

Cloud Computing Overview

June 2016

Table of Contents

- Disclaimer..... 4
- Introduction 4
- What is Cloud Computing? 4
 - Essential Characteristics 4
 - Resource Pooling 5
 - On-Demand Self-Service 5
 - Rapid Elasticity 6
 - Broad Network Access 7
 - Measured Service..... 7
- Service Models 8
 - Infrastructure as a Service (IaaS) 8
 - Platform as a Service (PaaS)..... 8
 - Software as a Service (SaaS) 9
- Deployment Models 10
 - Private Cloud 10
 - Public Cloud 11
 - Community Cloud 12
 - Hybrid Cloud 13
 - Overview 14
- Benefits 15
 - Faster Development of Applications 15
 - Cost Saving 15
 - Improve Operations (Agility and Scalability) 15
 - Disaster Recovery and High Availability 16
 - Modernization..... 16
 - Technological Advantage or Competition 16
 - Security 16
- Risks 16
 - Cost - No economies of scale 16
 - Vendor Lock-In 17
 - Infrastructure 17

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Cloud Computing Overview

- Migrating Applications 17
 - Structure 17
 - Dependency 17
 - Connectivity 17
 - Reliability..... 18
- Virtualization 18
 - Overview 18
 - Sizing 20
- Conclusions 20
- Public Cloud Vendor Comparison 21
- Glossary..... 27
- References 30

Disclaimer

The Toolkit is a diagnostic and planning tool intended to provide recommendations for action based on existing good practice. It does not constitute technical or legal advice and no inference should be drawn as to the completeness, adequacy, accuracy or suitability of the underlying assessment or recommendations. Without limitation to the immunities and privileges of the Bank under its Articles of Agreement and other applicable laws, the Bank shall not be liable for any loss, cost, damage or liability of any kind as a result of this Toolkit or its use.

Introduction

Cloud computing is still a relatively new concept and one that is rapidly evolving to meet ever changing technological demands and needs. This document provides a high level overview of cloud computing as well a comparison of two large public cloud vendors.


Vendor selection can be one of the most challenging parts of migrating to the cloud. The number of vendors and their various attributes can be overwhelming. In addition, vendors frequently do not provide the same metrics and attributes making comparisons even more challenging. The vendor comparison compares two of the largest public cloud vendors in terms of size, global reach, and variety of offerings. This section is intended to provide a guideline for vendor comparisons that governments may undertake.

What is Cloud Computing?

According to the National Institution of Standards and Technology, cloud computing is a model for enabling ever present, convenient, on-demand network access to a shared pool of configurable computing resources (i.e., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction (U.S. Department of Commerce, 2011). In other words, cloud computing can also be referred to as on-demand computing. It is a way for users to get continual access to shared computing resources, such as servers, storage, and sometimes services, as needed.

Essential Characteristics

There are five essential characteristics that define the cloud, as shown in the schematic below.



Resource Pooling

The cloud provider pools all computing resources to serve multiple customers (U.S. Department of Commerce, 2011). These customers can be both external, in the case of a public cloud provider, who might be serving multiple organizations, or internal, in the case of a private data center which may be serving multiple departments. The pooled computing resources are assigned as and when needed, but released and reassigned for other purposes when not being used. Instead of the traditional approach of allocating a single server or amount of space to an application, computing resources are dynamically allocated as needed. This optimization of the infrastructure typically reduces overall infrastructure costs and limits risks such as server failure.

However, the downside to resource pooling is that you have multiple users, groups, or organizations using the same computing resources. This concurrent use of shared computing resources by multiple users, also known as tenants, is referred to as multitenancy. As part of multitenancy, applications still need to be isolated from each other so that problems in one application do not affect others. In addition, access to one application does not mean access is provided to other applications using the same computing resources.

On-Demand Self-Service

Cloud services are provided on request (U.S. Department of Commerce, 2011). Users can request computing resources, such as server time and network storage, as needed, automatically, without requiring human interaction with the service provider. This automation is generally considered more efficient and less error-prone than traditional provisioning processes where requests must be submitted and servers manually set up and configured. The downside is that individuals may request resources whenever they need them, but may not release them when they no longer need them. Automated tools can help with this as well.


This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Cloud Computing Overview

Rapid Elasticity

Computing resources can be elastically provisioned and released, in some cases automatically, enabling applications to scale rapidly in line with demand. The computing resources available for provisioning may be requested in any quantity at any time. This enables more effective utilization of the available infrastructure (U.S. Department of Commerce, 2011). To better understand this concept, it helps to understand what it means for an application to scale.


An application can scale either vertically or horizontally. Vertically means the existing application instance is using more of a specific resource, horizontally means adding additional instances of an application or nodes. An example of scaling horizontally would be going from one web server to three and an example of scaling vertically would be going from 4 GB of memory to 16GB.



Traditionally, computing resources have been allocated with additional contingency in case it is needed. Elasticity refers to the ability for a platform to be dynamic and adaptable as opposed to static. A cloud platform is elastic and can adapt to increasing and decreasing utilization by rapidly expanding and shrinking computing capacity for a given application or application service. In the diagram below the overall application infrastructure that is used is significantly less in the elastic, cloud based approach.

Traditional (Data Center) Approach

Cloud Computing Overview



Broad Network Access


Computing resources are available over the network and accessed through standard devices such as computers or mobile phones (U.S. Department of Commerce, 2011). It is important to keep in mind how a cloud will be reached and what the network availability and bandwidth capacity is before choosing a particular cloud solution.

Measured Service

Cloud systems automatically control and optimize resource use by tracking usage at a level appropriate to the type of service (i.e., storage, processing, network bandwidth, or active user accounts) (U.S. Department of Commerce, 2011). Payment for these services are based on this usage. This is also known as “pay per use”.


Service Models

There are three service models in cloud computing: (i) Infrastructure as a Service (IaaS), (ii) Platform as a Service (PaaS), and (iii) Software as a Service (SaaS).



Infrastructure as a Service (IaaS)

Infrastructure as a Service provides the capability to request (or provision) processing, storage, network, and other fundamental computing resources; the requester is able to deploy and run operating systems and applications (U.S. Department of Commerce, 2011). The requester does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and limited or no control of the networking components (i.e. host firewalls).




Platform as a Service (PaaS)

Platform as a Service provides the capability to deploy onto the cloud infrastructure, user-created or owned applications created using programming languages, libraries, services, and tools **supported by the provider** (U.S. Department of Commerce, 2011). The requester does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.


Cloud Computing Overview

settings for the application-hosting environment. If an application currently resides on an unsupported operating system i.e. UNIX, the application will need to be updated to run on a supported operating system i.e. Linux or take advantage of an IaaS offering where any operating system can be installed.




Software as a Service (SaaS)

Software as a Service provides the capability to use the provider's applications running on a cloud infrastructure (U.S. Department of Commerce, 2011). The applications are accessible from various user devices through either an interface, such as a web browser (i.e., web-based email), or a program interface (i.e. Office 365). The requester does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage. Individual applications cannot be altered but there may be user configuration settings that can be adjusted.



Deployment Models

There are four deployment models to choose from when considering migrating to the cloud.



To understand when to use a particular deployment model as the preferred choice, the models have been compared across five categories – Security, Reliability, Flexibility, Cost, and Vendor Lock-in (degree of difficulty to migrate to a different model if needed in the future). These comparisons are primarily for legacy applications. For each category there is a description and a general score. The score is in relation to the other models.

The table below describes the scoring used in this section.

Icon	Meaning
✓	In comparison to other deployment models, this model is particularly strong in this area.
-	In comparison to other deployment models, this model is neutral or average in this area.
✗	In comparison to other deployment models, this model is weak in this area.

Private Cloud

A private cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple users (i.e. departments). It may be owned, managed, and operated by the organization, a third party, or some combination, and it may exist on or off the premises.

Category	Description	Benefit?
Security	Private clouds are typically more secure than alternatives as the servers are controlled and no other organization has access to them (Pham, 2011).	✓

Cloud Computing Overview

Category	Description	Benefit?
Reliability	Depending upon the infrastructure within a country, a private cloud, especially if there is a direct line connecting the cloud to the government buildings, may be more reliable than alternatives. For example, if the Internet is frequently slow or unavailable during the day during times of high traffic, then making the internet the primary method of reaching key applications may impact day to day business activities.	✓
Flexibility	A private cloud can be geared towards a particular government's needs. It can be built based on the specific requirements that an agency or department needs.	✓
Cost	Higher setup costs, as all hardware (servers, storage, etc.) must be repurposed or purchased. In addition, all future server maintenance would be performed by the government or third party vendor.	✗
Vendor Lock-in	Once an application is virtualized, it is much easier to move from platform to platform. However, a specific virtualization software must be selected when creating a private cloud. This will create a certain amount of lock-in to a specific vendor, but not significantly more or less than any other cloud option.	-

Public Cloud

A public cloud infrastructure is provisioned for use by any organization that wishes to pay for computing resources (U.S. Department of Commerce, 2011). It may be owned, managed, and operated by a business or outside organization. The infrastructure exists on the premises of the cloud provider rather than the users.

Cloud Computing Overview

For the purposes of this toolkit, there is also a deployment model called local public cloud. This term applies to a local public cloud provider whose premises are within the country's borders. This may be the only option if a government has strict laws or policies around the storage and transport of data.

Category	Description	Public	Local Public
Security	For governments in particular, there is a risk of having classified or sensitive data located outside the country's borders. There is also the risk of an external threat (cyber-attack). However, there is also the benefit that cloud providers typically have more skilled employees to dedicate to cloud security.	-	✓
Reliability	Depending upon the infrastructure within a country, a local public or public cloud may be more unreliable than alternatives.	-	-
Flexibility	Local public or public cloud providers may limit the operating systems or databases that they provide. This may require that applications be upgraded to a more recent version of some components before being migrated.	-	-
Cost	Minimal setup and maintenance costs as hardware does not have to be purchased or maintained by the government. There will; however, still be licensing fees.	✓	✓
Vendor Lock-in	While there are companies that specialize in enabling users to move from one cloud platform to another, it does require effort. In addition, once the government gets rid of hardware or requests more capacity than they currently have purchased, it is difficult to move all applications back to government data centers without investing time and money. Thus, going with a public cloud provider results in a certain level of vendor lock-in.	✗	✗

Community Cloud

The community cloud is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (i.e., mission, security requirements, policy, and compliance considerations) (U.S. Department of Commerce, 2011). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises. Community clouds are frequently used by government or educational institutions that consist of a number of different entities (i.e. departments or colleges).

Community cloud is a form of private cloud with multiple tenants where all of the tenants are part of the same parent organization. For the purposes of this toolkit, if multiple departments or ministries decide to utilize the same private cloud then private cloud and community cloud are equivalent. For example, if both the Ministry of Finance and the Ministry of Defense want

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Cloud Computing Overview

to use the same private cloud, but the Ministry of Defense does not want employees from the Ministry of Finance to have access to the defense data, then you have a private cloud with two tenants. This is now a community cloud. The addition of another tenant does impact the security and flexibility of the offering in relation to a private cloud that is dedicated to a single tenant. A private cloud with multiple tenants must be able to offer the technical architectures both need. For example, if the Ministry of Finance has primarily .Net applications running on Windows servers and the Ministry of Defense has primarily Java applications running on Red Hat Linux, the private cloud must now offer both platforms. In addition, appropriate security needs to be in place to ensure that access is restricted to the appropriate individuals. This is especially true if any database consolidation takes place.

Category	Description	Benefit?
Security	Typically all the organizations sharing a community cloud have similar types of data and restrictions. It also enables the organizations to combine their skilled employees. However, the more individuals with access to the cloud from other agencies or departments, the greater the risk of an external attack.	-
Reliability	Depending upon the infrastructure within a country and who owns the community cloud a direct line connecting the cloud to the government buildings, may be more reliable than alternatives.	✓
Flexibility	A community cloud can be geared towards a particular group's needs. However, if a large amount of variety is seen in terms of architecture and technologies across the community, some limits and standardization may be required.	-
Cost	Cost is greatly dependent upon whether the community cloud is owned by a member of the community or a third party. Also, if a large amount of effort is required to standardize the platform and applications across the organizations the upfront cost will be higher.	✗
Vendor Lock-in	Whether owned by one of the members of the community or a third party, any time you standardize options across a group you have a certain amount of vendor lock-in, but not significantly more or less than any other cloud option.	-

Hybrid Cloud

A hybrid cloud infrastructure consists of two or more distinct cloud infrastructures (private, community, or public) that remain separate, but are bound together by standardized or proprietary technology which enables data and application portability (U.S. Department of Commerce, 2011). A hybrid cloud is almost always a combination of public and private and is the combination considered in this section. The most common scenario is a predominantly private cloud that “borrows” computing resources from a public cloud when it experiences spikes in data. One example is taxes. Most people submit their taxes within a one month period of time. During the rest of the year there is minimal use of those tax applications.

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Cloud Computing Overview

Revenue agencies must have enough computing resources to handle the peak demand before taxes are due. In a hybrid environment, that additional demand is handled by public cloud computing resources. This enables the agency to not have to maintain all those additional computing resources on a day to day basis.

Category	Description	Benefit?
Security	A hybrid approach can combine the strengths of both models, allowing the government to keep data under tighter control, but still get some of the benefits of the public cloud.	✓
Reliability	Depending upon the infrastructure within a country, a private cloud, especially if there is a direct line connecting the cloud to the government buildings, may be more reliable than alternatives. Since the public cloud is only used when needed, infrastructure issues will be minimized.	✓
Flexibility	If applications are also using public cloud computing resources, they typically must be compatible with the public cloud. Since public cloud providers may limit the operating systems or databases that they provide, a hybrid approach may require that applications be upgraded to a more recent version of some components before being able to use the public cloud.	-
Cost	Future setup and maintenance costs will be lower than with a purely private cloud approach, since excess capacity will be freed up. Rather than keep computing resources on hand to deal with peak demand, that additional demand will now spill over to the public cloud enabling temporary increases in capacity (Savvas, 2014). However, setting up the hybrid cloud requires expertise in integration and standardization, which can be expensive in the beginning.	-
Vendor Lock-in	Private clouds still require virtualization software. Moving applications from one software to another is difficult and can be costly so the government could be "locked-in" to the vendor of whatever software is chosen. Changing the public provider once a hybrid solution is setup can also be challenging.	✗

Overview

All four deployment models have different attributes making them better fits for some organizations than others.

Category	Private	Public	Local Public	Community	Hybrid
Security	✓	-	✓	-	✓
Reliability	✓	-	-	✓	✓
Flexibility	✓	-	-	-	-

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Cloud Computing Overview

Category	Private	Public	Local Public	Community	Hybrid
Cost	✗	✓	✓	✗	-
Vendor Lock-in	-	✗	✗	-	✗

It should be noted that not all organizations should move to the cloud. Before selecting a deployment model, an organization first needs to consider the benefits and risks of moving to the cloud in the first place.

Benefits

Cloud computing has opened up new possibilities and enables numerous potential benefits, including significant cost savings, faster innovation, and greater flexibility. The following are the common benefits gained from cloud system implementation.

Faster Development of Applications

Cloud computing allows applications to be created and implemented faster. For many governments and organizations it can take weeks, if not longer, to order new servers, set them up, and then build a new application. A cloud system would enable computing resources to be available within hours instead of weeks (Rodier, 2011).

Cost Saving

Infrastructure is expensive to purchase, to operate and to maintain. Cloud services are typically pay as you go, or “on-demand”, which allows end-users to utilize computing resources as needed. It maximizes the utilization of computing resources and reduces the operation and maintenance costs especially during non-peak times. Cost savings are impacted by current IT expenditure, current hardware life cycles, and which deployment model is chosen.

Improve Operations (Agility and Scalability)

Limited computing resources can prevent applications from running as quickly as they could or from running at all if the resources are needed for other applications. For example, a government has a processor intensive census program that runs once every ten years and runs on the same server as an application that shows who is eligible to vote. It may not be possible to generate a list of voters and process the census results at the same time. The cloud can help by automatically supplying additional computing resources during heavy system use.

Growth can also exceed a system’s capabilities. Perhaps in the past most citizens went to their local government office to apply for benefits or get a driver’s license, but with the growth of mobile phones, they can now reach these applications online. This sudden spike of usage may require more processing power than was originally planned for or purchased. Without the cloud, such a spike of usage might cause the system to crash or become inaccessible. With the cloud, additional computing resources are added as needed and removed when no longer required. (Microsoft, 2011)

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Cloud Computing Overview

Disaster Recovery and High Availability

Many public cloud service providers have data centers located in multiple locations. This provides a failover location in the event that the primary location becomes unavailable due to a security event, natural disaster, or human error. This capability keeps the government operating seamlessly.

Modernization

Many governments have servers with a variety of software components on them. There may be multiple versions of Linux or Windows operating systems, the same for different versions of databases, or even programming languages. Moving to the cloud typically gives governments the opportunity to standardize their technology architecture across the government or across a department. This increases the ease of maintenance and the ability to add additional features and functionality to applications going forward.

Technological Advantage or Competition

Governments have a mandate to provide services to their citizens. As part of pursuing this mandate, government may consider implementing a cloud strategy. Alternatively, a government may consider implementing a cloud strategy in order to gain or maintain a perceived technical advantage. This advantage could be in either the public or private sector. A government may work to build demand or skills in the area of cloud computing in order to encourage the development of certain skills or products in the private sector.

Security

Major public cloud service providers have their own security protections against internal and external threats. They also support top-line security protocols commonly used. While anything you put on a public server is at higher risk than a computer not connected to an external network, public cloud service providers have security expertise, operation expertise, and are typically up to date on the latest security technologies.

Private clouds have a certain level of security, especially if they are directly connected to the users they serve rather than accessed via the Internet. However, organizations using private clouds generally have a smaller skilled security team than a public cloud provider would.

Risks

Cost - No economies of scale

There are economies of scale that come from owning an entire data center. Adding one more server is cheaper than the first one was. In the cloud, every CPU and GB needed will cost the same, whether you use 200 or 200 million. Savings are greatest if there are large spikes in usage that cause storage or servers to sit idle when not in use. In the cloud, you only need to pay for those additional computing resources when used. This can also make it more

Cloud Computing Overview

challenging to predict monthly costs. Sudden increase in usage of an application can result in a sudden jump in costs.

Vendor Lock-In

Whether the decision is to build a private cloud or go to a public cloud, there will be a certain amount of vendor lock-in. The degree of lock-in varies, particularly when it comes to deciding to move out of a public cloud. Once you exceed existing computing resources, it is much harder to leave the cloud. This should be considered if you think you might need to make changes in the future due to data or other concerns.

Infrastructure

If the network infrastructure is unreliable or is already highly utilized then moving to the cloud may be too much of a burden on the existing infrastructure. It could cause applications to crash or be inaccessible. In such situations the network infrastructure must either be upgraded before considering a move to a public or hybrid cloud or, alternative, a private cloud on a dedicated line should be considered.

Migrating Applications

An important step in planning for a cloud implementation is deciding which applications to move. Not all applications should be moved to the cloud. There are many attributes that are considered in the application assessment, but some of the most important categories to consider are structure, dependency, connectivity, and reliability.

Structure

A large, single-tiered legacy application typically isn't a good fit for the cloud. In a single-tier application the user interface, business logic, and data storage are all located on the same machine. While these applications are typically the easiest to design, they are also the least scalable. Efficiencies are gained when an application is scalable and the load can be spread over several instances. This also helps with disaster recovery as it enables a failure in one part of the system to be mitigated without affecting other parts of the system.

Dependency

Applications that depend on specific hardware—such as a particular chip set or an external device such as a fingerprint reader—might not be a good fit for the cloud, unless those dependencies are specifically addressed. Similarly, if an application depends on an operating system or set of libraries that cannot be used in the cloud, or cannot be virtualized, that application should not be moved to the cloud.

Connectivity

Applications that interface with or use computing resources that will not be reachable from the cloud, including other applications or storage, are typically poor candidates for migration. For

Cloud Computing Overview

example, if tax data cannot be moved to the cloud, you might not move an application that accesses the tax data frequently throughout the day. In some situations, these issues can be resolved with a custom network setup, but how well this works depends on the chosen cloud environment.

Reliability

Applications by their nature are not perfect, but the more reliable an application is, the longer it can run before encountering a problem. Applications that are known to be unreliable should be reviewed as a possible candidate for rewriting or replacing, since known functionality issues may become worse when migrating an application to a new platform. Trying to migrate an unreliable application may not only increase the effort required to perform the migration, but also fail to achieve the benefits of moving to the cloud.

Virtualization


Overview

Cloud computing is built upon the ability to virtualize applications, regardless of the deployment model selected. Understanding virtualization is key to understanding how pricing works in the cloud. A high level knowledge of this area will enable the creation of more accurate estimates and thus better, and more cost effective, utilization of cloud computing resources. It will also assist with the building of a business case around implementing a cloud computing system.

When researching cloud providers and other various cloud service offerings there will be frequent references to virtual central processing units (vCPUs) and virtual cores (vCores). These components differ from their physical counterparts in a manner that is not always very straight forward. NOTE: Amazon Web Services (AWS) uses the term vCPU whereas Azure uses vCore. Conceptually, they are the same.


The main goal when virtualizing a server is to be able to run multiple applications on the same server. Each application has its own space, or virtual machine, on the server. One way to look at this is to think of a physical server as a house. Each room is a virtual machine and each member of the family, or application, gets their own room. The software that enables the creation of these rooms is called a hypervisor. A hypervisor is a piece of software, hardware, or firmware that creates and runs virtual machines. The hypervisor can either be installed directly on the server or on top of the operating system running on the server. The following diagram shows how three applications running on a virtualized server might look, depending on where the hypervisor is installed.

Cloud Computing Overview



Once a hypervisor is installed, either on the host operating system or directly on the server then the hypervisor manages the physical computing resources (i.e. CPU, memory, etc.) and allocates computing resources to create virtual machines instances upon user request. The guest operating system is installed on the virtual machine instance. Applications can then be installed on the guest operating system and accessed by users.

An increasingly common practice is to take virtualization to the next level and build containers that can easily be moved from server to server. Instead of rooms, the server now has multiple houses and each house can be picked up as a single unit and moved somewhere else as needed. The following diagram shows how container-based virtualization is delivered from a physical server.



Unlike traditional virtual machines, containers do not have a guest operating system installed, but it does require that the physical server have a host operating system. The container itself contains the application in addition to all the components needed for that application and uses

Cloud Computing Overview

the host operating system. This means there is less wasted computing making for a more efficient system, and is also easier to move when needed.

Sizing

When taking applications that currently reside on a physical server and moving them to a virtual machine, it can be challenging to determine how much of various computing resources (i.e. storage, memory, CPUs, etc.) to assign to the application. The recommended approach is to determine what your peak utilization of your current resources over a period of time (ideally 12 months). If that is not possible, then request the same cloud computing resources as the current physical server and monitor the application for the next 12 months to determine utilization, and refine any budget estimates. Based on computing resource usage, the computing resources can be scaled either up or down.

Conclusions

Increasingly, citizens expect that they can complete tasks online rather than going into an office and waiting in line. In addition, the amount of digital data is growing across the globe and is expected to continue to do so. The ability to take advantage of this data and use it to help improve efficiencies within the government and provide better services to citizens is driving many governments to consider cloud platforms. Cloud has the possibility to enable government employees to work from anywhere and citizens to get access to information from their phones or homes. It can enable governments to quickly deploy applications and new functionality.

While cloud has the power to connect, it also comes with risks. Moving data outside of secure locations opens it up for attack. This can be especially true if a limited number of employees with skills in security has led to the development of applications that are particularly vulnerable. Legacy applications that were not originally designed for the cloud may have to be updated, a potentially time consuming undertaking.

It should also be noted, that while much focus is placed on the potential cost savings of cloud, much of those savings are difficult to quantify. Many benefits of cloud enable governments to avoid costs in the future. For example, the implementation of a scalable infrastructure can reduce future capacity costs, and faster development of applications reduces development costs. However, these costs do not reduce the current IT budget, and are sometimes overlooked (Neville Cannon, 2015).

The preferred deployment model and path to implementation will be different for every country, and possibly even differ by departments or ministries within the same country. It may be that an agriculture application can move to the public cloud, but a finance application should consider a private cloud. Then the government must decide if everyone should use the same solution or if there should be multiple solutions. A Cloud Readiness Assessment will provide insight into the current state of a country, and will help provide insight into where a country is now, and what recommendations there are for the future.

This paper, created by The World Bank in collaboration with Accenture, is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Public Cloud Vendor Comparison

If a government decides to go with a public cloud setup then the next step is to determine which vendor to select. Vendors usually have multiple offerings, and it can be challenging to compare vendors. Comparisons are typically further complicated by different vendors using different terminology and units. It is recommended that, even if deciding to pursue a private cloud, governments still assess public cloud vendors to determine a baseline of offerings and service level agreements that they may wish to provide. In order to assist with any future comparisons governments may undertake, a vendor comparison can be found in this section for reference.

Azure and Amazon were chosen due to their breadth of services and geographic offerings. This report is not recommending one vendor over another, but only providing an example of a vendor assessment to provide guidance to governments on developing vendor requirements for their own vendor assessments.

Price is a key factor, especially as it can differ per region. Unlike private clouds, public clouds are not fully customizable. Pricing can vary depending upon the components and services provided by the public cloud service provider and also how the government utilizes those computing resources. A rough baseline for public cloud pricing can be found in the tables in this section.

At this time, there is no data center in Africa for Azure or Amazon, so it is recommended that African countries consider either using a data center on the European continent or a local cloud provider. If a local provider is selected, it is recommended the provider be assessed based on the general concepts and specific recommendations outlined in this report. Please be aware that actual pricing can vary based on utilization and contracting (i.e. predicted infrastructure usage, upfront payment, transaction volume, sizing, etc.).

The tables in this section are a representative list of various options and pricing for Amazon and Azure at a specific point in time, it is not comprehensive and further investigation should be done before selecting a provider. January 29, 2016

Key Differences – Azure and Amazon

Type	Amazon	Azure	Advantage
Availability	Amazon supports high availability across data centers. Services such as load balancing, virtual network, and auto-scaling spans the region	Azure supports high availability within a data center Services such as load balancing, virtual network, and auto-scaling spans the region	Amazon

Cloud Computing Overview

Type	Amazon	Azure	Advantage
Load Balancing	Supports load balancing based on IP address (layer 4) and application performance (layer 7) and provides metric-driven load balancing	Supports load balancing based on IP address (layer 4) and application performance (layer 7) and provides sophisticated load balancing policies	Tie
Virtual	Virtual Private Cloud (VPC) which supports Flow Logs which logs relevant traffic for storage and analysis	VNet to VNet (virtual network)	Amazon
Network	Direct Connect provides faster port speed than Azure however Amazon charges extra for a redundant port	Express Route has redundant ports by default	Azure
Auto Scaling	Has auto scaling provisions, terminates instances based on configured policies, and replaces unhealthy instances automatically	Automatically replaces unhealthy instance (service healing). Auto-scaling also supports both time and load-based scale up and scale down.	Tie
Compute	EC2 is billed by the hour	Virtual Machine is billed by the minute, but is slightly more expensive on average	Tie
Storage	Allows requestor to choose the input/output operation per second (IOPS)	Has more predefined IOPS level	Amazon
Security	Provides both server-side and client-side encryption options	Provides both server-side and client-side encryption options	Tie

Map of Major Data Centers – Azure and Amazon



Regional	Amazon	Azure
Asia & Pacific	Tokyo, Japan Beijing, China Singapore, Singapore Sydney, Australia India (Coming soon) Ningxia, China (Coming Soon) South Korea (Coming Soon)	Hong Kong, Hong Kong Singapore, Singapore Saitama, Japan Osaka, Japan Sydney, Australia Melbourne, Australia Pune, India Chennai, India Mumbai, India
Africa	None	None
Europe	Ireland Frankfurt, Germany	Dublin, Ireland Amsterdam, Netherland
North America	Northern Virginia, United States Oregon, United States Northern California, United States Ohio, United States (Coming Soon) Canada (Coming Soon)	Iowa, United States Virginia, United States Illinois, United States Texas, United States California, United States
South America	São Paulo, Brazil	São Paulo, Brazil

General Comparison

Category	Description	Amazon	Azure
Container Support	Container is an image that contains the complete file system in order to run software. It includes code, runtime, system tools, system libraries and all other components you can install on a server. This will allow environment and component consistency.	✓ EC2 Container Service	✓ Azure Container Service
Analytics (Big Data)	This feature will enable the processing and analysis of large amounts of data to reveal patterns, trends, associations, and other information readable by human.	✓ Elastic Map Reduce (EMR)	✓ - HDInsight (Hadoop) - Azure Data Lake
Compute Service	This service provides the computing power. It comes with different operating system and other services such as storage and network.	✓ - Elastic Compute Cloud (EC2) - Amazon Elastic Beanstalk	✓ - Virtual Machine - Cloud Service - Azure Websites and Apps
Desktop Service	This service provides virtual desktop service where you have your desktop computer in the cloud and access it via the internet.	✓ Amazon Workspace	✓ Azure RemoteApp
Hybrid Cloud Storage	This allows on premise applications to access storage which is located in the cloud system. It makes data growth management, data management, and backup (disaster recovery) easier.	✓ AWS Storage Gateway	✓ StorSimple
Load Balancing	A load balancer distributes network or application traffic across a number of servers. Load balancers are used to increase capacity (concurrent users) and reliability of applications.	✓ Elastic Load Balancing	✓ Azure Resource Manager (ARM)
Managed Deployment	This service automates code deployments, enabling you to deploy reliably and rapidly. The service allows you to launch and track the status of application deployments.	✓ AWS CodeDeploy	✓ Visual Studio Team Services

Operating System Comparison

Cloud Computing Overview

Type	Amazon	Azure
Linux	CentOS 6.0+ / 7.0 Debian 8.0+ Red Hat Enterprise Linux 6.0+ / 7.0+ SUSE Linux Enterprise 11+ / 12+ Ubuntu 12.04 / 14.04 FreeBSD 9.0+ / 10.0+	CentOS 6.3+ / 7.0+ CoreOS 494.4.0+ Debian 7.9+ / 8.2+ Oracle Linux 6.4+ / 7.0+ Red Hat Enterprise Linux 6.7+ / 7.1+ SUSE Linux Enterprise 11 SP3+ / 12+ Open SUSE 13.1+ Ubuntu 12.04 / 14.04 / 15.04 / 15.10
Windows	Windows 2003 R2 Windows 2008 R2 Windows 2008 Windows 2012 Windows 2012 R2	Windows 2008 R2 Windows 2012 R2
Virtual Desktop	Windows 7 with MS Office, Trend Micro and utility bundles	Not Supported

Network Comparison

Type	Amazon	Azure	Remark
Virtual Network	Amazon Virtual Private Cloud (VPC)	Virtual Network	This service enables you to establish a private network (closed and security enhanced). This network is logically (rather than physically) isolated from other networks.
Direct Connection	AWS Direct Connection	Express Route	This service enables you to directly connect to the cloud directly from your premises (office or data center) over vLAN which means you can control bandwidth throughput, and keep a more reliable connection than internet-based connections
DNS	Amazon Route 53	Azure DNS	Domain Name Server (DNS) is used to translate domain names to IP address (like yellow pages). This feature enables users to quickly access applications and infrastructure in the cloud.

Database Comparison

Type	Amazon	Azure	Remark
Relational Database	Amazon Relational Database Service (RDS)	Azure SQL Database	Both Amazon and Azure provide Database as a Service (DaaS) options. Amazon provides more database options as part of their DaaS.
NoSQL Database	DynamoDB MongoDB	DocumentDB MongoDB	NoSQL databases use means other than tabular relationships to organize data and are mostly used to store large amounts of unstructured data.
Data Warehousing	Amazon Redshift	Azure SQL Data Warehouse	Data warehousing is used to run data analysis and produce reports. It stores current and historical data.

Operating System Pricing Comparison – Azure and Amazon

Data Center Location	Amazon – Linux	Azure - Linux	Amazon - Windows	Azure - Windows
Japan	\$0.08	\$0.11	\$0.10	\$0.158
Australia	\$0.08	\$0.116	\$0.10	\$0.186
Singapore	\$0.08	\$0.116	\$0.10	\$0.174
EU Region #1 – Ireland	\$0.056	\$0.094	\$0.076	\$0.15
EU Region #2 - Varies	\$0.06	\$0.102	\$0.08	\$0.162
Brazil	\$0.108	\$0.116	\$0.128	\$0.178
US West	\$0.052	\$0.094	\$0.072	\$0.154
US East	\$0.052	\$0.088	\$0.072	\$0.148

- Amazon EU Region #2 - Frankfurt
- Amazon - 2 vCPU / 4GB RAM
- Azure EU Region #2 – Netherland
- Azure - 2 cores / 3.5GB RAM

Storage Pricing Comparison – Azure and Amazon

Data Center Location	Amazon - Storage (500TB)	Azure - Storage (500TB)
Japan	\$0.0313 per GB	\$0.0228 per GB
Australia	\$0.0313 per GB	\$0.0251 per GB
Singapore	\$0.0285 per GB	\$0.0228 per GB
EU Region #1 - Ireland	\$0.0285 per GB	\$0.0228 per GB
EU Region #2 - Varies	\$0.0308 per GB	\$0.0228 per GB
Brazil	\$0.0387 per GB	\$0.0309 per GB
US West	\$0.0285 per GB	\$0.0228 per GB
US East	\$0.0285 per GB	\$0.0228 per GB

This table compares S3 storage on Amazon and Locally Redundant Storage (LRS) on Azure

- Azure EU Region #2 – Netherland
- Amazon EU Region #2 – Frankfurt

Network (traffic) Pricing Comparison – Azure and Amazon

Traffic	Amazon – DNS Query	Azure – DNS Query
First One Billion Queries / month	\$0.700 per million queries	\$0.540 per million queries
Over One Billion Queries / month	\$0.350 per million queries	\$0.375 per million queries

Traffic	Amazon – Health Check	Azure – Health Check
Internal	\$0.50 per health check / month	\$0.36 per health check / month
External	\$0.75 per health check / month	\$0.54 per health check / month

Health check is a process by which network traffic is sent to check if an instance or node is active. This is required in order to setup load balancing and high availability.

Data Center Location	Amazon – Gateway	Azure - Gateway
Japan	\$0.062 per hour	\$0.036 per hour
Australia	\$0.059 per hour	\$0.036 per hour
Singapore	\$0.059 per hour	\$0.036 per hour
EU Region #1 - Ireland	\$0.048 per hour	\$0.036 per hour
EU Region #2 - Varies	\$0.052 per hour	\$0.036 per hour
Brazil	N/A	\$0.036 per hour
US West	\$0.045 per hour	\$0.036 per hour
US East	\$0.045 per hour	\$0.036 per hour

A gateway is a network point that acts as an entrance to another network. It enables the end users to access the system over the internet or enable a hybrid cloud system. This table compares a NAT Gateway in a VPC on Amazon and basic VPN or ExpressRoute Gateway on Azure.

Glossary

The following terms appear in this document and in the assessments.

Category	Term	Definition
General	Multitenancy	The concurrent use of shared computing resources by multiple users, also known as tenants
General	Private Cloud	A private cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple users (i.e. departments). It may be owned, managed, and operated by the organization, a third party, or some combination, and it may exist on or off the premises.
General	Public Cloud	A public cloud infrastructure is provisioned for use by any organization that wishes to pay for computing resources. It may be owned, managed, and operated by a business, academic institution, government organization, or some combination. The infrastructure exists on the premises of the cloud provider rather than the users.

Cloud Computing Overview

Category	Term	Definition
General	Hybrid Cloud	A hybrid cloud infrastructure consists of two or more distinct cloud infrastructures (private, community, or public) that remain separate, but are bound together by standardized or proprietary technology which enables data and application portability. Normally, it is a combination of public and private.
General	Community Cloud	The community cloud is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (i.e., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
General	IaaS	Provides the capability to request (provision) processing, storage, networks, and other fundamental computing resources, but the requester is able to deploy and run anything they want, including operating systems and applications.
General	PaaS	Provides the capability to deploy onto the cloud infrastructure, consumer-created or owned applications created using programming languages, libraries, services, and tools supported by the provider.
General	SaaS	Provides the capability to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either an interface, such as a web browser (i.e., web-based email), or a program interface (i.e. Office 365). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage.
General	High Capacity Link	Also known as an internet gateway, this is a primary or backbone link outside of a country to the Internet.
General	Internet Service Provider (ISP)	An organization that provides services for accessing, using, or participating in the Internet.
General	Public server	A server that is owned by a third party and accessible via a public network, such as the internet (i.e. AWS or Azure).
Application	Criticality	<p>Critical - the application cannot afford to have more than 2 hours of downtime and there is no alternative for this application. Also, application that is classified as 'critical' by internal policy</p> <p>High - the application cannot afford to have more than 4 hours of downtime and there is alternatives for this application (i.e. manual entries)</p> <p>Moderate - the application can cannot have more than 12 hours of downtime</p> <p>Low - the application can have more than 24 hours of downtime</p>
Application	Single Tier	Single tier, sometimes called one-tier, architecture involves putting all of the required components for a software application or technology on a single server or platform. The alternative is multi-tiered architecture or the three-tier architecture that is used for some web applications and other

Cloud Computing Overview

Category	Term	Definition
		technologies where various presentation, business and data access layers are housed separately.
Application	Static Attribute	Any source code component that has been hard coded (i.e. hard coded IP address and hostnames).
Data	Personally Identifiable Information	Personal information is data that can be used to identify the individual (i.e. name, passport number, phone number).
Data	Sensitive data	Sensitive data refers to data that is deemed sensitive by the owner of the data (i.e. classified government documents).
Data	User Information	User information is data that belongs to an individual but cannot be used to identify them without additional information (i.e. ID, position).
Functional	Service Level Agreement (SLA)	An agreement that sets maximum or minimum targets for various metrics. For example there may be a service level agreement in regards to how quickly technology support must respond to defects of various severities.
Infrastructure	Demilitarized Zone (DMZ)	A physical or logical subnetwork that contains and exposes an organization's external-facing services to a larger and untrusted network, usually the Internet.
Infrastructure	End of Service (EOS)	The expected retirement date of a server based on internal policy or other methods.
Infrastructure	Input / Output Operations Per Second (IOPS)	A common performance measurement used to benchmark computer storage devices such as hard disk drives (HDD), solid state drives (SSD), and storage area networks (SAN)
Infrastructure	Virtual Machine (VM)	An emulation of a particular computer system. Operates based on the computer architecture and functions of a real or hypothetical computer, and its implementations may involve specialized hardware, software, or a combination of both.
Technical Architecture	Central Processing Unit (CPU)	The electronic circuitry within a computer that carries out the instructions of a computer program by performing the basic arithmetic, logical, control and input/output (I/O) operations specified by the instructions.
Technical Architecture	Horizontal scaling	Ability of an application to function across multiple instances or nodes.
Technical Architecture	Vertical scaling	Ability of an application to take advantage of additional computing power, when added (i.e. CPU, memory).
Technical Architecture	Hypervisor	A piece of computer software, firmware or hardware that creates and runs virtual machines. Sometimes called a virtual machine monitor (VMM).
Technical Architecture	Loose Coupling	Refers to designing a system in which each of its components has, or makes use of, little or no knowledge of the definitions of other separate components.
Technical Architecture	Random Access Memory (RAM)	A form of computer data storage. Stored information is lost if power is removed (computer is shut down).

References

- Cannon, N. (2014). *Key Skills Needed for Successful Deployment of Cloud Computing in Government*. Stamford: Gartner.
- Microsoft. (2011). *Business Agility and the Cloud*.
- Neville Cannon, G. A. (2015). *Government CIOs See Expected Cloud Cost Savings Evaporate*. Stamford: Gartner.
- Pham, T. (2011, September 15). *Benefits of Private Cloud Computing: Compliant & Cost-Effective*. Retrieved from Online Tech: <http://resource.onlinetech.com/benefits-of-private-cloud-computing-compliant-cost-effective/>
- Rodier, M. (2011, May 18). *Speed-to-Market Is Biggest Benefit Of Cloud Computing*. Retrieved from InformationWeek WallStreet & Technology: <http://www.wallstreetandtech.com/infrastructure/speed-to-market-is-biggest-benefit-of-cloud-computing/d/d-id/1264839>
- Savvas, A. (2014, May 14). *The benefits of hybrid cloud computing*. Retrieved from ITProPortal: <http://www.itproportal.com/2014/05/14/the-benefits-of-hybrid-cloud-computing/>
- U.S. Department of Commerce. (2011, September). *The NIST Definition of Cloud Computing*. Retrieved from National Institute of Standards and Technology: <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>