Athiphan Chunmano, 123RF

Overview of cloud platforms and appliances

Obscured by Clouds

The current trend toward cloud computing is obfuscated by a cloud of buzzwords and acronyms. Pushing the buzzwords aside, we take a look at the nitty gritty of the current crop of offerings.

By Jörg Fritsch

More or less every virtualization provider defines the term cloud to match their own offerings. These offerings are then promoted as Cloud Services or Cloud Servers. Public Cloud, Private Cloud, Hybrids, Software as a Service (SaaS), Platform as a Service (PaaS) – the list goes on and on.

In this article, we will attempt to peek through the mist of terminology and provide readers with an orientation on the topic. We will provide an overview of the properties of real-life cloud offerings that can be found on the Internet.

Definition and Features

Administrators typically will not be interested in a theoretical definition of cloud computing. Instead, they will want to know two things: Do I already have something similar to a cloud in my server room or data center that I can use? How can cloud computing supplement my existing technology and IT platforms and help me solve current and future problems in my environment?

On the one hand, clouds can help support new business models and services for consumers and startups that previously required a huge amount of effort or enormous financial risk. For example, if you have programmed a new web application in Ruby, you can simply launch it in the cloud; if it takes off, your scaling options are virtually unlimited, and you can add Content Delivery Networks (CDN). As an example, SlideShare [1] integrates

Table 1: Admin Cloud Index for Various Providers

Product	Static IP Address (0.2)	Software Appliances and Images Used (0.4)	Web Storage (0.2)	CDN (0.2)	IPv6 (0.2)
AWS EC2	No (elastic IPs)	Yes	Yes	Yes	No
Rackspace	Yes	Yes	Yes	Yes	No
Kamp Virtual Core	Yes	Restricted	No	No	Yes
Strato MultiServer	Yes	No	No	No	No
Microsoft Azure	No (similar elastic IP)	No (no MS products)	Yes	Yes	No

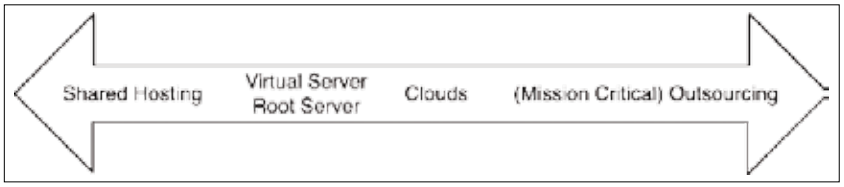

Figure 1: The continuum between simple servers and full-fledged clouds.

document downloads and Flash file hosting with Amazon S3 and CloudFront Services. But even if your blog suddenly takes off and becomes a global event, you can migrate it to the Amazon or Rackspace CDN with just a couple of clicks – or so says the theory. The cloud means a technological revolution that requires major rethinking to put it to optimum use. At the end of the day, clouds are just the ongoing development of virtualization technology. Providers see themselves somewhere between shared hosting and outsourcing (Figure 1). In an ideal world, cloud computing would free administrators from the hardware headaches (e.g., scalability, availability, maintenance contracts) in a geographical sense, giving them time to get on with running their applications. Internet platforms and server rooms currently on the LAN would be partially or fully virtualized and then run on the cloud provider's technology platform.

The Mother of All Clouds

When comparing cloud computing platforms, the reference is normally Amazon's Elastic Computing Cloud (EC2) [2] by virtue of its pioneering role (Figure 2).
Amazon's Web Services comprise virtual servers (EC2), a web-based storage service (S3), and a CDN (CloudFront), all of which are fully integrated and available as self-service features. By credit card, you can pay as you go only for the resources you actually use. For example, you can configure and launch one or more virtual servers in the cloud and pay for them only until you delete them again. The price for a couple of hours' use of a virtual server will typically be far less than a dollar. Self-service is an important component and not restricted to the web GUI. Larger automated applications (e.g., SaaS third-party business models) can use both the GUI and documented APIs. Some cloud customers see the APIs as an easier approach into the cloud and back out again. Theoretically, it should be possible to develop software that uses the programming interface to copy content from Amazon to Rackspace and thus avoid dependence on any one service provider. Emerging standards and API frameworks for cloud computing, such as OpenNebula [3] or Deltacloud [4], are a big help.

The Admin Cloud Index

The *Admin* Magazine Cloud Index (AMCI) evaluates 10 differently weighted features of commercial public clouds. In the AMCI, Amazon's Web Services (EC2, S3, and CloudFront) serve as the baseline, with an index value of 100. More innovative services can achieve scores of more than 100, and less complete services will obviously score lower (Table 1). The index doesn't say much about the quality of the tested service – after all, there are no benchmarks – but it does tell something about the "cloudiness" of the offering – that is, whether it covers the full spectrum of options (see Figure 2).

Highs and Lows

In contrast to the Amazon cloud, Rackspace has a clear-cut GUI that is free of legacy restrictions (Figure 3) [5]. Because of the user friendliness of the Rackspace Cloud and the static IP addresses that any Rackspace Cloud Server is assigned until you delete the server, the Rackspace Cloud scored 112 in our index.
Provider Kamp doesn't advertise its Virtual Core product as a cloud and doesn't claim to offer cloud services [6]. Surprisingly, the Virtual Core Web GUI (Figure 4) is based on the Typo 3 Content Management System and not on a web application framework such as Ruby on Rails.
Virtual Core doesn't currently implement some of the features offered by cloud services, but it is more innovative in some other areas. For example, you can easily set up a software appliance with IPv6 addressing on the Virtual Core platform within about four minutes.
Currently, Virtual Core lacks sophisticated self-service features; however, administrators can contact support if they have special requirements, such as additional VLANs for cluster heartbeats or special Service Level Agreements (SLAs).
In comparison, Amazon EC2 offers a standard SLA that guarantees

Pay As You Go (0.2)	Usability (0.2)	Self-Service Scalability (0.2)	Self-Service Scope (0.4)	API (0.2)	Hardware Features (e.g., NICs, VLANs) (0.2)	Total
Yes	-	++	++++	Yes	+	100
Yes	++	+	+++	Yes	-	112
No	+	+	++	No	++	67
No	+	-	+	No	-	22
Yes	+	-	++	Yes	-	61

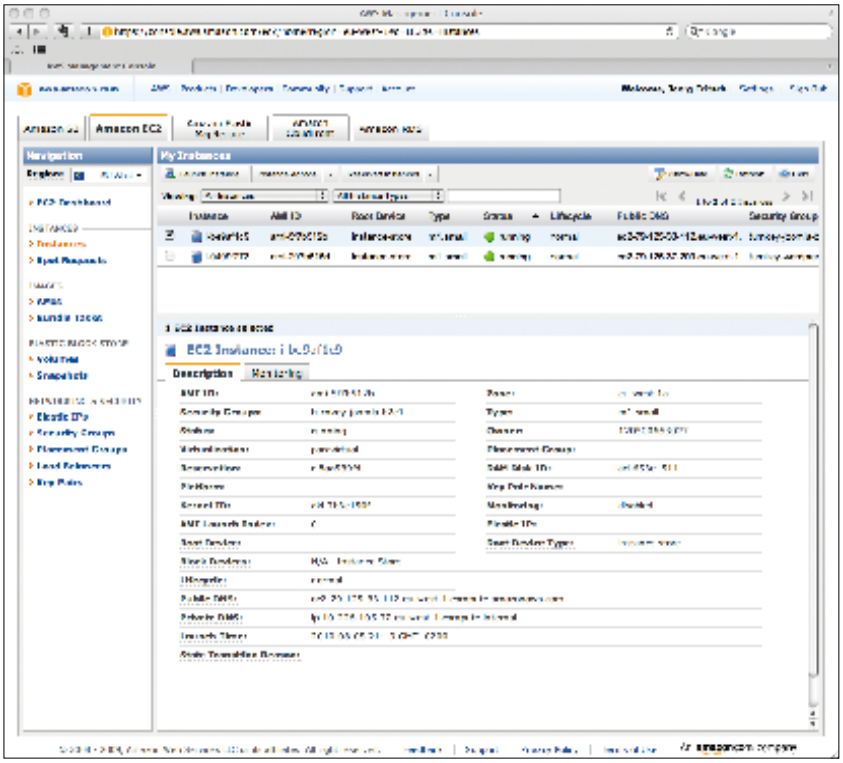

Figure 2: The Amazon Cloud. The mother of all clouds sets the bar for all others. Notice that the AWS GUI is slightly long in the tooth.

99.95 percent availability without any restrictions. Rackspace offers 100 percent but does not count planned maintenance with a notification period of 24 hours. Basically, Rackspace's Cloud SLAs only apply to power or heat issues. The Strato MultiServer didn't fare quite so well in our index. Although the product is advertised as "Dedicated Cloud Hosting," it is not a genuine cloud according to the *Admin* magazine definition [7]. Instead, it refers to a dedicated server on which the Xen hypervisor is preinstalled. The product itself is eminently usable, but has very little to do with a real cloud, which explains why it only scores 23 in the AMCI.

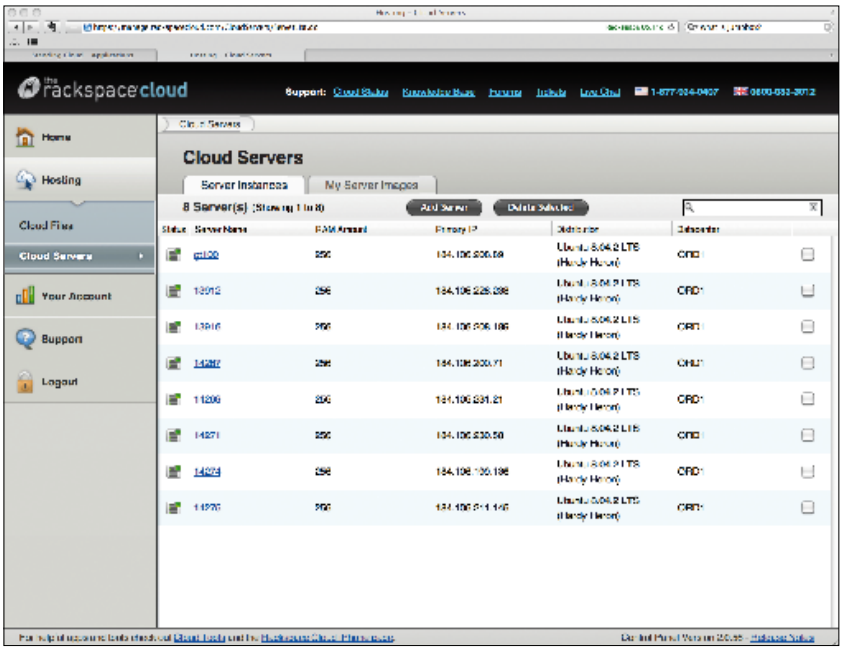

Figure 3: The Rackspace Cloud. Despite the minimalist GUI, you can achieve much with just a couple of clicks.

MS Azure

Microsoft has its own special view of the cloud world. The Microsoft Azure Online [8] Cloud Platform didn't fare too well in our index because its focus is on (Microsoft) applications and services and not on providing virtual servers.

MS Azure is a cloud operating system on which users can run Microsoft, PHP, and Ruby on Rails applications. Rather than managing the virtual server, you manage the service that you provide (or the application). Of course, Microsoft relies totally on products from its own universe – MS SQL Server, Dotnet frameworks, Live/Presence technologies, and so on – which are tightly integrated to the exclusion of all others. For example, you can't use Azure to run an Exchange or OCS server in the cloud. Microsoft offers a different product for this (Business Productivity Online), which is more in the vein of Application Hosting (i.e., SaaS).

Even Microsoft critics have to admit that the vision behind Azure is seminal and future-proof. Administrators don't want to run clouds just for the sake of doing so; instead, they want to launch applications in them. They view maintenance (installing patches and so on) on a virtual operating system platform as an unnecessary overhead.

On top of this, every application on an Azure Cloud operating system has its own IP address, a concept that has potential for terminal devices (hosts, laptops) on IPv6. On the other hand, a service- or application-oriented technology like this is very close to shared hosting.

The developers of the Amazon Cloud recently founded Nimbula (whose largest investor is VMware) and announced their own Cloud OS [9]. What the Nimbula Cloud OS will offer and how, and whether it will be similar to the MS Azure Cloud OS, is still uncertain. Because both products aim to be cloud operating systems, it

is quite possible that they are based on similar visions.

Operation Questions

VMware has an advantage in the virtualization market because of its history [10]. Currently VMware accounts for the largest share of the market, with around 68 percent of the end users.

Most VMware administrators would thus instinctively want to create a cloud with the vCenter Lab Manager, because it is the commercial product that comes closest to being a cloud. Having said that, commercial clouds typically aren't part of commercial products but are based on the community version of Xen or on KVM. If the inside of the cloud depends on a single vendor, likely the flexibility of self-service and pay as you go will be difficult to offer to end users. Linux certainly gives providers the flexibility and economy they need.

Additionally, clouds have a couple of technical (and some less technical) requirements of their users. Whether a cloud offers static IP addresses or not, applications that you run in the cloud should always be designed in such a way that DHCP or NAT doesn't bother them.

Even though innovative cloud providers (e.g., Rackspace or Kamp) support static IP addresses that do not require NAT, you shouldn't rely on this, especially if you might need to migrate between clouds some time in the future, which would automatically change the IP addresses.

No problems are expected with Linux, but Microsoft products are a different case. Any administrator who has tried to modify the IP address of an Active Directory Domain Controller will be aware that this is not a trivial task. Even with the less invasive changes caused by the cloud – such as one part of the infrastructure (or an application) running in your own server room and the others running in the cloud – you must take several obstacles in stride when using Microsoft technologies. Fortunately, Microsoft has seen the issues and has announced changes through blogs and at the Next Generation Active Directory (NGAD) developer conferences, for example. NGAD will make Microsoft Windows, including all of its components (SharePoint, Exchange, OCS, etc.), as fit for the cloud as Linux systems.

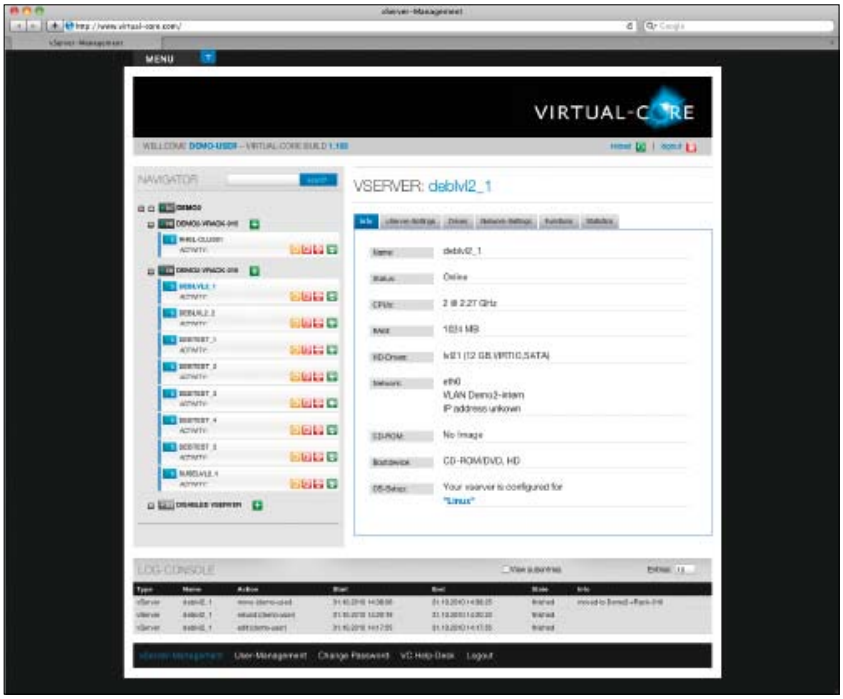

Figure 4: Kamp Virtual Core (developer version). Kamp plans to release its English product in the first quarter of 2011.

Right now, the cloud and cloud-based services (such as Software Appliances, virtualization, licensing, and so on) appear to offer a historic opportunity for Linux.

Cloud Fuel

Most cloud service providers offer launching new virtual servers from preconfigured images (Ubuntu, CentOS, etc.) or from ISO images that you insert into the virtual server for the install. This gives you the bare

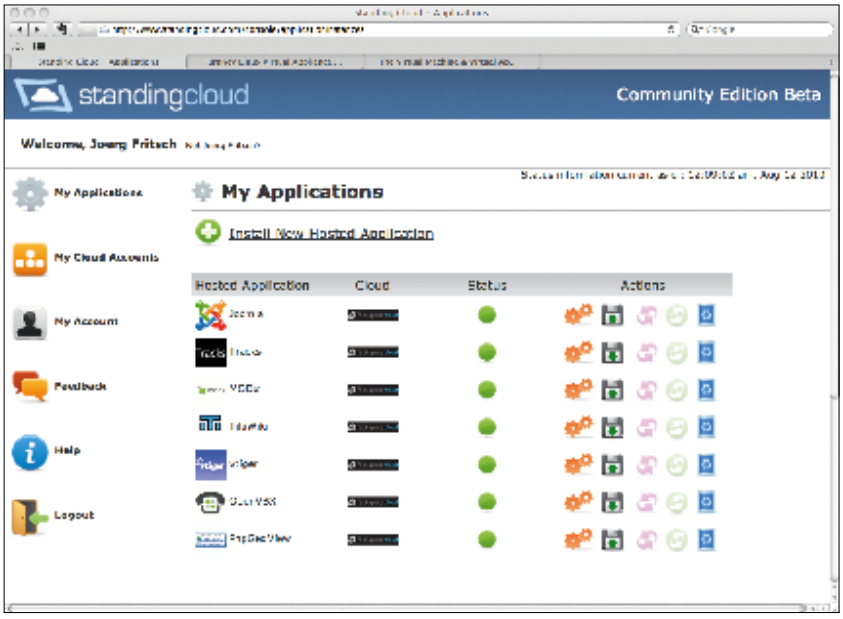

Figure 5: Standing Cloud managing seven software appliances all of which are running in the Rackspace Cloud.

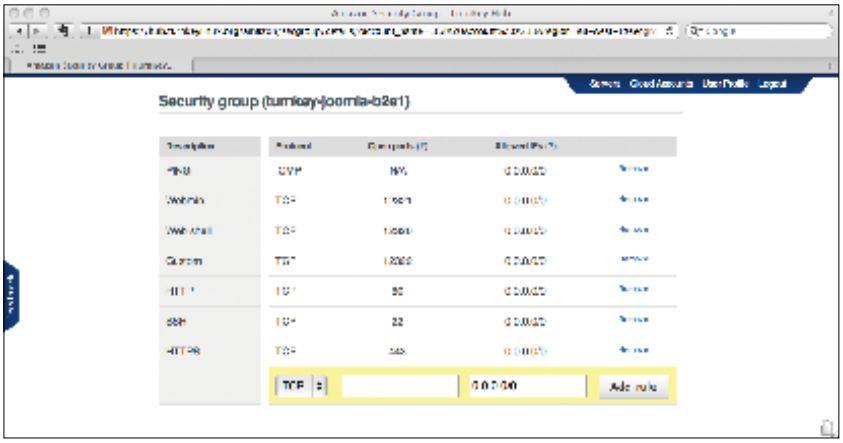

Figure 6: The firewall settings for the running software appliances can be modified in the TurnKey Linux GUI.

bones of a basic platform on which you need to take steps to ensure security and keep the system updated. This means you must know how to install the required applications.

The Amazon Cloud has accumulated more than 7,000 Amazon Machine Images (AMIs) in recent years, and you can use them to set up new servers in the cloud. The AMI library has grown historically, like the whole Amazon Cloud platform, and is the kind of jungle you might expect. The quality of the images differs, and you have no way of knowing where they came from and what bugs they possibly contain.

In these conditions, providers like JumpBox, TurnKey Linux [11], or Standing Cloud (Figure 5) [12] can enter the market. They offer image libraries and preinstalled applications (i.e., "software appliances") that users can launch in the cloud directly from the vendor's website.

Unlike AMIs, software appliances have quality assurance and include firewalls and update vectors that reflect the applications in question. Turnkey Linux and Standing Cloud are strictly open source appliances, whereas JumpBox [13] also offers commercial applications.

The differences between the providers are huge right now. Some only support the Amazon Cloud. Others leave it up to users to handle operating system, software, and security updates. Good software appliances should give the end user an easy option for running them (e.g., in VirtualBox) on their own LANs, while being able to migrate to the cloud and integrating more or less seamlessly when the need arises. Although this sounds complex, in practical terms, it means the ability to modify the firewall ruleset or the SSH keys for the software appliances either in the provider's GUI or in the cloud GUI (Figure 6). Currently, TurnKey Linux does this best, but it also has the smallest selection of software appliances, numbering just 49. TurnKey Linux is currently working on a migration path from a software appliance to physical hardware in your own data center and from a software appliance as a virtual server in your own data center to a software appliance in the cloud; this plan is exemplary and doesn't force the end user down a path they might not be able to back out of.

In comparison, Standing Cloud supports most clouds and makes it easy for customers to migrate, offering the best pricing as well. Standing Cloud customers can, for example, back up a software appliance currently running on EC2 then relaunch it on Rackspace or GoGrid. Most version 1.1 JumpBoxes also have the option of writing an off-site backup (e.g., for Joomla, Cacti, or Drupal systems) to the S3 cloud.

Table 2 provides an overview of the four software appliance providers

Info

[1] SlideShare: [http://www.slideshare.net]
[2] Amazon EC2: [http://aws.amazon.com/ec2/]
[3] OpenNebula: [http://www.opennebula.org/]
[4] d-cloud: [http://deltacloud.org/]
[5] Rackspace Cloud: [http://www.rackspacecloud.com]
[6] Kamp Virtual Core: [http://www.virtual-core.de/]
[7] Strato Dedicated Cloud Server: [http://www.strato-pro.de/de/multiserver/]
[8] Azure: [http://www.microsoft.com/windowsazure]
[9] Nimbula: [http://nimbula.com]
[10] Virtualization market share: [http://www.cio.com/article/586713/]
[11] TurnKey Linux: [http://www.turnkeylinux.org]
[12] Standing Cloud: [http://www.standingcloud.com]
[13] JumpBox: [http://www.jumpbox.com]
[14] JeOS: [http://wiki.ubuntuusers.de/Jeos]
[15] SUSE Studio: [http://susestudio.com]
[16] rPath rBuilder: [http://wiki.rpath.com/wiki/rBuilder]

Table 2: Software Appliances

Vendor	Download	Freely Copyable	Cloud Integration	Supported Clouds	(Auto)Update of OS and Application	Quality	Number of Appliances
TurnKey Linux	Yes	Yes	Tight	EC2	Yes	High	49
Standing Cloud	No	No	Little	EC2, Rackspace, GoGrid, Slicehost	No	High	54
AMIs	No	Partially	Not consistent	EC2	No	Low	7,000+
JumpBox	Yes	No	Untestable	EC2	No	Medium	60

we tested and what you can expect of them. Administrators who are not currently scouting for a cloud are advised to set up an account with TurnKey Linux, download a couple of appliances, and experiment on their own LANs to gain some experience. International Data Corporation (IDC) evaluated the total market for software appliances (not just in clouds) in 2010 at EUR 225 million (~US$ 303 million) and sees it rising to EUR 2.3 billion (~US$ 3.1 billion) by 2014. The last time a market grew 100-fold was shortly before the dot-com bubble burst, and right now this market is attracting surprisingly few startups (compared with the positive forecasts).

Conclusions

"Cloud" is not just a new word for IT outsourcing. Clouds will dramatically change the way we experience IT and the architecture it uses in the next few years. Whether IDC's relatively optimistic forecasts actually come true, clouds will not go away overnight.

Right now, cloud builders are, without exception, relatively new names in the computer industry: names such as Amazon, Rackspace, Kamp, Nimbula – and not IBM or Oracle. Whether new major players evolve or the big fish end up swallowing the small fry, the future developments in the world of cloud computing will be exciting to see. ■

"Juice": Do-it-Yourself Software Appliances

Admins and developers can build their own software appliances, but this will involve some rethinking. Software appliances are designed for applications and not for hardware. Just Enough Operating System (JeOS, say "Juice") is the buzzword [14].

The developers took into consideration that their software appliance might need to be virtualized (although this is not a requirement) and took a hardware-agnostic approach to development. Performance is achieved by a lean basic operating system platform and by optimization in the cloud. JeOS developers don't attempt to make the most of any specific SCSI driver (or even ASICS). Instead, JeOS does without anything that could get in the way of virtualization and portability between individual clouds.

In its image collection, TurnKey Linux offers a 68MB JeOS Ubuntu image that any registered user can download and use as the basis for their own applications in the cloud. Some Linux distributors also offer bespoke services that let users more or less automate the process of putting together minimal or customized distributions. Examples of this are Novell's SUSE Studio [15] or rPath's rBuilder [16].