GLASSHOUSE

Whitepaper

Implementing A Multi-tiered Cloud Service Strategy

Using a multi-tiered cloud structure to provide the protection required for critical workloads while reducing overall costs.

The confusion over how to adopt and realize the benefits of the cloud has been intensified by "Everything-as-a-Service", a result of "cloud washing" vendors trying to spin existing technology as something new.

Introduction

Cloud technologies open up new avenues of data and application availability. However, treat all workloads the same and you'll spend too much on expensive cloud technology that you don't need, spend too little and you put the business at risk. Done thoughtfully, a multi-tiered cloud strategy can provide all the protection required by critical workloads, and a cost savings for non-critical workloads.

Commoditizing the cloud

Historically business processes and practices commoditize in the same pattern: An innovation starts out being used by a small group while it's new, then over time consumers start to better understand its use and potential - so adoption increases. At this point, users begin to take advantage of the innovation in order to create value for other uses than what the innovator first intended; improvements are made in order to meet demand from users. Eventually, the innovation evolves to its highest potential and becomes a commodity.

The evolution of cloud services is no exception. Public cloud computing began by allowing consumers to rent and share services over the Internet. However, the confusion over how to adopt and realize the benefits of the cloud has been intensified by "Everything-as-a-Service", a result of "cloud washing" vendors trying to spin existing technology as something new. Still, the list of what is available as a cloud service continues to evolve, but now encompasses:

- Application Programming Interfaces as a service (APIaaS)
- Backup as a service (BUaaS)
- Data as a service (DaaS)
- Desktop as a service (DaaS)
- Disaster Recovery as a service (DRaaS)
- Infrastructure as a service (IaaS)
- Platform as a service (PaaS)
- Security as a service (SECaaS)
- Software as a service (SaaS)
- Storage as a service (STaaS)
- Test Environment as a service (TEaaS)

One of the original benefits of cloud services was inexpensive outsourced IT. However, the unpredictable market of the last few years has provided an added, timely and much-needed business benefit: increased agility. Organizations using cloud services have been able to respond quickly to market demand for their products and services (which in turn increases their IT consumption) or scale back IT consumption as demand wanes during economic declines. Combine a volatile market with cost savings and the potential to improve manageability of just about everything in the IT environment, and you've got a few hundred thousand early adopters nearly giddy with the possibilities of cloud.

The not-so-silver lining

However, the practical and broad adoption of cloud has run into whatever the opposite of a silver lining is: Organizations who have no problem sharing a network and storage with others are not interested in sharing virtual servers due to their vulnerability. While virtualization – the underlying foundation of cloud technology – is revolutionary, virtualization software is complex and in a constant state of evolution. The specific point of distrust for those evaluating cloud services is that hypervisors, the software agents that manage groups of virtual machines, are vulnerable to attack. Organizations with a higher tolerance for risk can fully enjoy the cost savings and efficiency that cloud services have to offer. Organizations that deal with sensitive data, that are regulated by compliance agencies, or that require 24x7 uptime, will have a low-to-zero tolerance for security risks or downtime.

The specific point of distrust for those evaluating cloud services is that hypervisors, the software agents that manage groups of virtual machines, are vulnerable to attack.

Technology and enterprise have responded to the varied requirements for cloud services with different types of clouds tailored to different situations. On one end of the spectrum, public clouds are available for those consumers who simply need the cloud to save them time and money. Those consumers who need workload availability and security can build private clouds. The third category of consumer needs an approach that combines the best of both – a hybrid cloud. However, the choice between cost savings and security isn't always black and white, and a hybrid approach is a complex structure.

Public, private, hybrid, and multi-tiered clouds

Public clouds

Public cloud services share infrastructure and hardware with many organizations and give the consumer little or no control over the underlying technology infrastructure, which is usually built on inexpensive or free proprietary software and therefore offers little service level assurance. Public clouds virtually come with a guarantee that they will fail and most of today's applications are not designed to keep themselves running somewhere else in case of a server failure. They promise the potential of a significant cost savings, yet our experience suggests this is not always the case.

Private clouds

Private cloud services do not share hardware or infrastructure and are the response to the demand for more secure, dependable infrastructure. Private clouds are internal to an organization and provide services over a private intranet or via a private data center. These private clouds can provide the appropriate level of fault tolerance and security required by your organization. However, they tend to integrate incumbent technologies and as a result can be expensive to build and maintain – a necessary evil if security and resilience must be guaranteed.

Hybrid clouds

Early in cloud's evolution there was an all-or-nothing approach to outsourcing IT. Today, there are more options; a hybrid cloud is a potential answer for organizations that want to take advantage of cloud benefits for non-critical or ad hoc workloads and maintain a certain resiliency for critical workloads. Hybrid clouds combine both public and private cloud options and allow more customized security policies and infrastructure; workloads can be allocated to the public or private cloud resource as necessary.

Multi-tiered clouds

A multi-tier cloud is a single public or private cloud built to provide the full range of cloud service levels with corresponding tiers of cost. Because all workloads are not equally important, a multi-tier cloud allows organizations to prioritize their application components and take advantage of the cost savings of lower-priority service levels.

How do you decide which cloud or combination of clouds is right for you? Start by determining how much data, time and revenue your organization can afford to lose.

Determining how much can you afford to lose

The two primary methods of measuring the criticality of IT systems are how much data and time you can afford to lose. Keeping other barriers to cloud adoption equal, they provide a useful mechanism through which to deconstruct cloud adoption; thus allowing you to avoid the cost or risk associated with one-size-fits-all and begin to unlock, more broadly, the true potential of the cloud:

Recovery Point Objective

The first, the Recovery Point Objective (RPO), is the threshold of how much data you can afford to lose after a disaster or major outage. Defining your RPO should take into account your organizations tolerance to risk – it's a top-down decision. However, figuring out your current RPO (if you don't already know it) can be done using the bottom-up method: Start by examining how frequently backup takes place. Since backups can be intrusive to systems, they are not typically performed more frequently than several hours apart. This means that your backup RPO is probably measured in hours or days of data loss.

When it comes to the cloud, an organization must evaluate required service level versus cloud protection price points.

Recovery Time Objective

The second, the Recovery Time Objective (RTO) is the threshold for how quickly you need to have your application service restored. Using these two primary measures will help you understand your cost of downtime and the risk that your organization can tolerate. Next, here's a simple business impact calculation, a way to estimate the average cost per hour of downtime:

Cost Per Occurrence = (To + Td) x (Hr + Lr)

- To = Length of Outage
- Td = Time Delta to Data Backup (How long since the last backup?)
- Hr = Hourly Rate of Personnel (Calculate by monthly expense per department divided by the number of work hours.)
- Lr = Lost Revenue per Hour (Applies if the department generates profit. A good rule is to look at profitability over three months and divide by the number of work hours.)

Finding the right balance of features and price to meet RPO and RTO requirements is one of the most critical things a business can do. When it comes to the cloud, an organization must evaluate required service level versus cloud protection price points.

Benefits of a multi-tiered cloud – unified, on-demand infrastructure

Cloud does not have to be one size fits all. A multi-tiered approach to cloud services allows businesses to take calculated risks while taking advantage of the spectrum of services the cloud offers. Additionally, a unified approach puts cloud service back under the control of the IT department; they're able to provide managers the cost savings and flexibility of cloud service that is appropriate for their workloads while keeping an eye on what's best – and most cost effective – for the entire organization.

A multi-tiered approach also solves the problem of non-cloud friendly application resilience by allowing organizations to take advantage of the n-tier architecture of today's applications. For example, an organization may put a stateless web service front end on a low capability and cost tier equivalent to the cheap public clouds. When it fails, it won't have significant impact because it's easy to quickly start another one. Next, the application logic resides on a middle capability and cost tier, because it has a state but is likely to be clustered and can scale out. Finally, on the backend the model element of the database application could reside on a safe, compatible physical server.

A multi-tiered approach to cloud services allows businesses to take calculated risks while taking advantage of the spectrum of services the cloud offers.

Matching application components and tiers to appropriate cloud service levels

While high availability and zero data loss are reasonable and necessary goals for some application components, protecting all components at the highest level of service required for critical applications can be an expensive mistake. For example, if a healthcare provider is protecting marketing data and applications at the same level they are protecting critical components such as electronic medical records and scheduling, they're wasting a good deal of money protecting non-critical applications and data at an RPO and RTO that isn't necessary. While a healthcare provider may not be able to do without electronic patient records for 24 hours, it's likely that marketing applications and data can wait longer for restoration.

Matching application components and tiers to cloud service level tiers is simply an exercise in analyzing how much data loss and downtime the business can tolerate for each application component. An IT manager must first analyze the environment as a whole, and then look at it in pieces order to ensure critical applications have adequate availability and less critical applications take advantage of less expensive tiers of cloud protection. Determining the RPO and RTO for each component increases the flexibility an organization has when integrating required service levels with the costs of delivering that service level.

Conversely, caution must be taken to ensure cost savings doesn't take precedent over the critical nature of an application component. Protecting mission-critical components at a lower cost service level exposes the business to risk in the event of an operation failure or disaster. Organizations that have to answer to regulatory or legislative bodies have an even smaller margin of error for artificially loose RPO and RTO as fines or penalties for non-compliance can quickly cripple operations.

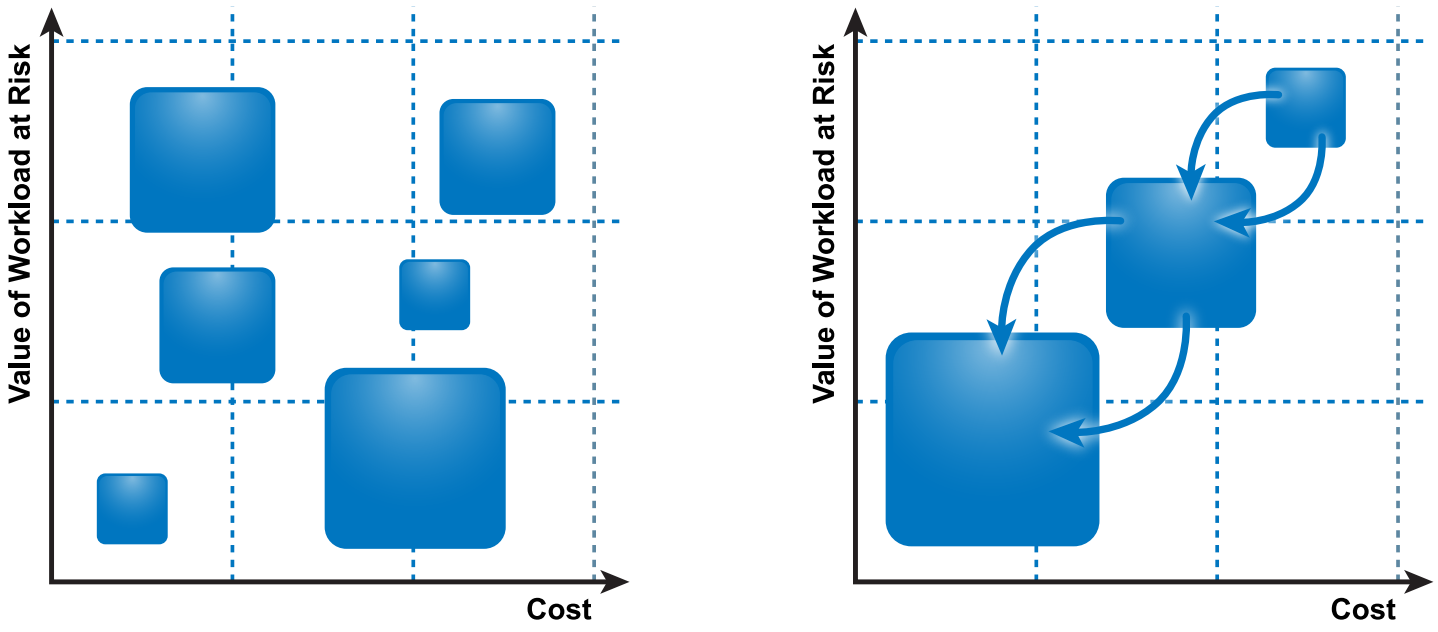

Figure 1: The workload value model

Matching application components and tiers to cloud service level tiers is simply an exercise in analyzing how much data loss and downtime the business can tolerate for each application component.

One of the most common reasons for a risky, artificially loose RPO and RTO estimation is a false sense of security, especially for organizations located outside of natural disaster hot zones. Keep in mind; most downtime is caused by simple human error, not natural disasters. Analysts agree that all businesses fall into one of two categories: those who have experienced an operational failure and those who will.

How to use a multi-tiered cloud for cost savings

Poll application owners about the criticality of their workloads and you're likely come away with the perception that all data and applications in the business are equally important and have a tight RPO and RTO, but that isn't the case.

In order to get the most out of a multi-tiered cloud strategy you must evaluate service levels for components based on the judgment of people who have a broad perspective of costs versus risk for the entire business. Risk assessments from the business side of the organization will identify critical components that affect revenue and reputation, two fundamental elements that successful organizations fiercely protect.

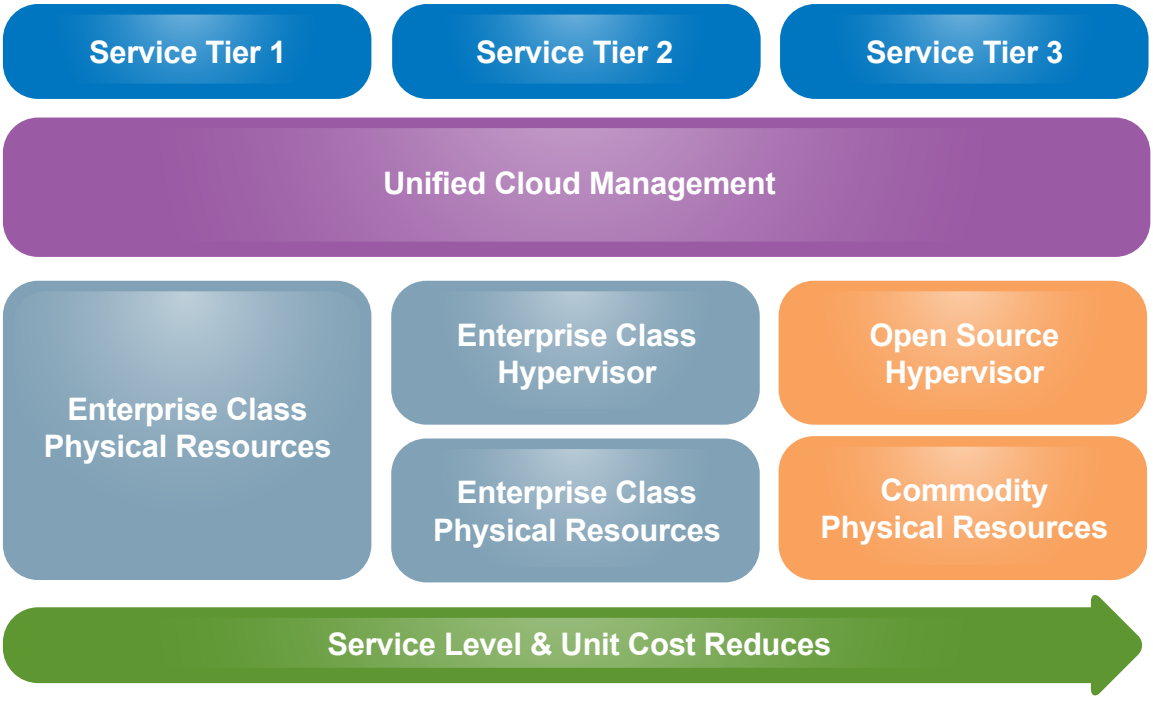

Figure 2: The Multi-tiered Cloud

Conclusion

Navigating the complexities of cloud and “Everything-as-a-Service” can be confusing when it comes to application availability. Do it wrong and you’ll spend more than is necessary or put the business at risk for downtime – causing revenue and reputation loss. Executing a multi-tiered cloud strategy can be the right-sized answer for organizations that are willing and able to prioritize application components based on individual RPO and RTO goals. Done right, a multi-tiered cloud strategy offers flexibility for the level of protection required for critical workloads and the cost-savings of inexpensive cloud services.

About GlassHouse

GlassHouse guides customers through the complexities of cloud, virtualization, storage, backup and security through vendor-independent data center infrastructure consulting and managed services. Glasshouse does not sell any product, a principle that enables us to provide objective recommendations and integration strategies. We consider the people, processes, policies and technology already in place while creating a customized plan that mitigates security and non-compliance risks, improves cost and service efficiency, and enables IT departments to become true service providers for their organization. The depth and breadth of our expertise has been developed through more than 17,500 engagements with more than 12,000 clients. For more information visit www.glasshouse.com or visit the GlassHouse blog for expert commentary on key data center issues. Twitter users can follow us at @GlassHouse_Tech.