

Student: Dan Tataru
Student ID: 2125053
Course: Cloud Computing Architecture
Module code: ACCA5025
Lecturer's name: Koushik Modak
University: UWTDS
Submission deadline: 06/04/2023

Assessment 1

Table of contents

1.Introduction.....	2
2.Description, Analysis and Evaluation.....	2
2.1. Evaluation of Cloud benefits.....	2
2.2. Evaluation of Availability.....	3
2.3. Evaluation of Security.....	6
Conclusion.....	7
References.....	8

Introduction

This report explores the use of Infrastructure as a Service (IaaS) in two scenarios with differing budgets and high availability requirements. It provides an overview of virtualization and cloud service models, comparing their benefits and drawbacks. The report also examines the strategies for managing resources in Microsoft Azure and Amazon Web Services (AWS), offering insights and recommendations for each scenario. By understanding these concepts, organizations can make informed decisions when migrating to IaaS and optimizing their resources based on budget constraints and availability needs.

Description, Analysis and Evaluation

Virtualization is the process of converting the physical IT resources to a virtual representation of them. Most type of this resources can be virtualised including: servers, storage, network, power etc. Some of the benefits of virtualization might include the cost reduction because of the sharing of resources, the isolation of the VM's , independence of underlying hardware and the portability it might provide.

Virtual machines (VMs) play a critical role in cloud computing, as they allow multiple instances of operating systems and applications to operate on a single physical server. Essentially, a virtual machine is a software-based replica of a physical computer, equipped with its own processing power, memory, storage, and networking features.

Evaluation of Cloud benefits:

In the context of cloud computing, virtual machines promote efficient resource utilization and flexibility in the distribution and management of computing resources. Cloud service providers like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) employ virtualization technologies to create VMs on their vast infrastructure and provide them as a service to clients.

Infrastructure as a Service (IaaS) supplies virtual computing resources over the internet, comprising virtual machines, storage, and networking. Advantages include cost savings resulting from reduced capital expenditures, increased scalability, and flexible resource administration. Drawbacks consist of potential security risks, the need for in-house expertise to manage the infrastructure, and the risk of vendor lock-in.

Platform as a Service (PaaS) delivers a platform for developing and deploying applications, with the provider managing the underlying infrastructure. Benefits involve reduced development times, simplified application deployment, and the capacity to focus on application code rather than infrastructure management. Downsides include limited control over the infrastructure, possible vendor lock-in, and potential compatibility challenges with existing applications.

Software as a Service (SaaS) presents software applications via the internet, with the provider responsible for hosting, maintaining, and updating the software. Benefits involve ease of use, lowered maintenance costs, and access to the latest software functionalities. Drawbacks encompass limited customization options, possible data security and privacy concerns, and reliance on the provider's uptime and performance.

When comparing IaaS, PaaS, and SaaS, each service type presents unique pros and cons when contemplating the migration of an on-premises application to the cloud. IaaS grants more control over the infrastructure but demands greater management expertise. PaaS eases application development and deployment but may restrict customization options. SaaS supplies a ready-to-use solution but may raise data security and privacy concerns. The optimal choice for a small or medium-sized enterprise operating an eCommerce website and related web service would depend on their specific needs, internal expertise, and desired degree of control over the infrastructure and application.

Evaluation of Availability:

In recent years, cloud computing has become a popular choice for businesses and organizations looking to streamline their operations and leverage the numerous benefits of virtualized resources. Infrastructure as a Service (IaaS) is one such model that allows users to access and manage virtual machines (VMs), storage, and networking resources over the internet. Major IaaS providers include Microsoft Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP).

This report will examine two distinct scenarios in which a company has decided to move to an IaaS model and is working on two large projects with multiple virtual machines requiring high availability. High availability means that there should be minimal or no data loss and services must remain accessible 24/7. The two scenarios involve different budget constraints:

1. A very high budget with no cost limit
2. A medium budget with tolerable service loss under certain circumstances

Both Azure and AWS offer various solutions for managing resources in different ways. This report will discuss general concepts related to these providers, compare their benefits and drawbacks, and suggest alternative solutions for maximizing efficiency in both scenarios.

To determine the best approach for managing resources in each scenario, it's essential to understand the general concepts related to locations, availability zones, and resource groups in both Azure and AWS.

1. **Locations (Azure) / Regions (AWS):** These are geographically distinct data centers that host cloud services. Deploying resources across multiple locations or regions can improve fault tolerance and reduce latency for users in different geographic areas.
2. **Availability Zones:** These are physically separate areas within a location or region, each containing one or more data centers with independent power, cooling, and networking infrastructure. Distributing resources across multiple availability zones can ensure high availability and fault tolerance against data center-level failures.
3. **Resource Groups (Azure) / Resource Tags (AWS):** These are logical containers for resources deployed within a cloud subscription. They allow you to organize, manage, and monitor resources based on their lifecycle and interdependencies, which can help manage and visualize resources based on their relationships, simplifying management and monitoring.

Scenario 1: Very High Budget with No Cost Limit

In this scenario, the organization can deploy resources across multiple availability zones within a location (or region) and across multiple locations (or regions) to achieve maximum availability and prevent data loss. This approach can protect against both data center-level and location-level failures, ensuring 24/7 service availability. Additionally, using services such as load balancers and traffic managers to distribute traffic across instances can further enhance availability and performance.

With an unlimited budget, organizations can take advantage of premium services offered by Azure, AWS, and other cloud providers for disaster recovery planning, dedicated connectivity, and high-performance storage and databases. Moreover, organizations can consider a multi-cloud strategy, using a combination of Azure, AWS, and other cloud providers to optimize availability, redundancy, and performance. This approach may involve using container orchestration platforms like Kubernetes to manage resources across multiple cloud providers seamlessly. Adopting a multi-cloud strategy can help organizations avoid vendor lock-in and leverage the unique strengths of each cloud provider.

Scenario 2: Medium Budget with Tolerable Service Loss

For organizations with a more limited budget, resources can be allocated across multiple availability zones within a single location (or region) to provide fault tolerance against data center-level failures. While this method is more cost-effective than deploying resources across multiple locations (or regions), it may not guard against location-level failures, which could result in occasional service loss under specific conditions.

Organizations can use various cost-saving features provided by Azure, AWS, and other cloud providers, such as reserved instances and savings plans for virtual machines, which offer significant discounts compared to on-demand pricing. Additionally, organizations can use autoscaling and spot instances to optimize resource usage and reduce costs. Autoscaling allows for the automatic adjustment of computing resources based on demand, while spot instances enable users to bid on unused computing capacity at a lower price than regular instances. To further minimize costs, organizations can consider using open-source solutions and tools that are compatible with Azure, AWS, and other cloud providers. These solutions can provide similar functionality at a lower price point or even for free.

Another cost-saving approach is to implement a hybrid cloud strategy, combining the use of public cloud resources with on-premises infrastructure. This approach allows organizations to maintain some data and applications on their own servers while leveraging the benefits of cloud computing for other aspects of their operations. By carefully selecting which resources to move to the cloud and which to keep on-premises, organizations can strike a balance between cost-efficiency and high availability.

Additionally, organizations should prioritize regular monitoring, auditing, and optimization of their cloud resources to identify and address inefficiencies, underutilized resources, or potential bottlenecks. Tools such as Azure Cost Management, AWS Cost Explorer, and third-party cost management solutions can help organizations gain insights into their cloud spending and optimize their resource allocation.

Lastly, organizations should ensure that they have a comprehensive backup and disaster recovery plan in place to minimize the impact of potential service disruptions. This plan should include regular backups of critical data, periodic testing of backup and recovery processes, and a well-defined incident response process.

In conclusion, both Azure and AWS offer various solutions for managing resources with varying levels of availability and budget constraints. By understanding the general concepts related to locations, availability zones, and resource groups, organizations can select the appropriate deployment strategy to deliver the desired level of service based on their budget and availability needs. In the scenario with a very high budget and no cost limit, organizations can deploy resources across multiple availability zones and locations, leveraging premium services and a multi-cloud strategy for maximum availability and performance. On the other hand, in the medium budget scenario with tolerable service loss, organizations can allocate resources across multiple availability zones within a single location, implement cost-saving measures, and adopt a hybrid cloud strategy to balance cost-efficiency and high availability. Regular monitoring, optimization, and a comprehensive backup and disaster recovery plan are essential for ensuring the success of both scenarios.

Evaluation of Security:

Authorization and authentication are two essential components of a secure system. While they are related, they serve distinct purposes in the context of cybersecurity.

3. Authentication: Authentication is the process of verifying the identity of a user, device, or system attempting to access a resource. This is typically achieved through the use of credentials such as usernames and passwords, tokens, or biometric data. Authentication ensures that only valid and recognized entities can access the protected resources.

4. Authorization: Authorization is the process of granting or denying access to specific resources based on the authenticated user's or system's permissions. Once an entity has been authenticated, the system checks their access rights and privileges to determine what actions they can perform or resources they can access. Authorization helps maintain proper control over sensitive data and resources within an organization.

The importance of the following security measures is as follows:

4. Multi-factor Authentication (MFA): MFA enhances security by requiring users to provide multiple independent credentials for authentication. This typically includes something the user knows (e.g., password), something the user has (e.g., a hardware token or mobile device), and/or something the user is (e.g., biometric data like fingerprint or facial recognition). MFA significantly reduces the risk of unauthorized access, as it is less likely that an attacker can compromise multiple independent factors.

5. Least-privilege permissions: Applying the least-privilege principle means granting users and systems the minimum level of access and permissions necessary to perform their tasks. This approach reduces the potential impact of

security breaches, as compromised accounts would have limited access to sensitive resources and data. It also helps to prevent accidental or deliberate misuse of privileges by authorized users.

6. Network Security Groups: Network Security Groups (NSGs) are used to control inbound and outbound traffic to resources within a virtual network. They act as virtual firewalls, allowing or denying traffic based on rules and policies. NSGs help to secure resources by limiting access only to trusted sources and preventing unauthorized access or data exfiltration.

7. Bastions (Jumpboxes): A Bastion host, or Jumpbox, is a specially configured and hardened server that acts as an intermediary between a user's local machine and the target resources in a secure environment, such as a virtual network. It provides a single, controlled point of entry for managing and accessing resources, reducing the attack surface and mitigating the risk of unauthorized access.

Given that cybersecurity is always a concern for public applications, it is crucial to implement these security measures when deploying a public service to the cloud. These practices help to protect sensitive data and resources while minimizing the risk of unauthorized access, data breaches, and other security incidents.

Conclusion

In conclusion, the adoption of cloud computing and Infrastructure as a Service (IaaS) can provide organizations with numerous benefits, such as increased flexibility, scalability, and cost savings. However, ensuring high availability and robust security are crucial aspects to consider when migrating to the cloud, particularly for large projects with multiple virtual machines. The two scenarios discussed in this report highlight the importance of understanding the general concepts of locations, availability zones, and resource groups when managing resources in Azure, AWS, or other cloud providers. In the high-budget scenario, organizations can focus on maximizing availability and performance by deploying resources across multiple availability zones and locations, and adopting a multi-cloud strategy. In contrast, the medium-budget scenario requires organizations to balance cost-efficiency and high availability, using various cost-saving measures and a hybrid cloud approach. Implementing strong security measures, such as multi-factor authentication, least-privilege permissions, network security groups, and bastion hosts, is crucial in protecting sensitive data and resources from unauthorized access and security breaches. These security practices can significantly reduce risks and ensure a more secure cloud environment for public services. Based on the findings in this report, we can recommend that organizations carefully assess their budget and availability requirements before choosing a cloud provider and deployment strategy.

Additionally, organizations should prioritize implementing robust security measures and continuously monitor and optimize their cloud resources to achieve the desired balance between cost-efficiency, availability, and security.

References:

- [1] M. Armbrust et al., "A view of cloud computing," Commun. ACM, vol. 53, no. 4, pp. 50-58, Apr. 2010. <https://dl.acm.org/doi/10.1145/1721654.1721672>
- [2] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Gener. Comput. Syst., vol. 25, no. 6, pp. 599-616, Jun. 2009.
<https://www.sciencedirect.com/science/article/abs/pii/S0167739X08001957?via%3Dihub>
- [3] J. Rhoton, Cloud Computing Explained: Implementation Handbook for Enterprises. Recursive Press, 2017.
- [4] T. Erl, Z. Mahmood, and R. Puttini, Cloud Computing: Concepts, Technology & Architecture. Prentice Hall, 2013.
- [5] Microsoft Azure, "What is high availability?" <https://docs.microsoft.com/en-us/azure/architecture/guide/design-principles/high-availability>.
- [6] Amazon Web Services, "AWS Well-Architected Framework." <https://aws.amazon.com/architecture/well-architected/>.
- [7] Google Cloud Platform, "Google Cloud documentation: Overview of solutions and best practices." <https://cloud.google.com/docs/>.