Hybrid Cloud Storage Visions and Potential

Author: Scott D. Lowe

Partner, ActualTech Media

November 2016

© 2016 ActualTech Media

Introduction

Companies don't want "cloud." They want the benefits that come with cloud. If that means they must adopt cloud services, they will, but it doesn't mean that the public cloud is the only option.

The Cloud as a public service inherently provides a number of benefits that have been discussed in this book. From turning traditional data center economics upside down to transforming how new services are deployed, the public cloud has changed how people think about the infrastructure upon which IT operates.

However, the outcomes are a journey, not necessarily a destination. If you're able to provide your company with improved IT economics and can feasibly implement the kinds of architectures that enable simple workload deployment, you will have done your job. With hybrid cloud storage systems offering robust APIs and a virtualization-first methodology, you will be well on your way to cloud success.

Let's dig a little deeper into the potential vision and outcomes that you can expect with the architecture described above.

Contents

Manage a Cloud

Once you have decided that a hybrid cloud is for you and that you need cloud-like characteristics embedded in your IT function, there are a number of achievements that you can unlock for your enterprise.

DevOps

We talked a bit about DevOps earlier on in this book, but it's emerging as a key consideration for how IT moves forward with data center architecture and infrastructure, or lack thereof (cloud). As you consider undertaking a hybrid cloud journey, consider the potential good that is imbued into your DevOps environment and processes. A fungible infrastructure solution, such as one based on cloud principles, enables an agile development environment for developers, who can now create and destroy virtual infrastructure on a whim in support of their development and testing efforts.

The goal here is to enable an agile environment to help push the business to meet and exceed its goals. This can happen with local infrastructure, but the environment can also extend to the public cloud with the development and support of cloud-native applications that seamlessly integrate with local applications.

Self-Service

Let's use a development example again. Developers are constantly in need of fresh infrastructure as they seek to develop bug free software that operates with high levels of performance. As they sully good virtual machines with test versions of their code, they need to continuously refresh the environment. It's doubtful that even small development shops want to constantly involve IT operations for these needs, so they need a way that they can create and refresh their own environments.

Likewise, individual business unit IT specialists may need ways to create centralized services that reside in the corporate data center. Although they may not be a part of the central IT staff, they want the ability to meet their own regular needs without always involving IT.

Such on-demand, self-service activities are increasingly important. On the storage front, enabling self-service capabilities requires the creation of templates that specify, among other things, quality of service (QoS) and data protection policies. As IT, you create a series of templates ahead of time and then allow users to simply choose from them. Behind the scenes, your API-driven automation and orchestration services do the heavy lifting.

In a multitenant environment, you want control of Quality of Service (QoS). But legacy storage can only set QoS at the LUN level, and that leaves all the resident VMs still fighting over assigned resources. Contrast that with QoS per-VM, enabled only by VM-aware storage. QoS at the VM-level allows you to guarantee resources for an individual VM; you can set a maximum ceiling on a rogue VM or a minimum floor for a critical VM. Cloud Service Providers can use QoS per-VM to establish performance tiers and premium services for their customers.

Chargeback / Showback

IT resources aren't free in any economic model, whether that's an OpEx-heavy cloud model or a CapEx-intensive traditional model. At some point, those services need to be paid for. For a variety of reasons, organizations often prefer to make sure that individual departments are charged for their data center activity. For example, if the sales group requires ten virtual machines in the data center environment, they cover the costs of those resources. This paints a much clearer picture from an organization-wide budgeting perspective than does the model where all data center purchasing is a part of the IT group's budget.

In a public cloud environment, it's easy to decide who pays for what since cloud provider environments are, by their nature, multitenant-oriented with comprehensive billing capabilities. In the virtualized world of a local data center, it can be a bit more challenging sometimes, particularly with legacy storage systems that aren't capable of neatly delineating who is using which resources. In a modern storage environment that leverages VM-aware storage, it can be much easier to identify in-use resources and map them to their consumers.

> **Chargeback vs. Showback**
>
> Chargeback is the process of actively billing business units for their resource utilization and is a common central IT cost recovery model. Showback is "chargeback light" in that it shows a business unit how much they're using in the way of

Autonomic Operation

Raw computing power continues to increase as Intel releases processors jam-packed with more and more cores. With all this compute power, there's plenty to spare for a core or two to keep a watchful eye on the data center environment. And, with an environment that has a lot of great APIs and automation and orchestration capabilities, we can do some really interesting things.

For example, if you're a typical company, you probably have people that work from 8 AM to 6 PM and that's your heavy period. Or, you may be an ecommerce company that experiences a surge at certain times of the day, but you're relatively quiet at other times of the day.

What if you could create an environment that simply managed itself based on the workload? This goes beyond the application-centric DevOps scenario that was discussed earlier. In this scenario, the infrastructure senses that there's no need for all thirty vSphere hosts to be operational since the current workload requires only eight. As such, the management layer migrates all of those workloads to a few hosts and shuts down the rest to save electrical and cooling costs. As demand dictates, hosts are brought back online and the virtual machine workload distributed accordingly. The same kind of process can be used with the storage layer, as long as that storage layer has some intelligence.

You can even create an environment in which test and development data is automatically refreshed from production at certain intervals so that your developers are always working with current information. This is sometimes referred to as

ActualTech Media

copy data management, but many providers have even more sophisticated interpretations that reduce operational overhead and storage consumption. Copy data management refers to a process in which the storage system manages copies of data by creating virtual copies of data as they are needed. New data copies aren't actually physically created, thereby reducing overall capacity needs. These capabilities are useful for development needs as well as for times when you need to create a copy of a virtual machine to, for example, perform point-in-time data analysis.

Summary

With a virtualization-centric architecture, VM-aware storage can help companies achieve these goals. By leveraging a public cloud-like web services architecture and clean, complete APIs, VM-aware storage enables greater automation, orchestration, scale and self-service than is possible with legacy storage solutions. Companies that are investing hybrid cloud are investing in these kinds of technologies to realize the full potential of their cloud vision.

The hybrid cloud isn't just a financial play, although the right technology *can* help to rein in capital and operational expenses. Rather, it's about the future of the business and the ways that you can best support your customers, employees, and partners. It's about possibilities and potential—the potential to propel your business ahead of the competition; the potential to eliminate any silos that may exist; the potential to transform the environment into an unstoppable force that can adapt as quickly as the business environment demands.