


computing



Hybrid cloud – the future is open source

By [John Leonard](#)

28 Oct 2014 [View](#)

[Comments](#)



Among respondents to *Computing's* recent data centre research programme, the hybrid cloud model is generating a lot of interest. Indeed, moving towards a hybrid model was the aim of 41 per cent of them (see figure 1).

Hybrid cloud implies a close interconnectivity between a private cloud (i.e. a collection of physical and virtual systems used exclusively by one company) and [the multi-tenant public cloud services exemplified by Google, Amazon and Microsoft Azure](#). This seamlessly integrated whole allows data, services and workloads to be moved between public and private clouds at will, with the administrator able to monitor and manage the whole system via a single dashboard.

Got some sensitive data concerning an ongoing contract? Then apply a policy to ensure it doesn't leave the private cloud. Contract completed? Go to your dashboard and de-classify it; now it's free to migrate to the public cloud making space for something new. Need some extra resources in the short term for analytics? Burst into the public cloud and pay as you go.

It's an easy vision to grasp. What's more [it's the way things are moving](#), with the increasing adoption by vendors of open APIs and the abstraction of hardware-based functionality into software (as in the software-defined data centre) breaking down proprietary silos and allowing the movement of data and applications to be decided by policy rather than as a result of technological restrictions.

Easy to grasp the concept may be, but implementing hybrid cloud is not so simple. Cloud computing platforms do not generally interoperate well. Moving services from one proprietary public cloud to another or from a private cloud to a public cloud is often an expensive and difficult process.

Of course, you can always stick with one vendor. All of the major technology firms offer hybrid cloud solutions, generally collections of linked products bundled together. So, if you're a Microsoft shop you might choose to go with Cloud Platform System (CPS), which links Windows Systems Center and HyperV with Azure and runs on Dell Hardware. Or there's the HP Helion system, Rackspace's Rackconnect and VMware's vCloud Air. Alternatively, especially for those without the internal infrastructure but who need to deploy a private cloud quickly, there's converged infrastructure from the likes of VCE, IBM, HP and Nutanix which also enables connectivity between public and private clouds.

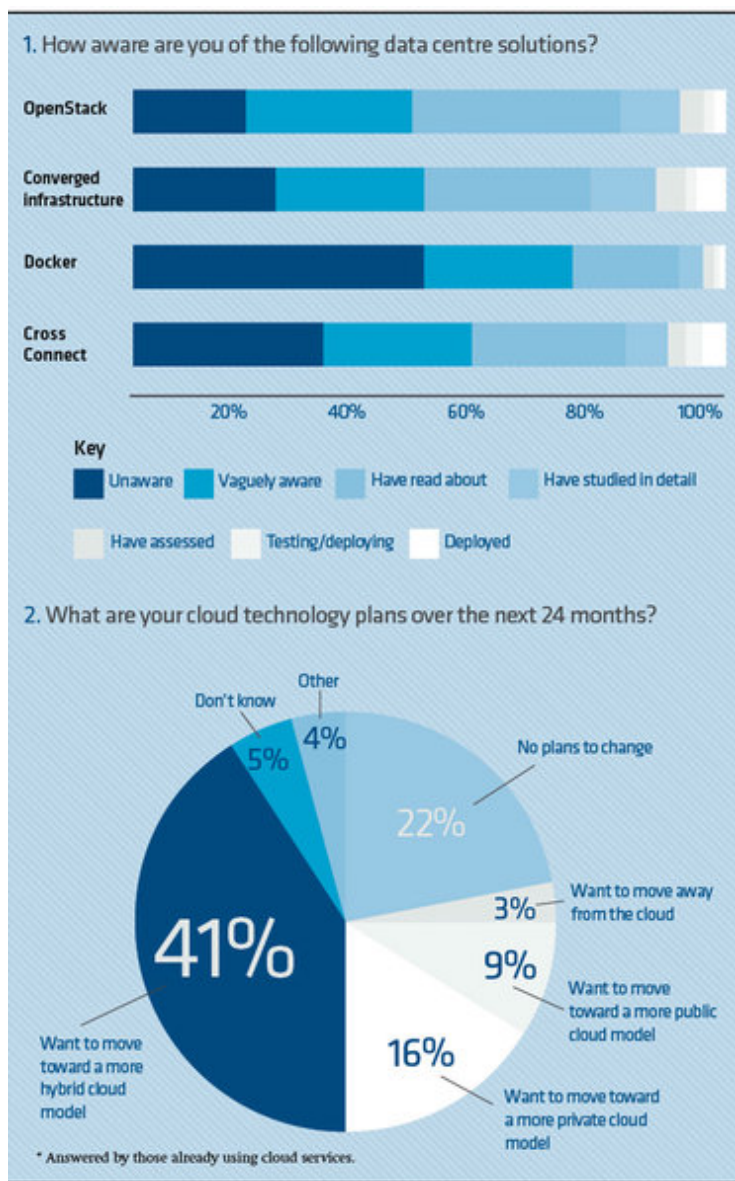
However, many organisations will prefer not to tie themselves down to a particular vendor. Arguably the most interesting developments are happening in the open source arena, both as individual projects such as Docker and also in multi-vendor collaborations such as Openstack.

OpenStack

OpenStack is an open-source cloud software platform that was originally created by Rackspace and NASA in 2010 and which now features contributions from hundreds of vendors including HP, Red Hat, Intel, Dell, AT&T and IBM, as well as thousands of individual developers. The fact that it has support from so many major technology companies is significant, in that it is designed, by virtue of its APIs, to be as platform-agnostic as possible, enabling it to scale out over a multitude of disparate environments. It also suggests that the vendors realise that it is a mistake to paint themselves into their own proprietary corner when it comes to something as big as cloud computing.

OpenStack is a collection of interconnecting services. Currently these include a database, dashboard, storage, networking, compute and tools for identity management, metering, orchestration and provisioning Hadoop clusters. In effect it is a cloud operating system that sits atop server clusters, storage and networking infrastructure and automates their operations, with virtual machines (VMs) and their workloads running on top. One of its key advantages is that it pools cheap commodity servers making them available on demand. As well as saving money on server

hardware and boosting agility, this can lead to more efficient usage of the infrastructure than when each server is dedicated to a specific purpose.



OpenStack is also compatible with public cloud services from the likes of Rackspace and HP, and with Amazon EC2 and S3 – and with a little work AWS and Google Compute engine. However, it is not compatible with Microsoft Azure or even fully with VMware’s public clouds.

Despite the latter’s stated support for OpenStack it is hedging its bets by continuing to deploy its own hypervisor in its OpenStack distributions.

Doubtless VMware is playing a game of wait and see. It may be only four years old but the OpenStack community is growing fast and many observers compare its growth favourably to that of Linux in the 1990s shortly before that operating system found its way on to the majority of web servers around the globe. If OpenStack takes off in the same way then no one wants to be caught on the wrong side of the fence.

Already it is used for hybrid cloud duties by companies such as Wikimedia, [CERN](#) and Cisco and in the private clouds of Ebay, PayPal, Wells Fargo and many more. However, as the survey results show, OpenStack still has some way to go before it hits the mainstream in the UK (see figure 2).

This aside, some of the IT heads we spoke to understood its value, not only to hybrid cloud but also to [the software-defined data centre](#).

“You will be able to get systems working together more collaboratively,” said the chief infrastructure architect of a major oil & gas firm.

“If you have converged infrastructure from IBM in the US and HP in the UK you can still use your OpenStack controller to report across both... OpenStack can potentially help you with vendor lock-in problems,” he continued.

“You could take OpenStack and a bunch of servers, OpenStack-compliant storage and networking and build your own software-defined data centre”.

However, others questioned its current levels of maturity, saying that open standards have not been implemented equally by all supporting vendors.

“The supportability is a risky area as you can’t sign up for a traditional kind of support,” said one interviewee.

“My feeling is that you should use the management tools that are provided by the vendor, if not you have to start downloading from the internet and using plug-ins,” said another.

There is also, as with all open-source projects, the potential for forking, with supporting vendors taking developments in different directions. And of course, if it fails to maintain and increase the support of the big public cloud providers it’s goals of ubiquitous computing will never be realised.

Docker and Flocker

Like OpenStack, but on an altogether smaller scale, Docker is another open-source project that’s making waves. Like OpenStack it seeks to grease the wheels in the journey towards ubiquitous computing, in this case by allowing applications to run on any Linux servers, without the need to provision virtual machines.

Linux is found in the majority of servers these days, and Docker makes use of features of the Linux kernel to create containers that hold the application and all of its dependencies. These “containerised” applications can be transported from Linux machine to Linux machine allowing the applications to function without any further tweaking being necessary.

The benefits of this approach are as follows. First, you don’t need different guest operating systems, and hence multiple virtual machines, to run the different applications as they will all run happily on a 64-bit Linux (or BSD) machine. This can save money on the cost of virtualisation licences.

Second, using Docker can ease the performance overhead caused by hypervisors.

And third, administration is simplified: with all applications running under the same regime, moving them from place to place, becomes a much simpler task. Thus, so long as private and public clouds are both running on Linux, creation of a hybrid set-up becomes simpler too.

However, there are obvious limitations, beyond its immaturity (the Docker project was only begun last year). Windows applications cannot be run on Linux without a hypervisor, and some believe security might be lessened in comparison with a virtual machine (VM). Another issue is what happens to the data associated with the applications?

The latter problem is addressed by another development that has happened alongside Docker: Flocker.

Flocker is a volume and container manager developed by ClusterHQ that allows applications in Docker containers to be deployed and moved with their data. It uses distributed local storage based on the OpenZFS file system along with proprietary management technology. Ultimately ClusterHQ hopes to expand the uses of Flocker to cover continuous replication, live migration and automated load balancing.

As with OpenStack, VMware has formed a partnership with Docker while continuing to push its own solutions, promoting it as making it easier for customers to run containerised applications on their VMware infrastructure and hybrid cloud, and working to create better interoperability between Docker and VMware.

Awareness of Docker was pretty low among our research respondents, but as companies like VMware know, things move fast in the world of open source and it’s better to embrace it than try to resist the changes.

@_JohnLeonard

[blog comments powered by Disqus](#)

© Incisive Media Investments Limited 2014 , Published by Incisive Financial Publishing Limited, Haymarket House, 28-29 Haymarket, London SW1Y 4RX, are companies registered in England and Wales with company registration numbers 04252091 & 04252093