CLOUDADMIN

Nick Veitch opens your eyes to the technology behind the cloud server revolution

Openstack

The cloud at last

All DevOps is headed to the clouds, but that doesn't mean abandoning open source platforms.

People used to eye the sky nervously and then fix you with a slightly menacing stare when you mentioned you had anything to do with clouds. Fortunately it is now a well-known word bandied about by IT types. Last year, Forbes reported that over half of US businesses used cloud services in some form or another (**www.forbes.com/sites/reuvencohen/2013/04/16/the-cloud-hits-the-mainstream-more-than-half-of-u-s-businesses-now-use-cloud-computing**), and of course, indirectly many consumers do too, whether it is for simple file storage, backup of mobile devices or cloud-based music streaming. In fact it is pretty hard to escape clouds, even if you want to.

So what is cloud computing?

The basics behind cloud computing are that it is an elastic resource, a commodity, where you can consume, on demand, resources such as computing power and storage. As such it has been highly disruptive – it isn't by accident that the pioneers in much of this computing-as-a-commodity revolution are an online bookstore rather than one of the major IT players. Though of course, Amazon now is a major IT player.

The good news for open source in all this is that the up and coming platform of choice with which to build OpenStack. The project began life as a collaboration between NASA and the hosting company Rackspace. The idea was simply to build a suite of open source tools that could provide the infrastructure to deliver cloud services, whether that be a private, internal, cloud (such as the original NASA project on which OpenStack is based) or a public-facing cloud, like the one developed by HP (**www.hpcloud.com**).

> **"With our OpenStack cloud, we've been able to offer physicists the resources they need to analyse and model data from the Large Hadron Collider."**

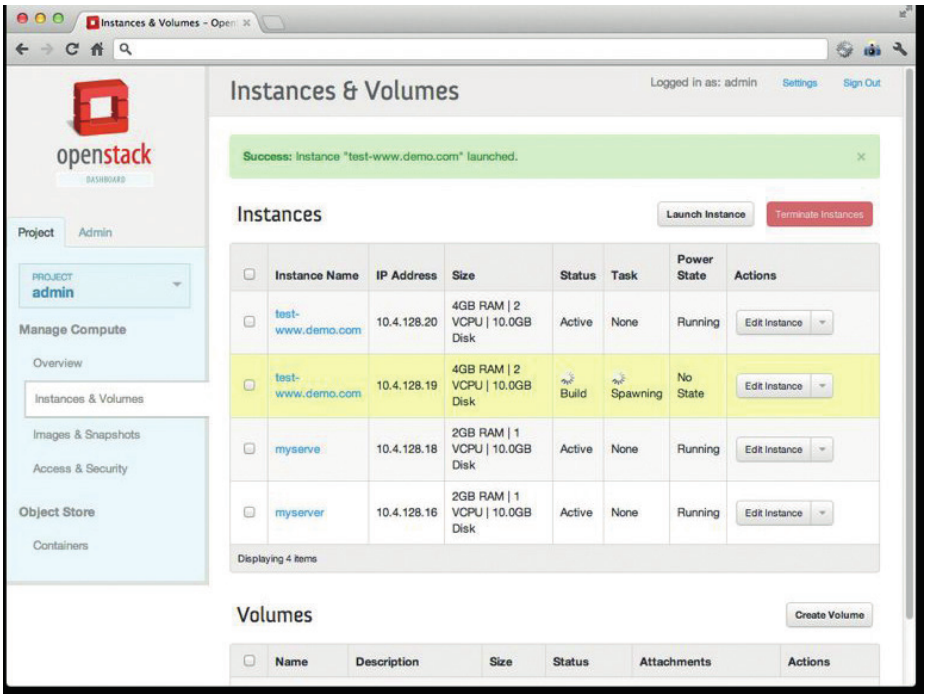

The OpenStack dashboard can be used to manage every aspect of your cloud deployments, and it's also possible to control it through an API, for great automation wins.

Why open matters

Openness in the cloud matters just as much as openness in an operating system. In fact, a cloud platform performs very much the same function as an OS – it enables the user to deploy software to perform some task. On a desktop, that might be image editing in Gimp; on a cloud it might be hosting a global payment platform or similar. With an open solution, it is at least theoretically possible to move your workload from one provider to another.

That may not seem to matter much (and indeed, it is currently difficult to achieve) but it will matter more as startups may need to shift workloads from local hardware to a more resilient cloud provider, or larger companies invest in hardware and pull services out of public clouds. It also provides a level playing field for innovation – many companies now contributing to OpenStack are developing a powerful and rich ecosystem of possibility. These companies include huge players in the IT field. The list is quite long and includes names like

Cloud anatomy

OpenStack is a suite of software tools designed for specific purposes. Each workload, and perhaps an entire cloud, may not need all of the components. New components are also being developed as the need for them arises. Here are the current ones (and their OpenStack names!)

Compute (Nova) One of the primary requirements of any work you may want to do in the cloud, the Nova component allocates compute resources (ie processing power to run software). Nova is designed to be very efficient at scaling horizontally (across many nodes, rather than adding more resources to a single node).

Object Storage (Swift) Object storage is storage that can store and retrieve blobs of data. They might be big blobs- a huge video file for example, or little blobs, like the *Linux Voice* payroll. The advantage of this sort of storage is that it is easily mirrored and replicated, and may be split easily across many actual nodes.

Block Storage (Cinder) This type of storage simulates a block-level device, such as you would find attached to your **/mnt** directory in Linux. As such it is transparent to applications accessing it and can be much faster than object storage, but it also requires slightly more overhead in resources. Like ordinary block devices, it needs a filesytem.

Image server (Glance) In dynamic clouds where nodes might be scaled out at any time, it is useful to have a store for images of particular devices (eg. a backup of a virtual server or a filesystem). This is a sort of specialised version of Swift (indeed, Glance often uses Swift as its storage medium), but with extra functionality, such as an API for querying what images are stored.

Networking (Neutron) Software-defined networking was first cooked up in 2008, and was joined later by Network Functions Virtualisation. Simply put, in an age where complex workloads are being defined in software, it makes sense to create the network topologies there too (Virtual LANs), and indeed virtualise other functionality (intrusion detection, DNS etc).

Identity (Keystone) Managing access to cloud services is a key concern. Keystone is a directory mapping users (and software) access to the services deployed. Authentication can be configured to accept a variety of methods – passwords, tokens, etc.

Dashboard (Horizon) The dashboard or control panel is just an easy to use interface for tweaking the cloud (or your part of it) as well as controlling the automation of some services. It is of course possible to use software tools to access all of this functionality through the API.

Metering (Ceilometer) Openstack describes this as a telemetry service, but in simpler terms it is basically the equivalent of an electricity meter, collecting data on what was used when and by whom. Although designed primarily for billing users, it provides auditing information for plenty of other purposes.

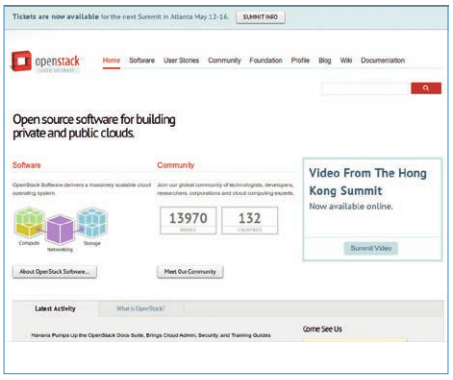

There is a very excellent website for OpenStack at www.openstack.org, and a very useful wiki at wiki.openstack.org.

Rackspace, HP, IBM, AT&T, Canonical, Dell, Cisco, SUSE, Intel, VMWare, RedHat, NEC, Hitachi and more.

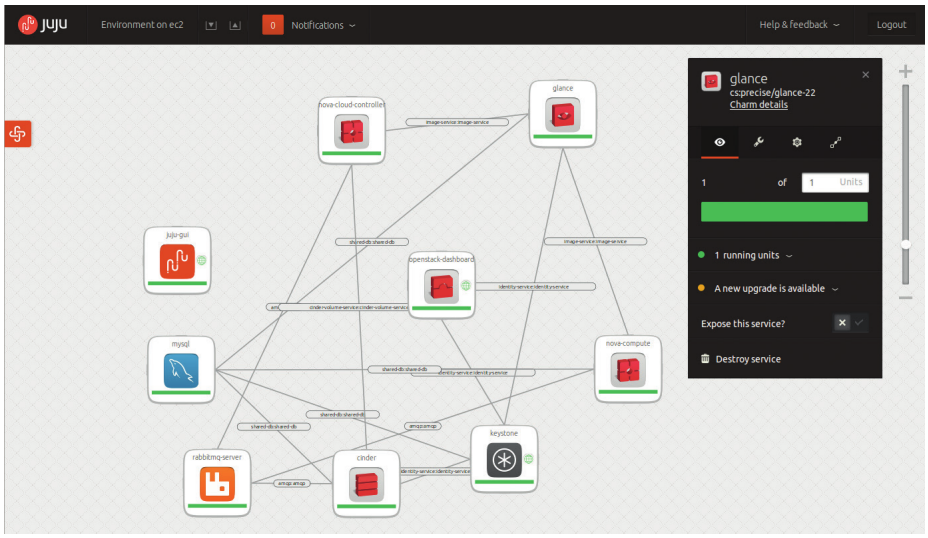

The Canonical's Juju GUI view of a deployed OpenStack instance. Other configurations are possible! Visit https://jujucharms.com and search for OpenStack bundles.

Loads of work

In many ways it is not the platform that is so interesting as the workloads that are deployed on it. There are inherent advantages to deploying even traditional server solutions on a cloud. If you run a simple LAMP stack on the cloud, it can automatically take advantages of features that the cloud provides for it – the underlying platform can make it more resilient, and vertically scale resources (storage, processing power) if required. There are more benefits to realise for software that has some mechanism for horizontally scaling – distributed, highly-available applications that run on multiple nodes. In future workloads will be written specifically for cloud use and take better advantage of the provided APIs to dynamically control aspects of their deployment.

Not everyone is a fan of OpenStack. Analysts at Gartner have continually downplayed the success of the platform, suggesting that there is a lack of differentiation and confusion about the services offered by various participants, and that the business model isn't clear. Uncannily, that is almost exactly what they used to say about Linux...

It is true that OpenStack is maybe not as mature as Amazon's Web Services, but the project is just over three years old, and there are a ridiculous number of companies committed to driving it forward. If you want to play with the clouds of the future, OpenStack is a great place to start. LV