VIRTUALIZATION & Cloud Review

EBOOK

# THE JOURNEY TO THE CLOUD

By Trevor Pott

## IN THIS GUIDE:





# CLOUD ADOPTION ISN'T ALL OR NOTHING

## The cloud might be the future, but there's no rush–take the time to do it right.

**You're a smart systems administrator** who can see the writing on the wall: The cloud is the future. On-premises clouds might be part of the mix, but somewhere in there everyone is going to have to incorporate workloads that exist either on a services provider cloud, or on the cloud provided by the big four public cloud providers. How do you get there from here?

The initial assertion that the cloud is the future may strike some as bold. This is especially true of individuals and organizations that have yet to embrace the cloud. In reality, however, "no cloud" policies are increasingly rare. By 2020, analyst house Gartner Inc. predicts they'll effectively be nonexistent, and this is backed up by virtually all other analysts and polling on the subject that is publicly available.

Gartner also says that current public cloud spending is rather high ($219.6B USD), with predictions that the amount recorded in 2016 will nearly double ($411.4B USD) by 2020. While this should be taken with a grain of salt – Gartner's predictions aren't exactly stellar when it comes to market size – the actual numbers don't matter so much. What matters is that both analyst predictions and real-world evidence show continued strong growth, with no end in sight.

Where the statement "the cloud is the future" gets hackles up is that a decade of tech marketing has left people with completely different ideas of what "the cloud" means. For some, mentioning the cloud conjures images of Infrastructure-as-a-Service (IaaS) solutions like Amazon Web Services (AWS). For others, the cloud means Software-as-a-Service (SaaS) solutions like Salesforce.com.

We all use cloud services–personally, if not professionally. Even if your worksite is inside of a hardened bunker in the middle of the high arctic, with no external network connectivity, somewhere in your life you are using an instant messaging application, Facebook, Twitter, or even just reading news articles online. Behind all of that, today, there is a cloud.


Icatnews / Shutterstock.com
Cover: jijomathaidesigners / Shutterstock.com

How we perceive cloud services determines to a large extent how we react to the idea of moving organizational IT resources to the cloud, and whether or not we realize we've already begun to do so.

### Phasing In

An important discussion point when planning your journey to the cloud is the difference between a bulk move and slowly phasing cloud solutions in. Today, the bulk of virtualized organizations run VMware's virtualization solutions. Most do so without any self-service portals or IT automation solutions. In their quest to become more efficient, however, organizations are slowly adopting on-premises cloud technologies.

VMware's vCloud Suite is a popular choice for the creation of an on-premises cloud, in large part because it can be layered over the top of your existing VMware installation. Microsoft has its own offerings that layer atop

its Hyper-V virtualization platform, and those invested in the open source world have OpenStack.

To varying degrees, each of these solutions offers customers with existing virtualization deployments a means to gradually ease into the cloud. Automation and self-service are simply added to what already exists and, voila: a wild cloud appears.

An on-premises cloud doesn't mean that the transition will always be gradual. If you purchase an on-premises cloud appliance and migrate all your workloads, that's a pretty blatant bulk movement of services.

This path to becoming cloud-enabled is less traumatic than a bulk migration to AWS, which is what most people think of when told to "adopt the cloud." Somewhere in the middle is the adoption of hybrid cloud solutions involving services provider clouds, (for example, the myriad vCloud Director-based VMware services providers,) or fully managed on-premises-compatible offerings like Microsoft's Azure Stack or VMware Cloud on AWS.

## Hybrid Clouds

Phasing cloud management software over the top of your existing virtualization infrastructure is the path of least resistance, but it only gets you so far. One of the touted benefits of cloud computing comes from the ability to outsource the underlying infrastructure entirely.

For most organizations this starts with the elimination of their disaster recovery (DR) site, and the adoption of a services provider or major public cloud provider as the new backup and DR destination. The right services provider can make all the difference for organizations cautiously adopting cloud computing, while a poor one can poison an organization against the cloud for years.

The major public cloud providers don't really have a great reputation for support or customer service, especially if you're a smaller organization. They stand up a fantastic cloud with a great UI and a top-notch API, but getting it running is left to the customer.

Smaller services providers, on the other hand, frequently specialize in meticulous support tailored to help organizations with no cloud experience take those first few cautious steps. Most smaller services providers put the time in to help customers configure their solutions appropriately, and develop testing and automation regimens that not only ensure solutions like DR to the cloud are working properly, they ultimately help an organization's IT team out by freeing them from mundane day-to-day tasks.

The helping hand of a managed services provider comes at a price. In terms of raw numbers, the major public clouds are always going to be less expensive. These hyperscalers operate at scale no other providers can even approach, and they can grind efficiencies out of their supply chain and operational management that the rest of us can only dream of.

The counterpoint to this is that most organizations aren't ready to play in the hyperscaler sandbox. If an organization was prepared to burn down their entire IT plant and rebuild it from scratch, learning entirely new skills, then the major public cloud providers are a great plan.

Phasing cloud management software over the top of your existing virtualization infrastructure is the path of least resistance, but it only gets you so far.

Most organizations don't have that luxury. They have to keep what they have running while they migrate to the cloud. A path that connects the two that doesn't involve jettisoning a career's worth of knowledge and experience is an attractive bonus.

For this reason, services providers clouds—especially those offering hybrid solutions that build on an organization's existing virtualization platform—remain popular.

## Technical Debt

On-premises workload management differs from hosted workload management. When you own your own gear, standing up a workload and leaving it running 24x7 is the norm. Organizations design their datacenters to handle the peak amount of workloads, so there's little incentive to turn off idle workloads. This is a terrible way to engage with hosted cloud services.

As you're charged by the hour for workloads that are active, cost minimization on hosted clouds requires that workloads be active as little as possible. This is best accomplished through the use of composable workloads.

At its most basic, composability has a workload's networking made dynamic and automated so that when a workload is moved from cloud to cloud it remains accessible.

A more complete version of composability separates the data from the application and the operating system environment (OSE). The OSE and application's configuration are laid out in a configuration file, and they can be instantiated as required, from scratch. In this case only the workload's data needs to be moved from place to place, not the entire workload. This makes standing the workload up and tearing it down as needed significantly easier.

The cloud might be the future, but there's no rush. Take the time to do it right. Automate what you can. Build a private cloud and evolve toward hybrid capabilities.

Most organizations are nowhere near composability. Workloads with static IP addresses are common, and the use of configuration management tools like Puppet, Chef, Ansible and Saltstack remains rare. Adoption of more advanced infrastructure automation solutions like Terraform is still in infancy.

It's this technical debt that is the real barrier to simply packing up your IT and moving it up to a major public cloud provider. A VMware administrator can learn the AWS interface if he has to. But it's a long road from pets to cattle.

The cloud might be the future, but there's no rush. Take the time to do it right. Automate what you can. Build a private cloud and evolve toward hybrid capabilities. Embrace services providers and rely on their experience. The major public clouds will still be there when you're ready.

In the meantime, one application at a time, your organization is probably already embracing the cloud. From Salesforce.com to Office 365, workloads that once were fully managed by an organization's IT are now consumed as a service.

Cloud adoption isn't all or nothing. Don't let anyone tell you otherwise.

# CLOUD-BASED DISASTER RECOVERY: HOW TO AVOID THE 'GOTCHAS'

Be sure to familiarize yourself with the limitations and capabilities of cloud DR.

**Backups and disaster recovery (DR)** are the standard first step into the cloud for many organizations. As a result, it is where organizations are most likely to encounter — and have to overcome — the myriad of non-technical problems that accompany outsourcing IT. What are these cloud gotchas, and how do they manifest for those attempting backups and DR to the cloud?

Cloud-backed data protection is an attractive solution. So attractive that is has infiltrated the daily lives of most of the western world. Sync-and-share solutions such as Dropbox, OneDrive nd Sync are often the first exposure users have to the concept.

Sync-and-share solutions allow individuals the ability to place files in a local folder and have those files automatically be uploaded to the cloud. Once in the cloud they're versioned and offer a limited ability to search through the history of a file in case a mistake was made and a previous version needs to be recovered. Sync-and-share solutions aren't what most IT practitioners would call a "proper" backup solution, but for the vast majority of personal users, they're the "better than nothing" solution that's most likely to get used.

## Expectation Management

Sync-and-share solutions are a great opportunity to examine one of the key problems of cloud computing: Namely, that expectations rarely match reality. IT practitioners have a pretty good idea of how sync-and-share solutions work. As a result, it rarely occurs to us to sit users down and have the talk with them about the cloud not being magic. Unfortunately, many users encounter cloud-based solutions with erroneous preconceptions, and this "knowledge" leads to errors.

One common belief is that sync-and-share solutions adequately protect users against ransomware. They don't. Modern ransomware makes numerous changes to files over time, ultimately running out the number of versions of files kept by the sync-and-share solution. There's even a little game of cat-and-mouse going on with some of them where the sync-and-share vendor tries to add some level of ransomware detection based on access patterns, and the ransomware evolves new access patterns.

Proper backups — solutions designed specifically for the purpose of backing up data — have different approaches to the problem. They'll set files read only once initially uploaded, or allow unlimited versions combined with alerting systems administrators of an infection.

Joe Techapanupreeda / Shutterstock.com

Sync-and-share solutions are only profitable because they sell users way more capacity than most of them would ever use. This oversubscription means that one good malware infection that increases file version numbers — and thus storage usage — could put a sync-and-share operation out of business.

Proper backup to the cloud solutions don't have this problem. They make the customer pay for usage. As a result, they have the ability to focus on data retention instead of data availability. Two different approaches, but used by many individuals to solve the same problem. One of the two — sync and share — is emphatically not suited to the task, but a decade of popular misconception will mean it continues to be used in this way.

All cloud solutions suffer from this issue to some extent. It makes expectation management one of the first — and most difficult — hurdles for new cloud administrators to overcome.

During a crisis that triggers a failover to the cloud is not the time that systems administrators want to discover that their DR solution is useless because key workloads aren't composable.

## Disaster Recovery

If expectation management related to simple backups can get messy in a hurry, the issues surrounding DR quickly become downright dangerous. Unlike the nearly universal sync and share, systems administrators are not widely familiar with the limitations and capabilities of cloud DR.

At first glance, cloud DR should be no different than DR to an organization-controlled second site. The No. 1 issue administrators are going to encounter in both cases is ensuring that workloads function on the DR site the same as they do on the primary.

Here, the most common problem is network configuration. On the production site workloads are set up with specific network configurations that allow them to interoperate with other workloads without conflicts. This may include static IP addresses, and usually involves DNS and firewall rules.

Making the workloads ready to be usable during a DR scenario involves setting the workloads up with dynamic IP addresses, configuring other workloads to access services offered by DNS instead of IP, and configuring the rest of the network infrastructure to automatically — and rapidly — update DNS, firewall, intrusion detection, threat protection and other services to accommodate the shifted workload.

This is the first step toward making workloads composable, and it's not easy. When an organization owns both sides of the DR solution, then the infrastructure is a lot more forgiving of mistakes made in making workloads composable. If an admin forgets to make an IP address manually, for example, administrators can simply log in to a virtual machine's console and change the IP address.

Not all cloud technology offers this ability to manually restore workloads to operation. For many cloud solutions, if the workload doesn't come up in such a manner that it can be remotely accessed using the operating system environment (OSE) native remote access tools, then the workload is irretrievable.

During a crisis that triggers a failover to the cloud is not the time that systems administrators want to discover that their DR solution is useless because key workloads aren't composable. Unfortunately, DR solutions are often sold using perfectly configured, best-case demos, creating false expectations with administrators, or simply never introducing them to potential problems in the first place.

Lack of time to thoroughly research the solution, lack of training in the intricacies of the offering, lack of budget to engage more full-serviced cloud offerings, and above all a lack of testing combine to make this first step into the cloud a disastrous one for many organizations.

## Cloud Business Practices

The problems involved in moving to the cloud are human as much as they are technical. They involve our susceptibility to marketing and sales fluff, a decade of preconceptions and overcoming the innate desire not to be bothered with difficult problems.

Outsourcing your IT to a cloud provider is supposed to make IT easier, but it can only be easier if you accept that the cloud isn't somehow a magic talisman, but is, in fact, nothing more than outsourcing. If you hired a custodial service to

take care of the office and grounds after hours, you wouldn't simply click a few buttons on a Web site and never think about it again.

When outsourcing custodial services there are practical issues to be concerned with: how keys will be made available, the insurance and reputation of the custodians, what to throw out and what not to throw out. So it's odd — but somehow nearly universal — that organizations don't have the same level of engagement or consideration with outsourcing their IT to a cloud provider. This is one place where appropriate businesses practices can make a big difference.

The cloud doesn't come with a license to disengage your brain. Cloud services providers, however, do come with training wheels.

IT practitioners with a great deal of cloud knowledge and a lot of time to prepare on-premises workloads for the transition to the cloud can simply saunter on up to Amazon Web Services (AWS) with a credit card and make magic happen. 10 years after AWS was born, the number of IT practitioners with the skills to do this remains small.

There are, however, a number of smaller regional providers available that offer a more managed service, and it's worth mandating their use, at least in the beginning. Use of these solutions can solve a number of problems, starting with being more forgiving of mistakes with composability, but extending into other areas.

Smaller services providers are more likely to have VMware and Hyper-V hosts available, eliminating the added complications of converting workloads as part of the DR process. Smaller services providers often have experts available to help with the really tough problems, like ensuring that bare metal workloads can effectively use a cloud solution for DR purposes.

Mandating rigorous and thorough testing is another business-level decision that must be taken before embracing the cloud. Whether the cloud is being used for backups, DR or for production workload testing, more testing and testing all over again needs to be a constant requirement.

It's all too tempting to treat the cloud like a fire-and-forget solution that magically solves problems with which you don't want to deal. Adapting business practices to the cloud needs to begin with countering exactly that tendency, and these business practices need to stay in place even as comfort and experience with cloud solutions grows.

The cloud doesn't come with a license to disengage your brain. Cloud services providers, however, do come with training wheels. And for all your cloud needs — from backups to DR and beyond — they're a great place to start. At least until you're ready to ride on your own.

VIRTUALIZATION & Cloud Review

# GDPR AND THE CLOUD: WHAT YOU NEED TO KNOW

The level of attention to privacy, security and data sovereignty won't be easy. Companies that haven't already started down this path are potentially in deep, deep trouble.

**When exploring IT outsourcing** of any kind it doesn't take long before privacy, security and data sovereignty concerns pop up. These issues are complex, and the same rules don't apply to everyone. By now, even systems administrators uninterested in the cloud know that there exists some nebulous concerns around data sovereignty and privacy, but what are they?

Concerns about privacy, security and data sovereignty rules are especially top of mind for organizations in Europe, Canada and other countries that have adopted meaningful privacy laws. Even for organizations in more lax jurisdictions, such as the United States, however, regulatory compliance efforts can seem overwhelming.

Many regulations have required special care and consideration from organizations and IT practitioners. And a number of these regulations have been widely ignored because enforcement has been lax, the consequences amounting to nothing more than irrelevant fines.

However, the next generation of regulations, which are soon coming into effect, have teeth. Unlike previous data regulatory regimes, the privacy, security and data sovereignty regulations coming into play in many jurisdictions around the world are increasingly willing to pierce the corporate veil and send individuals to jail for noncompliance.

From a practical standpoint, this means that security teams are suddenly relevant. Systems administrators and cloud architects have to put serious thought into privacy, security and data sovereignty, lest their designs get shredded by the security team. So what do IT practitioners need to keep in mind?

## The GDPR

Top of mind during any such discussion is always the EU's <a href="https://www.eugdpr.org/" target="_blank">General Data Protection Regulation (GDPR)</a>. Passed on April 27, 2017, the implementation of the GDPR begins May 25, 2018. Implementation marks the start of enforcement activities for noncompliance, and fines top out at 20 million euros, or 4 percent of global turnover, whichever is higher.

The reach of the GDPR is global. It applies to every individual and organization that processes data about EU citizens, even if those individuals or organizations are not located in the EU. Any organization, anywhere in the world, that collects data on EU citizens is required to comply with the GDPR. If it does not, it faces the same consequences as an organization located in the EU.

In practice, there's little the EU can do directly to small businesses located in, for example, the United States that has

jijomathaidesigners / Shutterstock.com

no European assets. But the EU can block that organization from selling into the EU, prevent EU businesses from doing business with that organization and potentially go after any EU assets of that organization's executives, possibly even preventing them from travelling to the EU. Governments tend to get creative, given a reason to do so.

The GDPR gives EU citizens new rights, such as the right to be forgotten. This means they have the right to see whatever data an organization holds on them, have that data changed to be accurate or to have that data deleted.

These changes and deletions apply not only to production data, but to all instances of data for which an organization is responsible. That includes **all** backups that an organization has made, as well as any and all data that the organization has exchanged with contractors.

Larger organizations with more formal commercial ties to the EU will face more direct consequences. EU assets can be seized for noncompliance, and organization executives can be fined or jailed. If you want to do business in the EU, or sell to EU citizens, compliance with the GDPR is highly recommended.

## Practical Concerns

The GDPR is among the most stringent of the new generation of data regulatory regimes. While the crossover isn't perfect, if you've managed to fully comply with both the spirit and the letter of the GDPR, chances are you'll be covered for almost everything else.

The first thing to understand about the GDPR is that it isn't a regulation aimed specifically at how organizations implement IT. The purpose of the GDPR is giving EU citizens control over their personal data. EU citizens are expected to have complete control over their data regardless of which organizations are handling that data, how it's collected or for what purpose.

The GDPR is founded on two basic concepts. The first is privacy by design, and the second is security by design. In practice, this calls for a vigorous—even obsessive—application of the principle of least privilege.

No user—whether that user is a human being operating a computer under a specific user context, or a user used exclusively by computer programs to talk to other computer programs—should ever be able to access more information than is absolutely required to perform its duties. A janitor shouldn't have access to student records, nor should a teacher from one school be able to look up students from another. A commerce application shouldn't be able to look up Social Security numbers, and a financials server has no need to be able to access medical records.

Gone are the days of having a handful of broadly powerful users with high-level access under which IT practitioners can run scripts, backups and other automated processes. As we move into the 2020s, all data access must be deliberate … and restricted. All of the above applies to outsourced IT, as well as on-premises IT.

The GDPR gives EU citizens new rights, such as the right to be forgotten. This means they have the right to see whatever data an organization holds on them

Outsourced IT means not only organizations hired to manage e-mail lists or ecommerce sites. It means regional services providers, Software-as-a-Service (SaaS) solutions and anything that has been put into the major public cloud providers, as well. There's no hiding from the GDPR. If your organization collects or causes to be collected any data on EU citizens, it is responsible for ensuring the privacy and security of that data, wherever it may be in the world, and whomever may have been hired to process it.

Organizations will also need to implement some means of updating their backups when changes or deletions are requested. They'll need to be able to inform all contractors that have been engaged to store or process data that updates or deletions will be required. If an organization handles data for another organization (say a vendor or other member of the supply chain), then that organization will need to make a mechanism available to allow any companies that contract with them. These are not trivial undertakings.

As discussions about GDPR compliance demonstrate, modern compliance regimes require a bit more than token changes to IT security. For most organizations it involves a complete rethink of data handling, both within IT systems and without.

## Encryption

A large part of practically solving privacy, security and data sovereignty issues involves encryption, but "involves encryption" is rather vague. Encryption can be employed in multiple places, each with their own benefits and drawbacks.

Most modern applications, databases, operating system environments (OSEs) and storage solutions can encrypt data as it is stored (data at rest). Many of these can also encrypt data when it's transmitted over the network (data in flight).

> Suffice it to say that encryption should be something that all organizations should be doing, whether or not their workloads live in the cloud.

While describing the intricacies of encryption technologies could fill at least a small ebook, the most important practical concept is that of the Key Management System (KMS). The short version of a KMS is that it stores and manages encryption keys, allowing organizations to control the encryption of their data and workloads, even when they're outsourced to a services provider or public cloud. The long version of key management is rather a bit longer, and I strongly recommend reading The Definitive Guide to Encryption Key Management Fundamentals by Townsend Security in order to gain a more complete understanding.

Suffice it to say that encryption should be something that all organizations should be doing, whether or not their workloads live in the cloud. And unless your organization is a U.S. corporation that only sells to customers in the United States, isn't in a regulated industry, and doesn't process regulated data, then encrypting everything that your organization places in the cloud is an absolute must.

## The Narrowing Gap

Had this article been written two years ago it would have looked a lot different. About 30 months ago my colleagues and I began a journey that lasted almost a year wherein we catalogued every major regulatory regime that affects IT in the western world. We examined them in-depth and we tried to figure out what IT teams would have to do in order to be compliant.

And then along came the GDPR. This was followed by new privacy regulation in Australia, Canada and even in several U.S. states. Several things have become clear in the past two years.

The first is that over the next five to 10 years, non-U.S. regulators are going to find every single loophole that organizations can use to not secure their data and close it. There's a concerted international effort to this effect. The second is that while the U.S. government isn't ready to act on this at a federal level, individual states are more than willing to force the issue.

The net result is that there's no wiggle room left on this topic. Organizations must begin encrypting everything. They must begin implementing the principle of least privilege, even if that means buying new software or changing SaaS solutions to something that implements privacy by design. Organizations must stop collecting all the data they can on their customers and then selling that off or using it for unsolicited marketing purposes.

What's also become clear is that organizations that haven't already started down this path are potentially in deep, deep trouble.

This level of attention to privacy, security and data sovereignty isn't easy. At a minimum it will require having datacenter and cloud implementation designs vetted by professionals. Realistically it should involve annual audits and working closely with services providers to ensure that all data—both production and backup—is appropriately secured.

VIRTUALIZATION & Cloud Review

# 6 WORKLOADS TO CONSIDER WHEN MIGRATING TO THE CLOUD

If you've been tasked with embracing the cloud, how exactly do you get there? Here's a rundown of the basic workloads to consider.

**Contrary to a lot** of the marketing you might encounter, migrating your workloads to the cloud isn't easy. Workloads can migrate to the cloud in a number of different ways, and each has their own challenges. So if you've been tasked with embracing the cloud, how exactly do you get there?

To start with, it's worth understanding the six basic types of cloud workloads.

## IaaS

When IT operations teams think of cloud computing they're most likely thinking of Infrastructure as a Service (IaaS). IaaS allows IT team s to rent resources from a cloud provider that provide only basic functionality, requiring configuration and oversight from operations teams.

IaaS would include the provisioning of a virtual machine (VM) or allowing for the ability to spawn containers. These workload environments are presented in a relatively raw state. A VM may have an Operating System Environment (OSE) installed, and a container will be instantiated on top of a managed underlying OSE, but it's up to operations teams to configure these execution environments.

Once an IaaS environment is configured, operations teams then have to inject applications, configure those applications, as well as configure networking to allow access to the newly created workload. Multiple workloads can be combined to form a service, in which case networking access to the service from the outside likely will only engage a single execution environment's application (such as a load balancer), which will then provide access to other workloads that make up the service.

IaaS solutions might also include provisioning common workloads with a standardized configuration without requiring operations teams to configure the underlying execution environment. These workloads are usually elements of a larger service, such as a database.

Database as a Service (DBaaS) and similar offerings tend to get counted as IaaS, despite not requiring operations teams to configure the underlying execution environment in large part because they aren't of much use to end users on their own. They're considered to be a part of the invisible underlying infrastructure that users don't want to know even exists, and as a result are usually considered to fall into the IaaS definition. As with all technology definitions, however, IaaS is somewhat arbitrary, and the lines blur over time.

Vaclav Volrab / Shutterstock.com

## PaaS

Platform as a Service (Paas) can be thought of as "IaaS for developers."

The purpose of IaaS is to provide pre-canned stacks of workloads that are typically used together in a service. The traditional example is the LAMP stack: Linux, Apache, MariaDB and PHP. (Before MariaDB, it would've been MySQL.)

Today, there are a number of common platforms that provide a multitude of managed execution environments for developers of virtually any language. The goal of PaaS is not merely to make instantiating an environment to develop or run an application easier. PaaS also serves to provide a predictable standard set of execution environments.

Today, there are a number of common platforms that provide a multitude of managed execution environments for developers of virtually any language.

Developers who develop an application on a PaaS platform can spawn an identical version of that platform for production uses. PaaS platforms are typically also fully managed by the cloud provider; execution environments are regularly updated, and they come configured as secure by default.

PaaS environments differ from IaaS in that PaaS is used almost exclusively to make cloud-native applications. Cloud-native applications separate data storage and configuration from the application, application configuration and the underlying application execution environment.

A properly built cloud-native application or service is fully composable, meaning that the entire PaaS platform on which it runs can at any time be torn down and instantiated without impacting the data upon which the application or service operates. This is frequently used as a first-line approach to any detected problems: If an issue is detected, the application environments are destroyed, new environments are stood up, known clean application code and configurations are injected, and then the data is reattached.

## SaaS

Software as a Service (SaaS) is software offered ready for consumption by end users. No IT people need be involved. Enter your credit card information and start using the solution. In the case of SaaS, only application-level configuration (and usually a limited subset of that) is made visible to the end user.

The underlying execution environment, application and configurations are managed by the SaaS provider. They're responsible for security, updates and other basic tasks.

## Hybrid and Multi-Cloud Services

Hybrid services consist of multiple workloads, where some workloads are operating on separate infrastructure from one another. The canonical example of a hybrid service would be a service where some workloads are on-premises and others are located on a major public cloud provider. Individual, none of the workloads accomplish a complete task. Together, they form a single solution.

A multi-cloud service can be either a hybrid service (where data is processed on multiple infrastructures at the same time) or a solution in which data is processed on multiple clouds in sequence. What separates multi-cloud services from hybrid services is somewhat nebulous, but the definition usually involves the amount of data being accessed.

A hybrid service might be something like a data protection solution involving a cloud storage gateway. Data is sent to a local device on-premises (the cloud storage gateway), which then compresses and deduplicates that data, as well as acts as a buffer. The cloud storage gateway then unspools the compressed and deduplicated data to the cloud provider, allowing for backups that occur at a single point in time (say, every night) to take all day to be sent to the cloud.

A multi-cloud solution typically has to make vast quantities of data available across multiple infrastructures. The term is typically invoked when it's deemed economically or temporally infeasible to send the data from one cloud to the next, even if processing the data sequentially is an acceptable approach.

Multi-cloud services are focused on providing a centralized storage solution to multiple cloud infrastructures. This may be accomplished by keeping data storage in a third-

party datacenter that has high-throughput, low-latency connectivity to all of the clouds on which the production workloads in the service will operate.

## Serverless

Serverless is less a workload or service type than it is a glue which holds workloads together. Serverless can be thought of as somewhere between a batch script and a TSR.

Serverless apps are essentially scripts that IT practitioners write, which monitor some type of input, take data from that input when it arrives, pass that data through one or more proper workloads, and then direct the output to a destination. In and of themselves, serverless apps don't generally process data. They merely act as a sort of programmable conveyer belt, shuttling the data from one location to the next.

Now that you know how workloads can be used in the cloud, let's look at what migrating to the cloud might entail.

## Replacing Workloads into Cloud

Arguably the method of cloud migration with the most long-term success is replacing on-premises workloads with an equivalent cloud-provided version. In almost all cases, this means replacing an on-premises solution that has to be managed by on-premises IT teams with a SaaS solution that requires little to no management from on-premises teams.

Among the most popular targets for this sort of migration are financials application and human resources applications. In the small business world, Quickbooks Online has gained a cult following while Saleforce.com has charmed larger organizations and, as a result, has become one of the most powerful technology companies in the world.

Not all applications can be migrated in this fashion. For such a migration to even be possible, a viable cloud-based application must exist in the first place. This isn't always the case, especially when talking about industry-specific applications.

Migrating to a SaaS application is as complex as migrating from one on-premises application to another. In the case of financials applications, getting the data conversion right can take months, or even years. Regulatory compliance issues, logging, monitoring and data protection all need to be considered, as well.

SaaS applications may be presented as an easy-to-consume service, but this doesn't mean that they solve all problems. Care and attention must be taken—especially during the migration phase—to ensure that the new SaaS-based offering will have all the same security, privacy and data sovereignty protections as the on-premises solution did.

## Evolving Workloads into The Cloud

In some cases workloads and services can be "evolved" into the cloud. Microsoft Active Directory and Exchange, for example, can operate in hybrid mode. In the case of these applications, the SaaS offering and the on-premises offering work in lock-step, and allow the migration of users and data to the public cloud over a longer period of time. The goal in this case is to allow compatibility with older applications that are only capable of operating in concert with on-premises infrastructure while moving what can be moved up to Microsoft's Azure cloud.

Arguably the method of cloud migration with the most long-term success is replacing on-premises workloads with an equivalent cloud-provided version.

Not all hybrid solutions are so blatantly unidirectional. Composable workloads, for example, will generally work anywhere. This can allow for the nearly mythical "hybrid cloud bursting," where public cloud instances of a workload are created when on-premises capacity is exceeded.

Hybrid approaches rooted in composability rather than as a vendor-designed migration path have the advantage of being able to move workloads back and forth at will. This has notable benefits when it comes to cost control.

Organizations deploying compsability based hybrid solutions tend to think less about migrating to the public cloud as some sort of end goal and more about using the public cloud as an extension of their on-premises infrastructure. Infrastructure automation solutions like Terraform, as well as configuration management solutions such as Puppet, Chef, Ansible or Saltstack are usually part of such an exercise.

## Failover to the cloud

Data protection solutions can also be part of a cloud migration exercise. For data protection to work in most scenarios a basic level of workload composability is required. At a bare minimum, networking needs to be made composable for most disaster recovery (DR) solutions to work.

Because data protection solutions take care of moving both workloads and data from one infrastructure to another, they're a great way to initially seed a cloud solution to which you intend to migrate workloads. DR failovers are for all intents and purposes a migration of an IaaS (or sometimes PaaS) workload, meaning there's no reason they can't be used as something of an easy button.

> Few organizations will have a single migration path for all workloads, and many organizations will continue having on-premises IT for years to come.

While the failover approach has a lot to offer, there are also drawbacks. IaaS, if done improperly, can be expensive. Similarly, our IT infrastructures aren't simply workloads. There's often orchestration and automation that glues workloads together. Data protection is often part of this—and it should again be emphasized that workloads always need data protection, even when they're in the cloud—but most organizations also have a middleware layer that shuffles data between workloads.

This middleware layer sometimes integrates with on-premises hardware such as printers, point-of-sales terminals or Internet of Things (IoT) devices. In some cases, the middleware may be responsible for taking data out of one workload and presenting it to a system not owned by the organization, but which lives on-premises. The canonical example here is a computer-provided courier service that must physically live in the shipping department.

DR solutions may be able to ignore some of these things because they're designed to operate in a disaster. You don't care much about the courier computer in the shipping department if the shipping department burned down, or is under water. When failing over production workloads using DR tools with the idea of using the destination infrastructure permanently, all of the little intra-workload "glue" really matters.

Over time, as technical debt is resolved, it won't be necessary to move entire workloads. Once everything's composable only the data storage and configuration needs to go anywhere. Of course, once your workloads are fully composable, you're not likely to be thinking about workload migration as a single process or a specific event, either. At that point you'd simply run your workloads wherever makes the most sense at the time.

Like everything else in IT, how your organization might migrate to the cloud depends on a number of circumstances. Few organizations will have a single migration path for all workloads, and many organizations will continue having on-premises IT for years to come. Regardless of whether you see a migration to the cloud as a one-time event, or as a broadening of your infrastructure options, understanding what those options are is always important.

That, and backing up your data. No matter where the production copy happens to live.

VIRTUALIZATION & Cloud Review

# HYBRID AND MULTI-CLOUD ENVIRONMENTS: WHAT YOU NEED TO KNOW

Beyond management applications there are more practical concerns that need to be borne in mind when thinking about engaging in hybrid and multi-cloud IT deployments.

**So you've solved how to get** your workloads into the cloud without getting sued. But now you have workloads in a managed cloud and on your own on-premises datacenter. You might even be looking at using multiple public clouds. How best to go about this?

The answer is that there is no good answer. The solution that we all want doesn't seem to exist. In a perfect world, there would be multiple hybrid multi-cloud management vendors all offering fully independent management capabilities, unified Role Based Access Control (RBAC), as well as inter-cloud workload migration, cost control, backup and disaster recovery. Ideally, they'd also do security and best practice audits against all infrastructures under management, and offer up rich reporting and monitoring, as well.

While there exists a number of solutions that try to accomplish this task, they all fall down in one way or another. Putting to one side that it's difficult to find even one provider that can meet the feature list, infrastructure diversity support is sorely lacking among the extant multi-cloud management vendors.

Every independent multi-cloud management vendor appears to support the two biggest public clouds: Amazon Web Services (AWS) and Microsoft Azure. Most support the third-largest public cloud, the Google Cloud Platform, with a smaller number of vendors supporting the fourth-largest public cloud, the IBM Cloud.

If IBM has trouble getting love, you can imagine what support for smaller cloud providers looks like. Many of the multi-cloud management solutions have hand-waving-class support for OpenStack, but when you drill down into it, they don't have much of a partner ecosystem actually signed up behind that supposed support.

VMware Cloud on AWS makes the occasional showing, but this makes sense, given that VMware is dumping incomprehensible resources into promoting VMware Cloud on AWS. Support for the much more vibrant – and arguably important – vCloud Director (vCD) services

provider ecosystem is virtually nonexistent, as is support for CloudStack-based providers.

## Yes, No, Maybe, Could You Repeat the Question?

Some solutions come close. Cisco CloudCenter appears to support vCD when deployed on-premises, but does not seem to support it from services providers. Morpheus added vCD support in version 3.1.0, though this isn't discussed anywhere on its main Web page, leaving room for questions about just how solid support really is.

In fact, the Web sites of the various multi-cloud management vendors are almost universally awful. Scalr, for example, has a pretty graphic that seems to indicate support for Azure, OpenStack, VMware, GCP and AWS. But VMware what? vCenter on-premises? VMware cloud on AWS? vCD ecosystem services providers? That information doesn't seem to be posted anywhere. (The Sclar Twitter team confirmed that it "has native support for VMware via vSphere [not via vCloud Director].")

The DevOps community has a few more options here. Infrastructure automation overachiever Terraform offers a number of providers, including vCD. And while it is possible to assemble every single multi-cloud management feature discussed here using tools common to the DevOps crowd, you'll end up needing a lot of tools to get the job done.

Naming, shaming and praising vendors in this space could be its own book. The relevant takeaway is that if your journey into the hybrid multi-cloud is predicated on finding a single pane of glass that solves all of your multi-cloud woes, then as of March 2018, you're chasing unicorns.

The net result of this is that in order to successfully do hybrid multi-cloud anything you have one of two paths: take a scripting-heavy DevOps approach, or know ahead of time exactly which infrastructure providers you wish to use, and ask every multi-cloud management application vendor you can find about support for your preferred infrastructure providers.

## Practical Concerns

Beyond management applications there are more practical concerns that need to be borne in mind when thinking about engaging in hybrid and multi-cloud IT deployments. Chief among these are data protection concerns, data lifecycle management and networking.

Cloud workloads need data protection, just like on-premises ones. Achieving this data protection can be tricky. Relying on cloud provider snapshots doesn't provide any more real-world protection than relying on nothing more than snapshots from your on-premises virtualization infrastructure.

Of more equivocal value is the ability of some cloud providers to offer data protection by copying snapshots to a datacenter in another region. There's validity in this form of data protection: The data is in two places, and those two locations are geographically distinct. Unfortunately, the data is not on two different physical mediums, and all data is controlled by a single entity (the cloud provider in question).

Relying on cloud provider snapshots doesn't provide any more real-world protection than relying on nothing more than snapshots from your on-premises virtualization infrastructure.

A better solution would be to use cloud-to-cloud data protection to back workloads up from one cloud provider to another. In this way, those workloads are resilient not merely against the failure of individual infrastructure components or loss of a datacenter. They're resiliency issues that affect the entire cloud provider.

## Data Charges

Moving data and workloads from cloud to cloud is something that should not be taken lightly, however. Many cloud providers — big and small alike — charge far more to remove data from their cloud than they do to put data into that cloud.

These data charges should be a serious consideration when considering cloud data protection, cloud migration, or having workloads located in multiple clouds operate on the same dataset. Not all cloud providers have this restriction. In this case, it's often best to run production workloads on providers that won't break the bank on outbound data, and back those workloads up to a different cloud provider.

Similarly, it's increasingly common practice to run your

cloud data storage from a central third-party services provider with flexible network pricing. These services providers often have high throughput connections to the big four public cloud providers, as well as many of the other large services providers.

If the latency is low enough, workloads can operate in on one public cloud, but store data in another. This is useful in cases where the same dataset needs to be operated on by workloads in multiple public clouds, but where exporting that data from one cloud and then importing it into another isn't economical.

An example of why your organization might do this is if it wanted to use the Bulk Data Computational Analysis (BDCA) tools from multiple public cloud providers. BDCA solutions tend to be proprietary to each cloud provider, meaning that an organization must make the data available to each in turn in order to use them all.

Moving data and workloads from cloud to cloud is something that should not be taken lightly, however. Many cloud providers — big and small alike — charge far more to remove data from their cloud than they do to put data into that cloud.

## Networking

Interconnecting infrastructures rely on networking — whether those are on-premises, in a major public cloud provider, or a regional services provider. Software-defined networking (SDN) such as VMware NSX is increasingly taking the place of traditional VPNs as the ideal solution to stretch and seamlessly connect these infrastructures.

The key buzzword to keep in mind is layer 2 extensibility. In traditional networking each infrastructure would have its own subnets. A workload with the IP address 10.0.0.100 on one infrastructure might not be able to use that same IP on another infrastructure. This can cause a lot of problems for non-composable workloads.

SDN that incorporates layer 2 extensibility can solve this. The networking of all participating infrastructures can essentially be joined into one unified network. While not absolutely required for hybrid or multi-cloud computing, it makes everything much simpler. Especially when disaster recovery enters the mix.

Both hybrid and multi-cloud computing are entirely possible today. Things remain a little rough around the edges, but the requisite bits are there. As always, working directly with services providers makes all of this much easier. It will be several years yet before the multi-cloud management solutions remove the need to work with other humans to smooth out the kinks.

VIRTUALIZATION & Cloud Review

# THE IMPORTANCE OF TRUST WHEN CHOOSING YOUR CLOUD PROVIDER

Careful and considered selection of your cloud partners is essential. Start small, and build experience and trust over time.

**Using cloud computing** is not a license to disengage your brain. Workloads in the cloud still need to be configured, managed and secured. They still need backups. Proper architecture is as important to cloud deployments as it is to on-premises ones, and the cloud comes with the added problem of needing to trust your cloud provider.

The cloud removes a lot of the scut work from IT. Administrators don't have to worry about swapping out dead servers or designing storage solutions. Organizations don't have to build their own datacenters, with all the attendant costs and considerations that go with that. The cloud tends to have huge Internet capacity: the kind that organizations of any size can only dream of for on-premises deployments. The cloud offers real-world value.

But the cloud isn't magic. It's (mostly) run by humans, and humans are fallible. While the cloud does employ a great deal of automation, that automation is designed and implemented by fallible humans. From top to bottom, mistakes get made. Cloud outages occur. Data is lost.

As a result, organizations looking to engage a cloud provider for any reason need to understand that they will still bear a burden of responsibility regarding their own IT design. You need to plan for cloud outages, data loss and disasters just as you would for on-premises. The advantage to using the cloud is that while you may have to plan for these things, you don't have to build all the relevant bits.

ImageFlow / Shutterstock.com

## Recent Cloud Mishaps

Exploring some of the prominent whoopsies of the past year can help put the need to continue to be responsible in context. Cloud provider mistakes have come in different flavors over the past two years, here's merely a sampling of known events:

In mid-2017 Cisco lost customer data in its Meraki cloud. The lost data included customized interactive voice response greetings and undisclosed "enterprise apps." While considered to be a relatively minor data loss event, the outage is notable because the fact that Meraki gear is cloud-managed is the primary selling point. No on-premises management tools to contend with means less effort in getting things set up.

The Meraki outage is one of the more recent examples of the importance of trust in cloud computing. Trust, once lost, is nearly impossible to regain. As vendors ask organizations to move

more and more into the cloud — including the management planes of on-premises devices — having a cloud provider that you can trust becomes more important than ever.

Earlier in 2017 Amazon's S3 storage had a well-publicized outage that was caused by human error on Amazon's side. It reinforced the requirement to perform disaster planning, even when you use cloud computing.

Shortly thereafter, Digital Ocean saw an outage. In this case, someone on Digital Ocean's side used the production credentials in a script meant for testing. Fortunately, no data was lost. The same can't be said for a Netgear outage that occurred around the same time frame; this cloud outage managed to also delete data on customers' local, on-premises NASes. Whoops.

> The point of the cloud is to relieve systems administrators from occupying themselves with "keeping the lights on," and allow them to put their knowledge and expertise to use at higher levels.

A year earlier, Google managed to down its own cloud by applying a patch to the wrong routers. The incident was compounded by a monitoring solution on Google's end that couldn't identify the root cause, ultimately dragging the outage out. Here again end customers were fortunate in that data loss does not appear to have been a result.

Understanding your cloud provider is crucial. Consider the case of PC World KnowHow cloud backup users. In 2016 this cloud backup solution did not keep versioned copies of cloud backups, meaning that ransomware easily destroyed both production and backup copies of data. Others, such as Apple's iCloud, kept files that were supposedly deleted, revisiting this approach only when called out in the media.

The canonical example of not understanding how the cloud works, however, is Code Spaces. Code Spaces was forced to go out of business because it didn't understand how Amazon's cloud worked. A malicious actor compromised Code Spaces' Amazon EC2 user credentials and simply deleted all of the company's production and backup data. Its mistake is now a textbook lesson in why you must engage your brain before using the cloud.

## The Cloud Offers Value

Not all is doom and gloom about the cloud. The cloud does offer very real value. Everything, however, hinges on trust. Well ... trust, and a rational assessment of your own IT capabilities.

Consider GitLab. GitLab walked away from the cloud and ended up with a catastrophe: five out of five of its data protection tools failed, resulting in the loss of a production database and a very embarrassing public debacle. Running away from the cloud in fear doesn't necessarily solve everything.

The point of the cloud is to relieve systems administrators from occupying themselves with "keeping the lights on," and allow them to put their knowledge and expertise to use at higher levels. Systems administrators freed from mundanity can focus on automation, on architecture and on meeting business challenges. This is the real promise of the cloud. When done right, the cloud delivers on this promise.

Trust comes into play in determining how much effort organizations have to put into various layers of resiliency. In a perfect world, cloud providers would be completely trustworthy and no one would ever have to worry about things like backup verification, cloud-to-cloud backup or building disaster recovery solutions for workloads that run in the cloud.

## Extracting Value Takes Effort

Individual cloud providers may be more trustworthy than cloud providers as a whole. This is because some cloud providers—especially regional services providers—are willing to work with organizations one-on-one to ensure that all of an organization's needs are met.

Key to this is transparency; the cloud provider needs to be willing and able to answer how they'll handle any failure scenario, and they need to be open about which scenarios they cannot protect against, meaning that the organization has to engineer their own contingency. To contrast, the responsibility of the organization is to understand enough about their IT architecture and it needs to be able to ask relevant questions, and adequately understand the answers.

There are technologies and approaches to IT that organizations can use to minimize their exposure to IT risk, both on-premises and in the cloud. Composable workloads, overlay networking, microsegmentation, automated incident response and more are all relevant IT concepts. Similarly, application modernization toward the goal of being "cloud native" allows organizations that develop their own to use the cloud in an efficient, secure, resilient and cost efficient manner.

This is a journey, however; one that takes years, if not decades. For organizations seeking to take advantage of all that the cloud has to offer, the best advice that anyone can offer is to be prepared to do the work.

Research, research and more research is required. Careful and considered selection of your cloud partners is essential. Start small, and build experience and trust over time.

Above all, however: Test everything, back up everything and then test it all over again. If your data doesn't exist in at least two places, then it simply does not exist. And if you cannot restore that data from its backups, then once more it does not exist. On-premises or in the cloud, risk management is key to IT, and risk management begins with adequate data protection.

Research, research and more research is required. Careful and considered selection of your cloud partners is essential. Start small, and build experience and trust over time.

VIRTUALIZATION & Cloud Review

# EBOOK: THE JOURNEY TO THE CLOUD

**Every organization's journey** to the cloud is going to be unique, because every organization is unique. Each of us face our own challenges, have a different mix of talent and experience available, and different ideas about how to solve the various problems before us. The broad strokes of such a journey are usually similar, however, and Iland can help.

It is a rare organization that doesn't have some level of technical debt today, and as technology continues to diversify this debt threatens to grow larger. Keeping up with the Joneses requires outsourcing at least some IT, and the most popular place to outsource this IT is the cloud.

The rationale behind cloud engagement is simple: quality IT talent is hard to find. IT talent that understands your business is harder still. Expending that talent with keeping the lights on is wasteful, especially when there are vendors willing to take the mundanities of IT off your hands. Your IT staff's focus is better spent elsewhere.

Like all IT, clouds come with security and privacy considerations. Also like all IT, clouds still require careful thought and consideration, deliberate planning, and judicious selection of vendors.

What the cloud offers, however, is the ability to rise above a focus on the underlying IT infrastructure and focus instead on higher level concerns. These higher level concerns include transforming IT delivery into a service model, managing via policies and profiles, and cost control through workload portability across multiple infrastructures.

In addition to making life easier for IT teams, clouds can offer organizations resiliency and reliability that would be difficult and expensive to obtain with on-premises IT. As more organizations adopt cloud computing, these increased levels of service availability will be the new normal, and those organizations that aren't able to provide them will be at a disadvantage.

With all that cloud computing has to offer in mind, let's look at what it takes to get there from here.

Travel mania / Shutterstock.com

## The Ideal Journey to the Cloud

Few organizations have the opportunity to approach cloud computing without also having to consider existing IT solutions. This, combined with the human factors such as learning to trust cloud providers, and acquiring relevant expertise, mean that embracing cloud computing is a long term process.

For most organizations, data protection is their first step into the cloud. Backups and disaster recovery are an ideal way to learn about cloud computing. Done properly, disaster recovery requires the ability to duplicate one's entire IT infrastructure, offering a great way to learn the ins and outs of cloud computing.

Data protection is also an ideal first step into the cloud because it is also more forgiving to work with than mission-critical workloads. Organizations can run their existing data protection solutions alongside the cloud-backed data protection initiatives they are looking to transition to without impacting running workloads.

Data protection to the cloud gives IT teams the opportunity to experiment, to make mistakes, and to refine approaches. A number of challenges—such as workload migration—have to be solved for data protection to the cloud to work. By the time IT teams are confident in their ability to protect workloads with cloud backups and disaster recovery, they should have the majority of the skills required to start using the cloud for production workloads.

Data protection is also an ideal first step into the cloud because it is also more forgiving to work with than mission-critical workloads.

From here, IT teams can start engaging Infrastructure as a Service (IaaS) solutions, transitioning dated workloads to Software as a Service (SaaS), and building new applications on top of Platform as a Service (PaaS). In the real world, very few organizations end up transitioning all their IT into the cloud, so the ideal journey to the cloud ends with IT teams being able to move workloads between infrastructures at will.

For most organizations, hybrid cloud computing is the future. Not for ideological reasons, or as some vendor-backed slogan, but for reasons of simple practicality. There will likely always be some workloads that organizations must run on-premises, or are simply more comfortable doing so on-premises. Some workloads will be better suited to the cloud.

## The Ideal Hybrid Cloud

The ideal hybrid cloud solution wouldn't require organizations to retrain their IT staff in completely new approaches to IT or completely new management tools. The ability to use what IT staffs already know, perhaps with the addition of a few new cloud-specific features, would greatly reduce the cost, time and stress associated with cloud computing.

Making the jump to cloud computing by tapping into an organization's existing skills base has second order benefits. Security and compliance endeavours become significantly easier, because the tools and experience used with on-premises workloads also apply to cloud workloads. This is most likely to happen when the underlying virtualization platform is the same at the cloud provider end as it is on-premises.

A common platform can also make assessing potential cloud providers easier. Cloud providers will ideally use IT infrastructure that is as good (or better) than the organization uses on premises. Judging the suitability of a cloud provider's infrastructure for specific workloads is vastly simplified when one can compare apples to apples.

A common platform also makes it more likely that workloads can be migrated from on-premises infrastructure to the cloud without issue. Far more importantly, a common platform ensures that workloads can be migrated back on-premises as needed, and without significant effort. One of the most oft-repeated complaints about cloud computing is how difficult getting workloads and data back out of a cloud is once one has committed to it.

If the platform is the same on both ends, then creating a hybrid cloud solution should be reasonably straightforward. In some cases, public cloud infrastructure can even operate as the other end of a "stretch cluster", allowing for workloads to be live migrated between the on-premises infrastructure and the cloud provider. In others, workloads will have to be powered down in order to move, but can move back and forth without difficulty.

Much depends on networking. The ease with which workloads can move back and forth within a hybrid cloud solution depends on how integrated the networking between the two sites can be. Ideally, full layer 2 extensibility between sites would exist, so that workloads at the service provider side can be on the same subnet as workloads on-premises. This allows workloads to move without having to create fully

composable network solutions, which may present a problem for workloads that rely on static addressing.

The ideal cloud solution also has a robust Service Level Agreement (SLA). Cloud vendors which stand by their offerings are important, but so too are remedies that are relevant to the workloads in use. If the SLA states a high level of uptime, but offers no remedy for downtime except an irrelevant amount of discount off of future purchases of cloud time, this may not be adequate. Downtime can cost a lot of money. SLAs with penalties to the vendor high enough to ensure they stick to their marketing numbers are a necessity.

## VMware Enables the Hybrid Cloud

For most organizations, satisfying the requirements of the ideal hybrid cloud requires VMware under the hood. VMware is not only the most dominant data center virtualization provider, it is also the dominant solution in use for data center automation.

Most VMware-based cloud solutions use VMware's vCloud Director (vCD). vCD is reasonably easy to use, and has seen numerous feature upgrades, including the ability to integrate networking between sites. When examining the technical concerns of constructing a hybrid cloud solution, vCD-based clouds meet all practical requirements.

vCD provides administrators with a management interface that is simple, and grounded in the same approach to virtualization and workload management as the standard vSphere interface that administrators are used to. Migration between on-premises infrastructures and the cloud infrastructure is simple as both infrastructures use the same platform: VMware.

vCD APIs are well known, and very similar to vSphere's own APIs. The result is that there exists a diverse ecosystem of third-party software and services that support both on-premises vSphere infrastructures as well as service-provider-side vCD infrastructures.

These third-party solutions include most of the popular backup, recovery and workload replication solutions. In addition, numerous security, regulatory compliance, and auditing solutions exist that support both on-premises VMware infrastructures, as well as vCD-based ones.

The existence of a robust ecosystem around VMware-based hybrid clouds, combined with the ability to use standard VMware VMs creates competition between cloud vendors. This competition lowers prices, but more importantly causes cloud providers to compete on service quality, transparency and support.

VMware-based hybrid clouds have a huge advantage over any competing solution stacks because of VMware's NSX. Reliable, simple networking that spans infrastructures is the backbone of a nearly effortless hybrid cloud, and at the moment, NSX is the best there is.

## The Importance of Cloud Providers

Comparing cloud providers is often like comparing two bowls of mixed fruit. Each cloud provider offers a different set of features, services and support for different prices. Organizations will have to take careful note of what is offered, and at what prices, and determine the relative value of each provider.

Hyperscale cloud providers, for example, tend to make IaaS instances like t-shirt sizes; small, medium, large and so forth. Each size is a predetermined mix of storage, CPU and RAM, and doesn't necessarily reflect the resources actually needed by a VM. Need a VM with 2 vCPUs and 32GiB of RAM? You may be unable to get that without an extra 6vCPUs you don't need, or without a great deal of storage you don't need.

Cloud vendors which stand by their offerings are important, but so too are remedies that are relevant to the workloads in use.

Cloud Service Providers (CSPs), while not all the same, do tend to deliver more transparency and flexibility in their billing options. Greater cost visibility makes avoiding cloud costs easier, and when combined with the wide variety of different SLAs that CSPs bring to the table, makes adopting cloud computing less painful.

CSPs delivering VMware-based clouds are the simplest, most risk-free approach to cloud computing available to organizations using VMware-based virtualization. CSPs differentiate themselves from the hyperscale clouds in part based on the more personalized level of service that they provide.

Unlike hyperscale cloud providers, some CSPs also offer colocation. Though typically a solution for organizations to place servers they own in a managed data center in order to take advantage of that data center's internet connectivity, power resiliency and so forth, colocation can also be important for networking, especially network security.

Many organizations building hybrid clouds still rely on hardware firewalls, intrusion detection systems or other

advanced security solutions. For various reasons—often including regulatory compliance—these organizations require the use of these solutions in whichever cloud environment they choose. This may not be possible with the hyperscale cloud providers, or may be economically infeasible, even if it is possible. The ability to collocate this hardware at an CSP can offer an economically viable solution to adding functionality that is not normally part of the CSP's standard offerings.

## iland Makes Hybrid Cloud Easy

iland is a global, VMware-based Cloud Service Provider. iland's Secure Cloud Services enable your IT team to utilize their existing skill sets and begin taking advantage of the agility and scale of cloud almost immediately. iland's cloud infrastructure is built on enterprise class hardware and supports a number of VMware partners for data protection, network security, encryption, disaster recovery and more. iland offers an assortment of IT services including Backup as a Service, DR as a Service, Infrastructure as a Service and hybrid cloud connectivity.

iland offers both shared and dedicated clouds, and charges shared cloud usage at the virtual datacenter level, instead of the virtual machine level. Customers can choose to engage clouds using reservation, burst, as well as reservation + burst pricing models. Combined, these two approaches offer unmatched control over cloud costs, while allowing individual virtual machines to be assigned only the resources they need, when they need it.

iland's Secure Cloud ConsoleSM provides full visibility and control over the entire cloud-based environment. With the ability to rapidly spin up new virtual machines, modify existing machines, and connect to them from the interface, the cloud console allows administrators direct access to their servers as if they were still on premises. The benefits of the cloud console go beyond just virtual machine management though. Historical performance information, networking configurations, backups and recovery, disaster recovery, even billing and customer support are all accessed from this single pane of glass management console. With a comprehensive set of APIs, everything that can be done through the UI can be invoked through automation scripting allowing existing automation frameworks to carry into the cloud.

Security and compliance are always top of mind at iland. Utilizing the international standards of ISO 27001, ISO 20000, PCI-DSS, ITIL, CSA STAR, conformance with European General Data Protection Regulation (GDPR) utilizing BS 10012:2017 as well as United States standards such as CJIS, ITAR, SOC2 and HITRUST certifications and attestations, iland has ensured that whatever industry you work in, the proper documentation and controls are in place. Coupled with built-in security reporting around vulnerability, network intrusion, malware and virus scanning you can rest assured that the iland Secure Cloud environment is as robust as your own.

Networking is seamless by leveraging VMware NSX edge gateways or your own device and includes firewalls, DNS management, VPN connectivity and everything you need to create a seamless on-premises to cloud experience. With iland Zipline, accelerated speeds to all of the major clouds is available allowing for extremely fast hybrid cloud applications.

iland's world class support, based in Houston, TX and London, UK, is there with you for every step of your journey. In-depth consultative sales process along with a deep onboarding ensures you are as comfortable with your new cloud environment as you are with your own data center. Support is always included and available through the secure console, email or a phone call. iland cloud engineers can help you with everything from managing DNS to invoking recovery and DR.

iland Catalyst is a robust analysis tool that can help you plan for your migration to the cloud. From understanding which virtual machines need to be migrated to knowing how much in the way of resources to allocate, iland Catalyst can provide you an in depth report so you can accurately plan your resource consumption.

The journey to the cloud is a long one, but iland can make it pain free. Go to iland.com to chat with us and learn more.

---

*Trevor Pott is a renowned technologist with 20-plus years of experience serving organizations of varying sizes as an IT practitioner, adviser, consultant and writer. A prolific tech writer and highly sought-after vendor consultant who knows both technology and the people who apply it, Trevor is also a systems administrator with eGeek Consulting Ltd.*

**Check out iland today at www.iland.com.**

iland®