking and storage to users in the form of virtual machines. The main engine of a virtualized environment is called a hypervisor. The hypervisor is the operating system that runs on top of the hardware that allows the creation and execution of multiple sandboxed or logically separated spaces where operating systems are hosted. There are two types of hypervisors that are defined by how they run on the native hardware: bare and hosted. In a bare environment, the hypervisor is executed directly on the hardware. In a hosted hypervisor, it runs on top of a host operating system. Examples of bare hypervisors are Microsoft HyperV, Xen, VMWare ESX and Oracle VM and in the case of hosted hypervisors we have examples such as VirtualBox, Parallels and Microsoft Virtual PC.

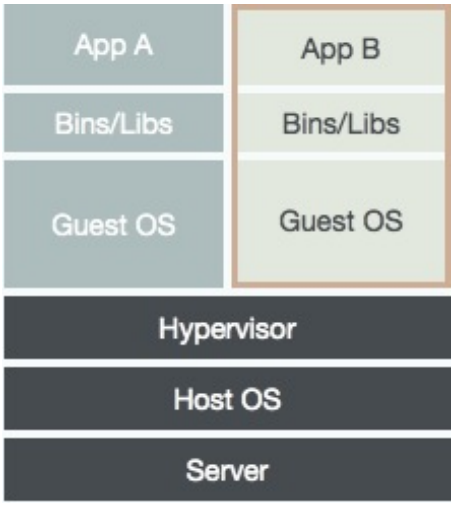

Figure 18: Typical Virtual Machine Stack on Hypervisor(26)

There are several reasons why virtualization technology is so popular and widely adopted. First, the increased resource use efficiency, that allows for more optimal use of computing resources and flexibility for developers to run constrained environments without requiring dedicated hardware or affecting other applications. Second, running in virtualized environments allows for flexibility in case of failure and allows for maximization of uptime, an important requirement of highly available infrastructure. This is achieved through the concept of elasticity, which means that the application environment can scale horizontally with the needs of the business while having the capability of quick disaster recovery. All of these benefits have allowed for more efficient use of resources, therefore a reduction in the cost of operating complex environments.

All major cloud providers use virtualization to power their infrastructure, and while there are more options in terms of the hypervisor operating system and the architectures. There have been considerable improvements in efficiencies for Virtual Machines and creation of lighter Linux Operating systems. For Example, CoreOS is a minimal OS that cosumes 40% less memory than the average Linux distribution that is created to run in clustered environments while running applications in containers(27). Virtualization technology has also improved considerably in efficiency in providing CPU and memory usage optimization while the biggest challenge is around disk throughput especially over the last few years with the adoption of higher speed Solid State drives. To address some of these shortcomings, developer applications have turned to solutions such as running on containers or on bare-metal hardware.

4.2 Containers: A case-study on Docker

Docker is an open source container platform that enables developers to ship, build and run applications in different environments. The goal of this platform is to minimize the number of components required to run distributed applications and to reduce the friction between development, testing and deployment.

This platform is different for virtualized environments as it is composed of only the application and the binaries required without requiring a complete host operating system and kernel that would add overhead, in particular the hypervisor. The main advantage of containers, and docker, is that they don't require a full operating system to execute the application, as they only need the application and the required binaries or libraries. Docker provides a standard, well documented and reliable way to manage the underlying containers and for this reason it has gained popularity since it's creation in 2013.

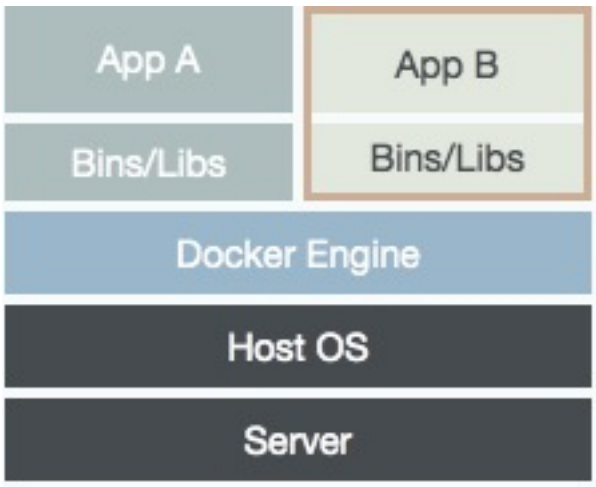

Figure 19: Container Stack(26)

Docker is composed of the **daemon** that sits on the server machine and accepts commands from the docker **client**. The daemon communicates with the underlying Operating System through a **libcontainer**, which is a library that is able to run commands to manage the containers. The user communicates with the daemon through a docker client, which can be a command line or a user interface. Finally, the **registry** is a service is provided by the docker platform on the cloud to host the libraries and the applications through an image. This image can be installed as a container through the docker daemon and these packaged images can be created by any user and shared through the registry.

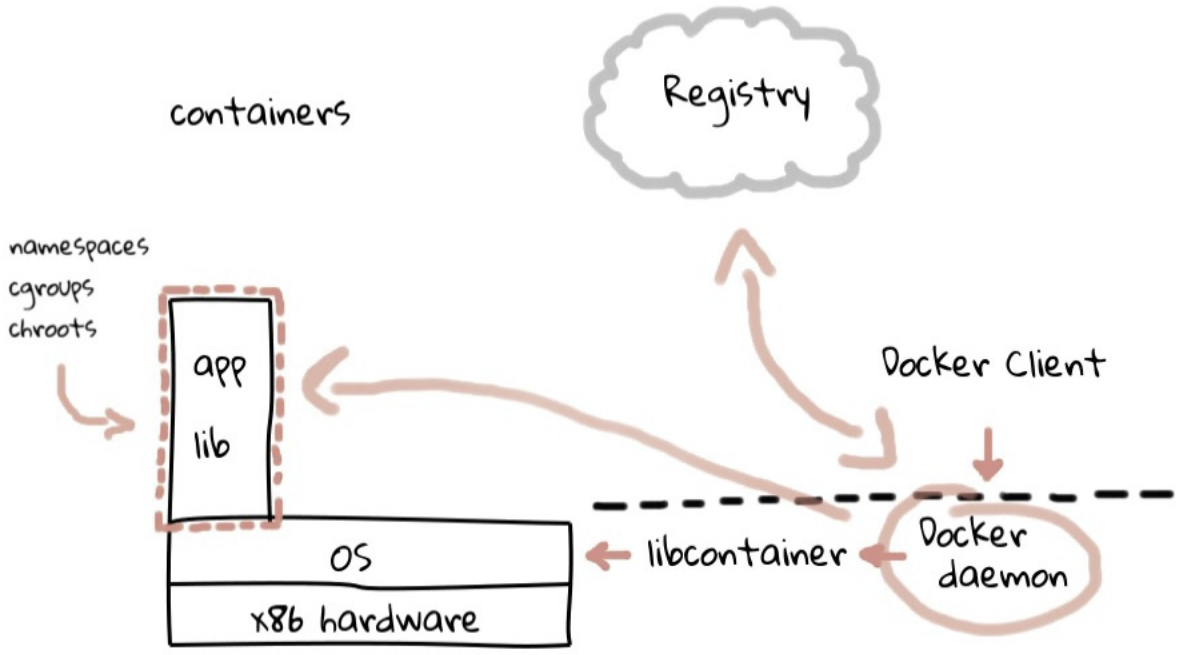

Figure 20: Docker Architecture(28)

Docker (and containers) is considered as a disruptive technology to traditional virtualization as it provides considerable gains in efficiency and resource utilization by removing the need of an operating system running on top of a hypervisor. Given that the project is Open Source major players in the industry such as Google, RedHat, Rackspace and Canonical are adopting it.

The startup company launched in March of 2013 as an open platform and it has gained popularity very quickly in the startup and developer community with already over a hundred meetups (developer gatherings), around the world and also with the support many large companies using it such as Gilt, Yelp, Google, Microsoft and Rackspace to name a few(26).

While Docker has the potential to change the IaaS space, it needs to evolve beyond a packaging framework since developers are still executing the same code in similar infrastructure, which is on virtualized environments such as AWS. One of the top arguments that this is about to shift is with the release of bare-metal servers, which are basically physical server nodes that can be provisioned in the same manner as virtualized environments, however, the cost of each of the bare-metal servers is considerably higher than virtual machines(29).

Early benchmarks have been run on docker containers comparing them to virtual machines and the results have shown that docker provides better startup and stop times. For example, in a test run by developers at a startup company called Flux 7, docker containers would start and stop in less than 50 milliseconds, while virtual machines would take between 30 and 45 seconds to start and up to 10 seconds to stop. In other tests, the memory, storage and CPU benchmark of Docker was similar to that

of KVM and bare-metal hardware. The main difference is in the network latency due to the routing mechanisms in docker that produced delays while executing the performance tests (30). In another test run on containers by IBM, the random IO throughput (IOPS), results showed that during random operations, the performance of virtual machines was considerably slower than to Docker and Native, therefore providing the most benefit (31).

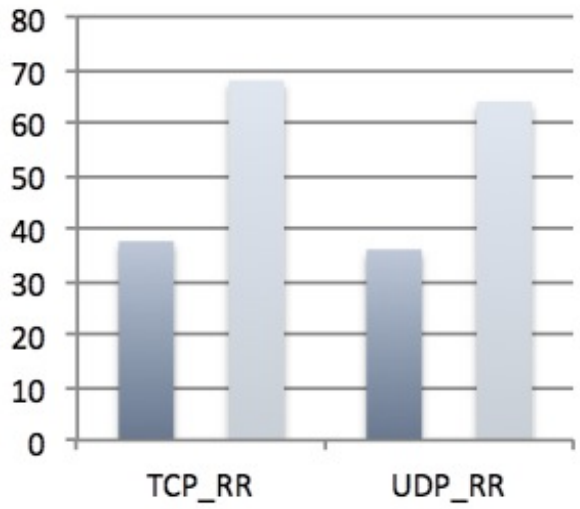

Figure 21: Network latency(31)

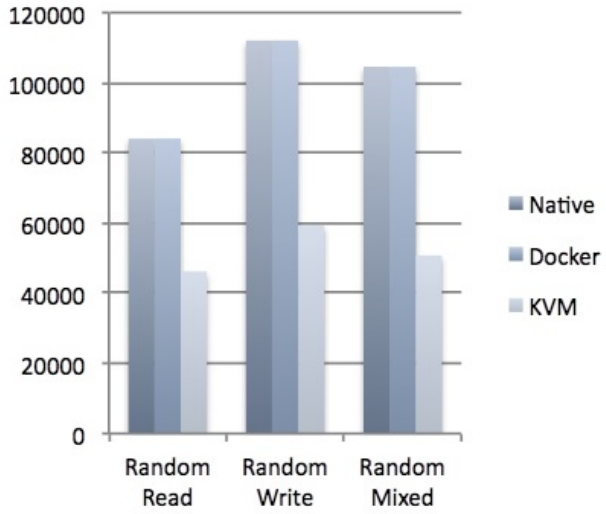

Figure 22: Disk Throughput (IOPS) (31)

Docker is considerably more efficient than virtualization technology as it doesn't require an operating system in each virtual machine. Also, docker containers are smaller so more containers can live in the same hardware. While the network latency could be impacted with the routing of networking packages to the corresponding container, this limitation will probably be overcome once docker evolves in a similar way to how a disruptive technology overtakes the incumbent as it is improved.

Conclusion

The cloud computing industry is rapidly evolving to more than just Infrastructure-as-a-Service. Over the past decades, cloud providers have focused on building a basic portfolio of services that offer the core building blocks of the cloud through compute, networking and storage resources. However, the increase in competition and the entry of powerful technology companies into the space have resulted in a pricing war in which these offerings have become commoditized in many ways. This will lead to some consolidation in the industry through acquisitions of smaller players by large technology companies.

The industry is quickly moving up the cloud pyramid to offer platform (PaaS) and Software (SaaS) services, in an effort to gain from the network effects of building communities and marketplaces around their services. Also, new companies in the space have been able to build their products on top of this infrastructure, filling the gaps that were left by the large infrastructure providers. All of this innovation has led to acceleration in the rate of innovation that has benefited the independent developers and startup companies, which can build their applications on top of all of these platforms and infrastructure at minimal costs and at a faster velocity than ever before.

The cloud market is growing rapidly as large corporations move their workloads off of their own data centers into the cloud. While the cloud has been growing at a rapid pace, most of the traditional workloads are still running outside of the cloud, mostly given that many of these are based on legacy systems and the need for increased compliance and security standards on the cloud. When some of these gaps are filled,

then the bulk of these applications will be moved over to the cloud to gain from the benefits of scalable and elastic clouds. Also, with the creation of open source cloud platforms, this shift will be easier as large companies won't have to commit to a proprietary vendor while having a solid set of features to have more control over their infrastructure.

Finally, recent developments in cloud computing such as the creation of container frameworks such as docker, and abstracted computing frameworks like AWS' Lambda, will result in acceleration in the creation of startup companies. The main benefit of all of these developments is the reduction in the friction between developing applications and making them available in production. Similar to what the Application Stores have provided for the smartphone industry, the creation of development frameworks and the abstraction of the underlying resources have made the most important advances in the Information Technology industry over the past decade.

References

1. Carr N. Cloud computing: BartonPlus [Internet]. [cited 2014 Nov 3]. Available from: http://eds.a.ebscohost.com.libproxy.mit.edu/eds/detail/detail?sid=28027a52-cce8-425b-a6de-1413a8c1b6e9@sessionmgr4004&vid=2&hid=4102&bdata=JnNpdGU9ZWRzLWxpdmU=#db=ers&AN=87323246

2. Hybrid Cloud Computing, Hybrid Hosting by Rackspace [Internet]. [cited 2014 Nov 7]. Available from: http://www.rackspace.com/cloud/hybrid/

3. VMware Virtualization Technology & Virtual Machine Software | United States [Internet]. [cited 2014 Nov 3]. Available from: http://www.vmware.com/virtualization

4. Pelkey J. 2.1-IntergalacticNetwork-1962-1964.html [Internet]. [cited 2014 Oct 27]. Available from: http://www.historyofcomputercommunications.info/Book/2/2.1-IntergalacticNetwork_1962-1964.html

5. Garfinkel S. The Cloud Imperative | MIT Technology Review [Internet]. 2011 [cited 2014 Oct 27]. Available from: http://www.technologyreview.com/news/425623/the-cloud-imperative/

6. Clark J. How Amazon exposed its guts: The History of AWS's EC2 | ZDNet [Internet]. [cited 2014 Oct 29]. Available from: http://www.zdnet.com/how-amazon-exposed-its-guts-the-history-of-awss-ec2-3040155310/

7. OnMetal - Single-tenant, bare-metal Cloud Servers from Rackspace [Internet]. [cited 2014 Nov 13]. Available from: http://www.rackspace.com/cloud/servers/onmetal/

8. Leong L, Toombs D, Gill B, Petri G, Haynes T. Magic Quadrant for Cloud Infrastructure as a Service [Internet]. [cited 2014 Nov 6]. Available from: http://www.gartner.com/technology/reprints.do?id=1-1UKQQA6&ct=140528&st=sb

9. D'Onfro J. Amazon Web Services Market Share - Business Insider [Internet]. 2014 [cited 2014 Nov 6]. Available from: http://www.businessinsider.com/amazon-web-services-market-share-2014-6

10. O'Grady S. The Implication of IaaS Pricing Patterns and Trends [Internet]. [cited 2014 Nov 30]. Available from: http://redmonk.com/sograd y/

11. Darrow B. Confused by cloud computing price comparisons? Here they are in one easy-peasy chart — Tech News and Analysis [Internet]. [cited 2014 Nov 29]. Available from: https://gigaom.com/2014/10/02/confused-by-cloud-computing-price-comparisons-here-they-are-in-one-easy-peasy-chart/

12. Cusumano M, Suarez F. The Role of Services in Platform Markets [Internet]. 2008. Available from:

http://ebusiness.mit.edu/research/papers/244_Cusumano_Platforms%20and%20Services%20Chapter.pdf

13. Varia J, Mathew S. Overview of Amazon Web Services [Internet]. 2014. Available from: https://media.amazonwebservices.com/AWS_Overview.pdf

14. What Is Google App Engine? - Google App Engine — Google Cloud Platform [Internet]. [cited 2014 Nov 30]. Available from: https://cloud.google.com/appengine/docs/whatisgoogleappengine?csw=1

15. Wayner P. Review: Google Compute Engine rocks the cloud | InfoWorld [Internet]. [cited 2014 Nov 30]. Available from: http://www.infoworld.com/article/2614917/cloud-computing/review--google-compute-engine-rocks-the-cloud.html?page=3

16. Google Cloud Computing, Hosting Services & Cloud Support — Google Cloud Platform [Internet]. [cited 2014 Nov 30]. Available from: https://cloud.google.com/

17. Google Container Engine — Google Cloud Platform [Internet]. [cited 2014 Nov 30]. Available from: https://cloud.google.com/container-engine/

18. Hauger D. Windows Azure General Availability - The Official Microsoft Blog [Internet]. [cited 2014 Dec 11]. Available from: http://blogs.microsoft.com/blog/2010/02/01/windows-azure-general-availability/

19. Woody B. Windows Azure – Write, Run or Use Software - Carpe Datum - Site Home - MSDN Blogs [Internet]. [cited 2014 Dec 11]. Available from: http://blogs.msdn.com/b/buckwoody/archive/2012/06/13/windows-azure-write-run-or-use-software.aspx

20. Software » OpenStack Open Source Cloud Computing Software [Internet]. [cited 2014 Dec 6]. Available from: http://www.openstack.org/software/

21. Stackalytics | OpenStack community contribution in all releases [Internet]. [cited 2014 Dec 11]. Available from: http://stackalytics.com/?release=all

22. Apache CloudStack FAQ [Internet]. [cited 2014 Dec 30]. Available from: http://cloudstack.apache.org/cloudstack-faq.html

23. apache/cloudstack · GitHub [Internet]. [cited 2014 Dec 30]. Available from: https://github.com/apache/cloudstack

24. openstack/openstack · GitHub [Internet]. [cited 2014 Dec 30]. Available from: https://github.com/openstack/openstack

25. Brief History of Virtualization [Internet]. [cited 2014 Dec 16]. Available from: https://docs.oracle.com/cd/E26996_01/E18549/html/VMUSG1010.html

26. What Is Docker? An open platform for distributed apps [Internet]. [cited 2014 Dec 16]. Available from: https://www.docker.com/whatisdocker/

27. CoreOS is Linux for Massive Server Deployments [Internet]. [cited 2014 Dec 17]. Available from: https://coreos.com/

28. Weissig J. Sysadmin Casts - Episode #31: Introduction to Docker [Internet]. [cited 2014 Dec 30]. Available from: https://sysadmincasts.com/episodes/31-introduction-to-docker

29. Shepherd D. Evolution of Docker and Its Impact on AWS - i build the cloud [Internet]. [cited 2014 Dec 28]. Available from: http://www.ibuildthecloud.com/blog/2014/08/12/evolution-of-docker-and-its-impact-on-aws/

30. Hussain A. Performance of Docker vs VMs [Internet]. [cited 2014 Dec 17]. Available from: http://www.slideshare.net/Flux7Labs/performance-of-docker-vs-vms

31. Felter W, Ferreira A, Rajamony R, Rubio J. An Updated Performance Comparison of Virtual Machines and Linux Containers. 2014.

Appendix A – Survey Template

Name

What are the top features that your customers would like Cloud Servers to support that are not offered today?

Feature 1

Feature 2

Feature 3

Feature 4

Feature 5

Are these features blocking them from using the Vendor X Public Cloud?

Yes

No

When your customers look for new cloud services, what are top things they generally consider? (Prioritize the list below)

Price

Brand

Support

Developer Community

Product Portfolio

Are there any other factors that are stopping new customers from using our cloud?

Add notes or comments here