Write Arm.

t: 0845 862 4646
e: info@writearm.co.uk
w: writearm.co.uk

https://stratogen.net/resources/our-blog/archive/how-openstack-private-cloud-works-with-containers/

How OpenStack Private Cloud works with containers

6 November 2015

The OpenStack Foundation's COO Mark Collier has described OpenStack as an integration engine that's agnostic. "That puts users in the best position for success. Just like we didn't reinvent the wheel when it comes to compute, storage and networking, we'll do the same with containers."

This move has taken a big step forward with the release of Murano, the OpenStack Application Catalog. This lets users find the applications they need based on category, tags and other attributes and, once found, it can be deployed immediately, or added to a Murano environment defining a multi-tier application, like a LAMP stack for web development.

Openstack and containers defined

OpenStack, the datacentre-level operating system, is becoming well-established in the datacentre for managing resources, including VMs. But there's a new kid on the virtualisation block – containerisation.

So, how does OpenStack work with containers?

Containerisation, as popularised by Docker, is a technology for wrapping applications in an operating environment that makes them portable. It provides many of the benefits of loading an application into a VM as the application can be run on any suitable physical machine without any worries about dependencies. Docker's sandboxed containers use a shared Linux kernel to support applications, which means that each container can be far smaller than a full-blown VM, so that, given the same resources, many more containers than VMs can be run.

Containers make it easier to try out experimental technologies, only deploying them to production when they're ready. They're also useful for the delivery of micro-services, and provide deterministic software packaging. And because they can be made stateless, it's easy to create as many as you want from a single image, and make changes to those instances in one operation.

Playing happily together

According to Joe Fitzgerald, Red Hat's VP and GM of cloud management, containerisation will help to simplify OpenStack installation and enable the building

of more complex solutions. And an OpenStack user survey found that most users were interested in deploying containers in conjunction with their OpenStack clouds for production use, so there's a lot of momentum to get the two technologies working together.

Find out more on how StratoGen OpenStack Private Cloud can help streamline your installation and manage complex requirements.

So in May 2014, an OpenStack team started working on improving support for container technology, with the aim of enabling users to create and manage containers on OpenStack just as they do with VMs. As a consequence, OpenStack's Magnum module now supports containers from Kubernetes and Mesos as well as Docker, using its container orchestration engine technology.

And following the release of OpenStack Murano, IT administrators can publish cloud-ready applications in an online catalogue. This enables OpenStack admins to offer a well-tested set of on-demand, self-service applications, and to find that these on-demand applications are deployed reliably and consistently. It simplifies configuration and provisioning, as well as managing the lifecycle of MongoDB by auto-scaling, self-healing, providing usage statistics for chargeback, and so on.

A work in progress

The relationship between OpenStack and containers remains very much a work in progress. For example, a standard for container images is under development by a group of 21 companies who have formed the Open Container Initiative. Docker has contributed its container format and runtime in support of this effort.

The OpenStack development community is committed to the project remaining hardware and software agnostic, and that OpenStack should retain its ability to manage anything, switching between VMs, containers and bare metal at will. Given that, expect container support to continue to increase and improve.

http://www.stratogen.com/resources/our-blog/archive/stratogen-cloud-backup-powered-by-veeam-cloud-connect/

StratoGen Cloud Backup Powered By Veeam Cloud Connect

Friday 10 July 2015

Backup just means making multiple copies, right? Not any more.

Backup now incorporates disaster recovery planning, and has wider ramifications for issues such as compliance and quality certifications. Accepted best practice is for backups to be stored in a 3-2-1 configuration, meaning three copies of data, on two different media, with one off-site. So off-site backups are no longer just nice to have: they're essential for disaster recovery planning, risk mitigation, and in many cases, compliance reasons.

Yet backup challenges you face are considerable and growing. It may be limited bandwidth, too expensive, you have burgeoning data volumes, or simply lack the resources to build or maintain a true off-site solution.

Off-site backup solutions

One way of instituting off-site backup is to mirror your live production sites, duplicating all systems. However, this approach is costly both to purchase and maintain, and requires capital expenditure that will need to be repeated at each hardware refresh.

So to help you avoid the risk of catastrophic data loss, whether you require native tape support, host-to-host replication, or a backup to a secondary site, StratoGen Backup as a Service is the right solution for you.

Backup as a Service

StratoGen's Cloud Backup leverages Veeam Cloud Connect to provide cloud-based backups for your local virtual guests and data. It allows you to keep an up-to-date copy or secondary copy of your virtualised applications in StratoGen's cloud, restoring files and virtual disks back to your local environment as needed. StratoGen uses Veeam Cloud Connect to enable you to use our global data centres for a backup plan that allows strategic location and optimised access.

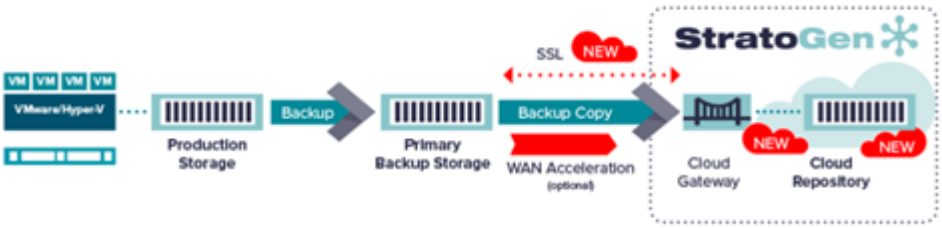
VM VM VM VM
VMware/Hyper-V
Production Storage
Backup
Primary Backup Storage
SSL
NEW
Backup Copy
WAN Acceleration
(optional)
StratoGen
NEW
NEW
Cloud Gateway
Cloud Repository

TOWARDS A MORE EFFICIENT CLOUD PLANNING PROCESS

Karl Robinson explains the strategies that IT management professionals can employ when developing a cloud computing master plan...

By Karl Robinson, Chief Commercial Officer, StratoGen

Challenges

As an outsourced IT option, cloud computing has an economic imperative that is second to none, with cost savings of between 40 and 85 per cent when compared to conventional `bricks and mortar' data centres. The actual cost saving, of course, is dependant on a number of issues, including whether you want a dedicated (private) cloud resource or are happy to use a shared (public) cloud system.

Other factors that influence the price - and therefore the cost savings - include the level and speed of access to the cloud resource, as well as the required `up time' of the service.

The economic imperative is perhaps best compared to the total cost of owning and running a company car, as compared against the cost of using a leased rental firm that supplies company vehicles on a pooled basis.

The key advantage of storing your organisation's data in the cloud, however, is that your business can then pay for the facilities it actually uses - rather than paying for the cost of data centre resources, whether or not you use them

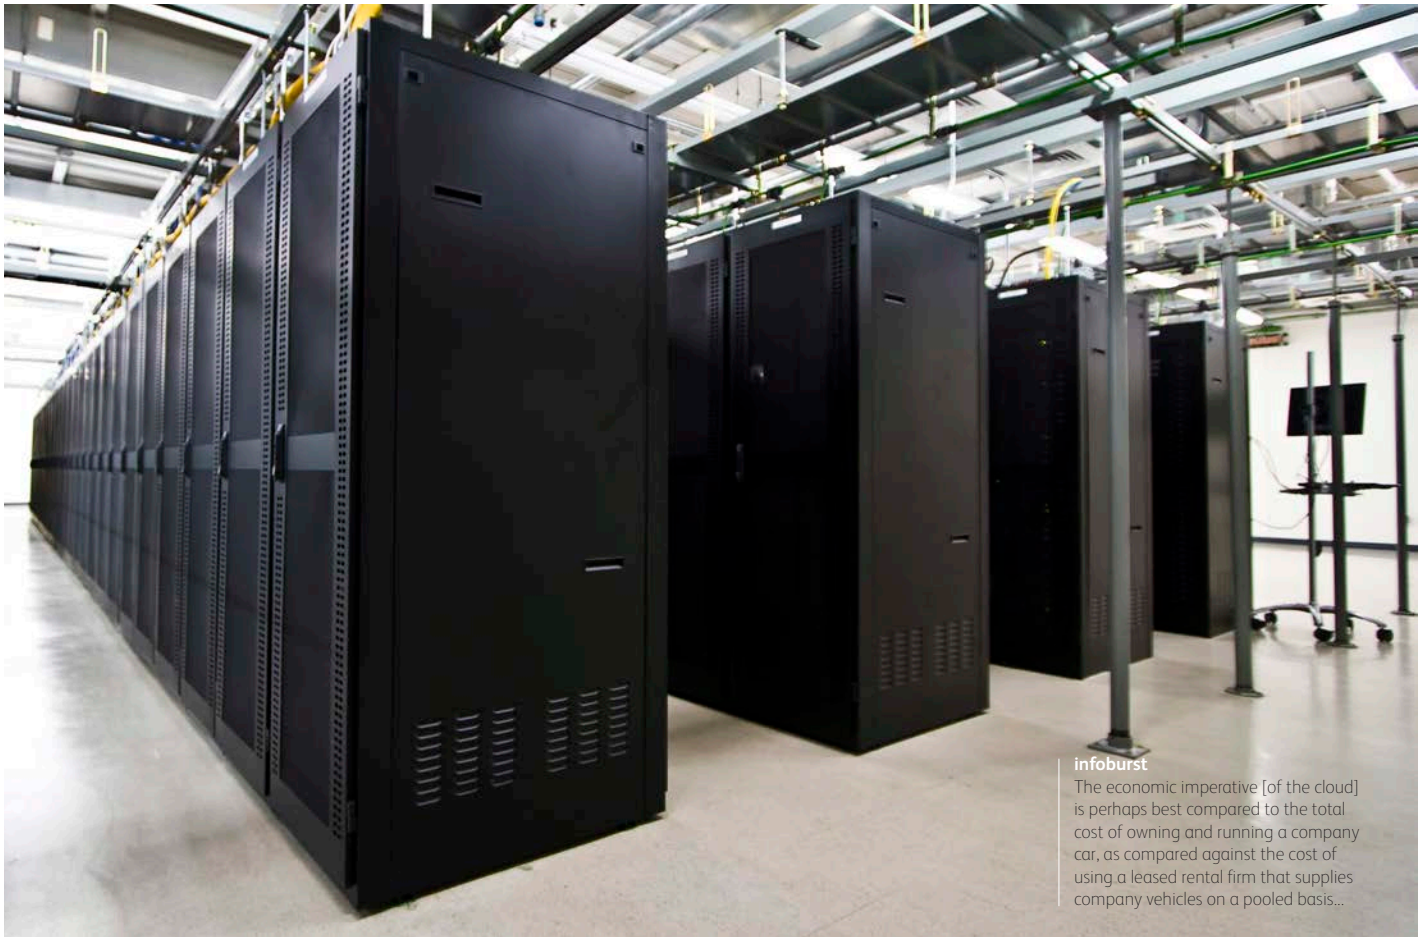

infoburst
The economic imperative [of the cloud] is perhaps best compared to the total cost of owning and running a company car, as compared against the cost of using a leased rental firm that supplies company vehicles on a pooled basis...

to their fullest extent. No small wonder a growing number of organisations are moving their data and IT resources over to the cloud.

The planning process with cloud computing, however, is the icing on the cake in terms of cost savings. If well executed, a well-planned cloud migration/implementation can mean the difference from truly saving money on the project in its first year, and only breaking even during the same period. And these cost advantages are not just a one-off, as they are recurring.

Groundwork

If you have carried out your groundwork - and other areas of due diligence - with regards to cloud computing, you will almost certainly have realised there are large differences in the cost of the various cloud facilities that are available.

These differences are not simply market-driven, but a reflection of the relative lack of maturity of the business models operated by many cloud computing service companies.

This is not a criticism, by the way, merely an observation - it is also a market differentiator that allows the clients of cloud services - that's you and your company - to select only those services they truly need, and to pay a fair price for those facilities.

It's worth noting that, when it comes to price differential factors, that criteria such as the ease (and speed) of access to your data - as well as where the data is physically stored - come into play here, as EC data protection laws often mandate that your company's data must be stored within the confines of the European Union countries.

There is a degree of pragmatism at work here, as whilst it is perfectly possible to host one's cloud-based data outside of the European Union, there are regulatory issues associated with this option. Increasingly, for example, many companies are discovering that legal issues such as the US PATRIOT Act come into play.

The PATRIOT Act is an Act of the US Congress that was signed into law by President George Bush in 2001. The stands for Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001.

Critics of the legislation in the IT sector point to the fact that the Act allows the US government and its many agencies easy access - without a court order - to data held on a IT resource that can be operated anywhere in the world, as long as the owning entity is a US company. This means that, whilst your cloud data may be held in a data centre in Dublin, if the owning company is based in the US, then the US government can request direct access to that data.

Some companies also elect to store their data in a UK cloud resource, either for financial regulatory reasons, or for ease of access.

Assuming that you have a cloud project in mind, the time then comes to complete the due diligence stage of the planning process, which typically involves scoping out your potential cloud service suppliers. This can range from discussing the suppliers with your colleagues at other companies, all the way through to requesting references from the supplier's existing clients.

This step is actually more complex than it initially appears, as there are a number of security and allied requirements that cloud service providers choose to meet, in order to better satisfy their client's needs.

These range from ISO 27001 compliance all the way through to compliance with PCI-DSS rules, a set of security standards that are mandated by credit card companies before a company is allowed to process credit and debit card transactions.

Where is your data?

The most fundamental question that potential clients ask their could service providers (CSPs) is whether their data is physically located. This can be a more complex issue than it first appears, as many CSPs choose to mirror (back up) client data across multiple data centres, except where the client has expressly elected to store their data in a specific territory.

Other questions that require asking include the speed of recovering data and the latency of the cloud service itself. It is no good relaying on a low-cost cloud resource if the latency is such that it takes several minutes to start downloading a given set of files or folders, and several days to download all of your data in its entirety.

Discussing this issue with your potential CSP will also reveal what levels of redundancy that the CSP's data centre resources actually offer.

This is an important issue, as a partially used CSP resource is usually a lot more responsive to data download requests than a data centre that is almost completely filled with data.

infoburst
The most fundamental question that potential clients ask their could service providers (CSPs) is whether their data is physically located...

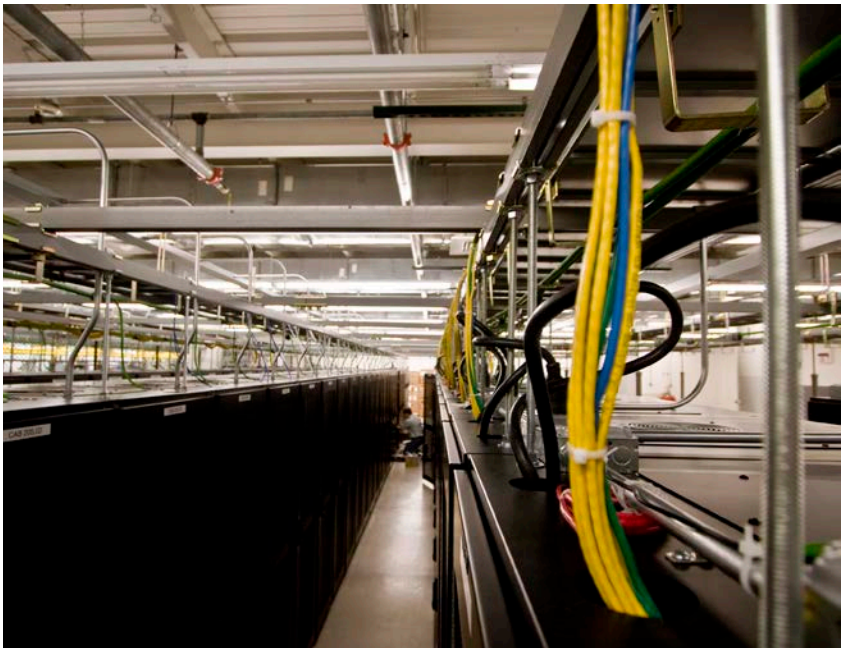

infoburst
It's worth noting at this point that StratoGen elects to use high-end - and known - vendor IT systems to maximise systems reliability and ensure the highest levels of support possible...

There may well be an argument to operate a cloud resource on a pooled or shared basis, since any resources that your organisation does not use can then be resourced by other users of the cloud facility.

This is a traditional cloud computing approach, and differs from the so-called `private cloud' facilities that a growing number of major companies now operate.

A private cloud resource is where the company concerned have full access - and control - over the cloud computing data centre, meaning that third parties are not involved.

Generally speaking, the larger a cloud computing data centre operation is, the better, as is the issue of whether the resource has a global portfolio of customers, as this strengthens the need for 24x7 active customer support.

There is also the issue of effective SLAs - Service Level Agreements. These are minimum set of service levels to which the cloud service provider agrees to, but are often drawn up before a given cloud service goes operational.

In our experience very careful attention needs to be given to SLAs, as there are signs that a few short-sighted CSPs try to include a number of limitations in their standard client agreements, hoping that this will - God forbid - allow them to side-step their responsibilities if something goes wrong.

Criticism

This is a not a criticism of CSPs generally, as it is important to understand that the nature of cloud computing service is such that it is perfectly possible to provide multiple redundancies for a given service with only a modest increase in costs.

Put simply, this means that the economic imperative of moving to cloud services is rarely affected to any major degree by pricing, so cost cutting is - usually - not an issue for most CSPs, except perhaps those providers operating at the lowest end of the price spectrum.

SLAs are also important, however, when it comes to dealing with what happens to the client's data at the end of the service contract or in the event that the CSP - for whatever reason - ceases operations, or curtails its services when (and if) it is acquired by a third-party company.

This leads us neatly into the questions that the diligent would-be cloud client should be asking of their CSP.

These questions are broadly in line with the due diligence questions that a company should be asking of all its IT systems suppliers, and centre on what type/quality of hardware the CSP uses - a growing number of service providers, we have observed, are opting for premium hardware systems for their cloud infrastructure, so reducing the MTBF (mean time between failure) of their systems, and helping to ensure that a given service is as close to 100 per cent uptime as possible.

It's worth noting at this point that StratoGen elects to use high-end - and known - vendor IT systems to maximise systems reliability and ensure the highest levels of support possible.

Conclusions

The adage that you get what you pay for applies in the cloud computing space. As with the provision of IT services generally, there is a fine line between operating a cloud services business profitably and at a break-even/loss-making level.

This is because profit margins in the CSP world are reflective of the growing maturity of the industry - as with the IT hardware industry of around a decade ago, profit margins are becoming commoditised to the point where a profitable CSP of today may find its service platform disrupted by new technology players in the future.

Unless the CSP concerned has deep pockets, this can cause the service provider to either seek acquisition by a larger company or focus its services on the more profitable clients.

In most instances, this means that smaller users of a cloud resource - through no fault of their own - may find the entry-level price of the service they are using starts to rise, making the cloud service less attractive in economic terms.

The good news is that the cloud industry has yet to reach this point on its evolutionary scale, although observers suggest that this issue will raise its ugly head in the cloud computing space sooner, rather than later.

Having said this, it is possible to prioritise/de-prioritise aspects of the various cloud services that a company uses in order to maximise the cloud resource's return on investment - without affecting/degrading the IT systems resources that are available to your company and its clients.

Your mileage - as they say - may differ, but our observations suggest that careful planning at all stages in a cloud services project can go a long way to avoid many of the pitfalls that early adopters encountered.

www.stratogen.net