



xterity
cloud services

Mission Critical Applications in the Cloud

The 5 Things Every

Enterprise Must Consider

CONTENTS:

- 1 The New Cloud Landscape . Page. 3
- 2 #1 Cloud and Data Security . . Page 5
- 3 #2 Availability Page... 7
- 4 #3 Virtualization Page.. 9. .
- 5 #4 Performance Page 10
- 6 #5 Cost Page 11
- 7 CloudRivo's Xterity for the Enterprise
Page 12

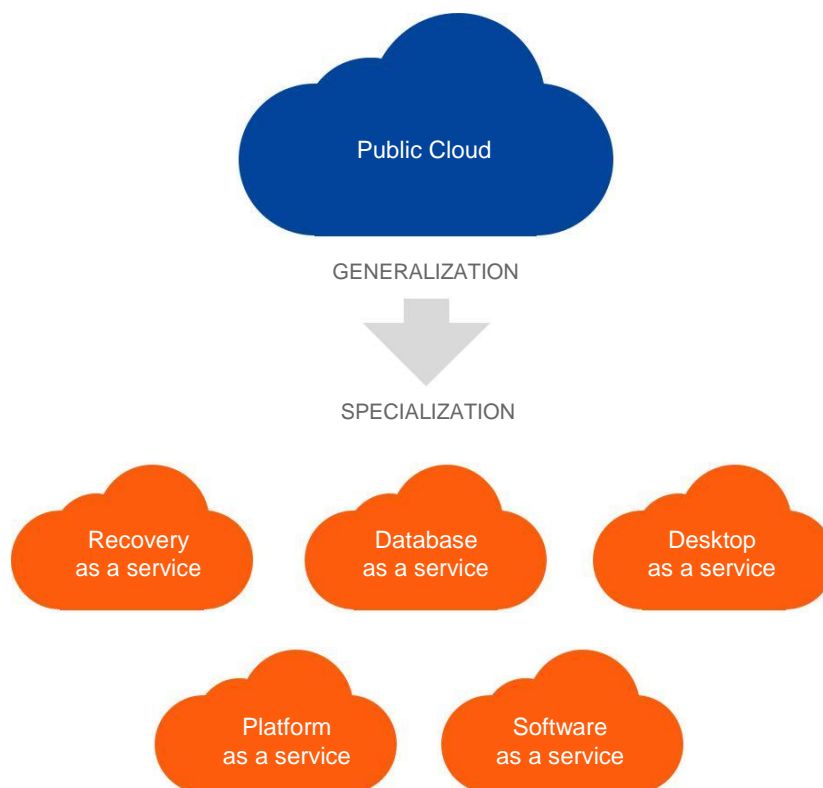
The New Cloud Landscape

1

The IT industry has undergone massive changes over the past decade. Virtualization, cloud computing, the growth of mobile and Big Data have reshaped user expectations and how computing is delivered.

With its ability to provide faster infrastructure access, make applications more affordable, reduce capital costs, and eliminate implementation cost and complexity, the cloud is becoming a key component of most enterprise's IT strategy.

Today, the public cloud market is moving from one of generalization to one of specialization, as the competition heats up and new cloud services providers look for ways to break into the market. There are an increasing number of interesting services, such as recovery as a service, desktop as a service, identity as a service, and mobile back end as a service - just to name a few.



We all know that cloud computing offers significant benefits, but is it ready for mission critical applications?

Mission critical applications are those that support your most important business functions. The failure of these applications can result in the failure of business operations. When this happens, not only does a company lose money but it can also lose customers and tarnish a brand image.

It is not surprising then that today only about 30-40 percent of large enterprises' applications are in the cloud – most of which are internal applications such as collaboration, conferencing, email and sales force automation, according to an IDG cloud computing survey. Concerns about security and availability are the barriers stopping most organizations from taking their mission critical apps to the cloud.

If these concerns can be allayed, there are benefits that the cloud can offer for these applications. The most important are scalability and reliability. In many cases cloud service providers can provide much faster access to additional computing resources, and have implemented a much higher degree of infrastructure redundancy than many individual enterprises have. These same providers can also elastically scale computing resources much more efficiently and cost-effectively, as resources can be shared across their customers.

How do enterprise users decide whether to trust the applications their business depends on to the public cloud? Are they considering the implementation of an internally hosted private cloud or a virtual private cloud from a provider who will dedicate systems to their organization?

Let's consider five important issues in the following chapters.



#1: Cloud and Data Security

Since the earliest cloud offerings, security has been the key concern for all organizations, especially when it comes to mission critical applications and the data it processes.

In fact, security is the primary reason that within the US Federal government, private cloud spending in 2014 is projected to be \$1.7 billion vs \$118.3 million on public clouds. However, in 2010, to increase the adoption of public cloud services by US government agencies, a program called FedRAMP was created to streamline the process of determining whether cloud services meet Federal security requirements. In May 2013, Amazon Web Services passed the FedRAMP cloud security assessment, making it one of the first commercial cloud providers to be certified. This should speed the adoption of public cloud services in the US federal government.



In the enterprise world, while security concerns remain, they are lessening as enterprises get more experience with cloud services. The Future of Cloud Survey from North Bridge Venture Partners and GigaOM of Cloud Survey from North Bridge Venture Partners and GigaOM showed that security as an inhibitor declined year-over-year from 55% of respondents in 2012 to 46% in 2013. That number should continue to decline as service providers focus on enhancing security in order to enhance their attractiveness to new customers and for new applications.

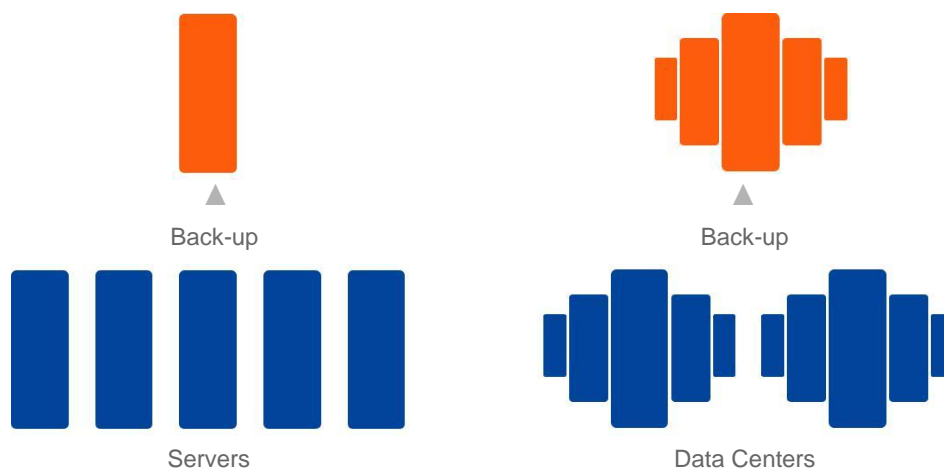
There are many different aspects to security that must be assessed for each cloud provider they consider to host mission critical applications:

- **Data/information security and privacy:** Key considerations include encryption of critical data and ensuring that only authorized users have access to it. Data backup is also an important aspect— both in terms of ensuring backups are done as often as the application demands, and ensuring the security of the backup data and infrastructure.
- **Identity management:** Whether the identity management solution is based on the enterprise's system or the cloud vendor's, it is critical to control access to information, applications and computing resources.
- **Legal issues/compliance:** While both service providers and enterprises must consider legal issues and laws that vary across different countries, enterprises have the bigger burden for ensuring compliance with regulations.
- **Application security:** Whether they use a SaaS application or an in-house developed application, application security must be architected in, and fully tested before deployment.
- **Physical and personnel security:** Beyond online access, it is crucial to know how the cloud providers secures their datacenter and infrastructure, and what personnel will have access to it.



#2 Availability

Availability requirements vary for different applications, but today most need to be available 7/24. There are of course a number of aspects of availability: from the network access, to the computing and storage infrastructure, to the software layers from the operating system to any hypervisor, and up to the application. While the highest level of availability could be achieved by duplicating all systems, storage, networking infrastructure and software, this is an extremely costly scenario that isn't realistic in most cases.



Going one step further, consider disaster recovery—what level of protection does a cloud service provider have against a catastrophic event occurring at their datacenter? How long would it take for applications to be back up and running if disaster strikes?

As noted, availability and DR come at a cost. However, the cloud vendor should have a cost-effective solution for both. An N+1 approach, where individual assets can be the backup for multiple systems, can provide the availability level required at a much lower cost. This is true for both high availability/failover of individual servers or virtual machines, and DR failover across datacenters.



Redundancy alone however is not enough. Delivering services that require 100% uptime requires the ability to know when there is a problem—or that there is a problem likely to come in the future.

This requires a management solution that can monitor and alert both the cloud provider's personnel and the enterprise customer when an issue arises. Users need to know about problems before they are reported on Twitter!

When assessing a cloud provider, it is critical to consider if they can provide the SLA the application needs. Consider whether an alternative provider can take over immediately in case of a failure or a backup system in place internally. In those scenarios, know what that requires in terms of having up to date data.



#3 Virtualization

Cloud computing provides the ability to transform limited, static infrastructure into an IT environment that dynamically scales and delivers the flexibility and self-service that users demand. While server virtualization has been the foundation for most cloud services, not all applications and tiers can run in virtual machines.

There are a number of reasons for the continued usage of “physical bare metal” servers. Many legacy applications weren’t written with portability or virtualization in mind, and may be tied to very specific environments that can’t be duplicated in the cloud. For example, they may run on older operating systems, require out of date drivers, use legacy databases or lack proper security.

Other applications or tiers, like databases or performance-intensive applications, may be best run on bare-metal. These include many financial services trading applications, where revenue is directly tied to how fast a transaction can be completed. Furthermore, regulations and industry practices may mean that public clouds are not yet a good choice for a subset of applications.

A converged infrastructure approach can provide the solution, as it goes beyond simple server virtualization by providing virtualized pools of compute, storage and networking resources. It can provide cloud flexibility for both physical and virtual resources, and enable applications that require full physical servers to gain the same advantages of simplified deployment and administration, automated high availability and disaster recovery (DR) and reduced infrastructure cost.

As with other specialized services, only a subset of cloud providers can provide non-virtualized servers with the flexibility you need in the cloud. Enterprise users need to determine whether the provider they are considering has such a capability.



Modern applications are written in tiers and each likely has a different performance (and availability) requirement.

The front end, or [web tier](#), can be easily distributed to provide the scalability and availability an application requires.

The [middle tier](#) that provides the business logic may also have been written to accommodate a distributed server model for easy scaling and high availability.

The back end, or [data tier](#), is the most likely to require a single physical computing environment, so it is imperative that a defined approach to hosting each tier in terms of scaling and availability be put in place.



Some of the questions user must ask therefore are:

- Can the cloud service enable dynamic scaling of compute nodes when required?
- Does the cloud provider have flexible, short term pricing options so users don't have to commit to a fixed timeframe when additional resources are needed sporadically?
- Will their data volume overwhelm the network pipe to the cloud?
- Is the service provider set up to handle complex tiered applications?
- Do database applications need a bare metal server and if so what does it take to provision and manage it?

Cost factors are always an important consideration. Given the level of systems and network redundancy and availability required for mission critical applications, an analysis of the cost of hosting these apps in house versus the ongoing cost of using public cloud services must be done.

Some of the most important factors to consider in the analysis include:

- **Internal hardware costs**—are new systems required or can users leverage existing infrastructure? Remember to consider the costs for failover and disaster recovery computing, storage and network infrastructure.
- **License costs** for third party software such as hypervisors, databases and applications including costs for any standby/failover instances that may be required.
- **Internal operational expenses** for managing internal infrastructure and software, power and cooling costs etc
- **For capital purchases**, the cost of capital leases and depreciation.

While a quick analysis may indicate that an in-house deployment is cheaper than ongoing cloud expenses, that may not be the case. Cloud providers are likely to have better economies of scale that translates into lower operating costs, higher availability and redundancy, and easier scalability.



The Xterity Cloud Service was built as an enterprise-class cloud service that MPSs/VARs and other IT services firms can use to offer differentiated, profitable cloud service offerings to their clients.

Xterity enables service providers to offer branded, highly available services for the most critical and complex multi-tier applications - without having to build, manage or own their own infrastructure. This frees them to focus on customer relationships, creating custom solutions and maintaining their brand with their end users.

The Xterity Cloud is built on Egenera Cloud Suite (ECS) software. The software provides multi-cloud orchestration and provides resellers with the ability to design, price and manage complex multi-tier and multi-cloud application environments for their end user customers.

The Xterity platform provides a complete business back end for managing the full business life cycle including pricing, managing margins, quoting, billing, monitoring and supporting end users.

Why CloudRivo's Xterity?

- Lets service providers build fully managed cloud services in minutes, delivering profitable, recurring revenue streams
- Enables them to brand their own cloud service and set their own prices and margins
- Expands relationships with current customers and attracts new customers

- Our reseller model means no capital investment
- We provide cloud management expertise and support
- Xterity Cloud services are sold only through Channel partners

Find out more.. Contact us today!



WE CAN HELP

CloudRivo's Xterity was founded with the vision of delivering a full range of dedicated, managed, private and hybrid cloud infrastructure as a wholesale cloud service to the IT reseller ecosystem.

Contact us today.

To learn more about adding cloud services please contact the CloudRivo Team at:

Email: sales@CloudRivo.com

Web: www.CloudRivo.com

Phone: 1-905-910-0094

