ActualTech Media

Tintri

Hybrid Cloud Requirements

Author: Scott D. Lowe
Partner, ActualTech Media
November 2016

Introduction

In previous papers, we've talked about the journey and how we got to this point. We've chatted a bit about the trends that are assisting with a move to the hybrid cloud. Now, let's talk a bit about what needs to happen for the hybrid cloud to actually come to fruition.

Contents

Storage Platform and the Public Cloud

In any environment, data is king. An organization's data is its lifeblood. As a result, storage is the only non-ephemeral resource in the data center. Compute and even networking frameworks can be built and destroyed at will, but storage requires much more care since it often needs to live forever (or, at least, for a very long time).

Regardless of environment, storage must be well-supported and there must be permanent mechanisms in place for its safety and retention. Moreover, because storage persists for a long time, it also needs to be protected from prying eyes. The financial and reputational backlash in failing to do so can be devastating.

Beyond ensuring ongoing security, storage in the hybrid cloud must be flexible when it comes to where it lives. After all, the very nature of the hybrid cloud is that it straddles both public cloud providers and private cloud environments. With that in mind, a basic assumption behind hybrid cloud storage is that it has at least some hook into the public cloud. For some vendors, this may mean running software-based versions of their operating system in a public cloud provider to enable a full feature set both on-premises as well as in the cloud. For others, it may mean having the ability to use the public cloud as a replication or data protection target. And of course, it may be a combination of these kinds of features or evolve into something else entirely.

Scalability

One of the key characteristics of cloud—whether it's public, private, or hybrid—is the ability to scale in whatever dimension makes sense for your workloads' needs. The traditional direction for storage scaling was up—adding more shelves of capacity as needed. But now that capacity generally comes with a lot of all-flash IO, and so the controller becomes the bottleneck. And piling

Inside the Data Center: Scale Up vs. Scale Out Systems

For years, a silent war has been waged inside your data center… one with the hallmarks of some of the greatest historical matchups: Coke vs. Pepsi, Mac vs. PC, Kirk vs. Picard. In this war, however, unsuspecting organizations are the losers.

Scale up vs. Scale out: To the untrained eye, both are just ways to add more capacity to an existing storage environment.

Scale up architectures allow you to add more disks to an existing array or add expansion shelves of disks atop a processing unit, which contains the processing and memory resources for the storage environment.

In a scale out system, every time you add capacity, you also add the underlying resources upon which that capacity depends. Add a shelf of disks and you also add more processors, RAM, and network/storage fabric ports.

This is really important and here's why: Predictability. As you add storage capacity to the data center, you shouldn't have to give up performance, but that's exactly what can happen in a scale up environment. Eventually, you overtax the shared processors and fabric connections and you begin to suffer from storage performance problems. When that happens, you must start being careful about where you're placing workloads and virtual machines so that you avoid the hot spots.

In the proper world of scale out storage, you don't run into such issues since you're adding all of the resources your storage needs to function. You can scale workloads and maintain predictable levels of performance. Even better as you add new workloads and virtual machines, the right system can optimize placement of these items without worry that necessary resources won't be available.

Happy users, happy business.

up shelves of capacity only creates a growing failure domain—lose the controllers and lose the data on ALL the shelves to which they are attached.

And so more organizations have shifted the direction of their scale plans to out. This entails adding more nodes—interconnecting arrays to join their controllers together, and adding both capacity and performance. This adds redundancy of controllers and increases your ability to put all that performance to effective use.

The challenge here is to keep storage scale-out as simple as say… scaling compute. If you need more compute resources for your cloud workloads, you simply install a new virtualized server, add it to the resource pool and automatic live migration optimizes the VMs across the pool.

Fortunately, the previously mentioned category of VM-aware storage has this visibility across your pool and capability to optimize the placement of every VM. Simply add another node of VM-aware storage, and intelligent algorithms re-distribute VMs to best balance capacity and performance requirements. Best of all, it all happens automatically without admin intervention.

Now remember that scale also means spinning workloads up and down. When we're talking about scale in the context of cloud, bear in mind that it should be possible to expand the size of a resource pool as well as shrink resource pools when workload demands begin to diminish.

But, there is a real difference in achieving scale in a public cloud as opposed to a private cloud, which means that the two sides of your hybrid cloud will have somewhat different capabilities here. On the public cloud front, you can temporarily provision a resource and then, when you're done with it, deprovision that resource on the fly without incurring any capital expenditure costs. On the private side of your hybrid cloud, if you run out of a resource, you'll need to buy more and deploy it before you can provision it for use. If you have a temporary workload, even if you deprovision the resource when you're done with it, the resource still physically exists in your data center and you've paid for it. In a private data center, by deprovisioning resources, you're making them available for use by other workloads.

Put concisely, on the public side, you get the full cloud experience, which includes both economics and operational efficiency. On the private side of your hybrid cloud, the focus is on making things easier to expand and contract and to make more efficient use of your resources.

Automation and Self-Service

In cloud, automation of routine tasks is a defining characteristic. In any hybrid cloud scenario, you should be able to automate the deployment of new resources, such as virtual machines, networks, and storage resources. Even if you're not a DevOps-drinking developer, you can still adopt some of the infrastructure-as-code methodologies that are bringing new possibilities to the application workload environment. For example, you can create self-service

systems that enable individual business units to provision their own resources. Or, consider this scenario: suppose you're working at a college and want to make sure every new student gets their own virtual desktop and dedicated storage space. You can write a routine that creates these resources automatically each time someone in the Admissions department accepts a new student and changes their record in the database.

If you haven't enabled deep automation in your data center, you don't have a hybrid cloud. Remember, self-service is one of the virtues of the cloud; without automation, you can't introduce self-service into the environment. This doesn't necessarily mean that every user in your company needs to be able to provision resources, but it does mean that the ability to use resources should not be limited to just the operations staff. In a DevOps world, this capability should extend to, at a minimum the development staff.

Summary

These are just some of key requirements for a hybrid cloud in general. In the next paper, we'll go deep in what's necessary for the storage resource.