

Microsoft Certified: Azure Fundamentals (AZ-900) Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

SAMPLE

- 1. Which Azure service allows you to build and host web apps, mobile back ends, and RESTful APIs without managing infrastructure?**
 - A. Web Apps**
 - B. Azure App Service**
 - C. JSON**
 - D. Azure Functions**
- 2. A resource group can contain resources from the same region as the resource group only. True or False?**
 - A. True**
 - B. False**
 - C. Only if specified by the user**
 - D. Not applicable**
- 3. A company is deploying a critical business application on two virtual machines (VMs). Which configuration should the company use to deploy the solution for high availability and low latency?**
 - A. Separate regions in a regional pair.**
 - B. Separate availability sets.**
 - C. Separate availability zones.**
 - D. Separate resource groups in the same region.**
- 4. You want to deploy your own web application quickly without taking any responsibility for the Operating System, Web server, or system updates. What is the right solution for you?**
 - A. DaaS**
 - B. SaaS**
 - C. PaaS**
 - D. IaaS**
- 5. What does the Azure Resource Manager provide?**
 - A. Access Control**
 - B. Resource Provisioning**
 - C. Monitoring and Diagnostics**
 - D. All of the above**

- 6. Can machine learning models from Machine Learning Studio be managed by Azure Machine Learning Service?**
- A. True**
 - B. False**
- 7. Which service is designed to monitor the performance of container workloads?**
- A. Azure Monitor for containers**
 - B. Azure Service Health**
 - C. Application Insights**
 - D. Azure Monitor for VMs**
- 8. Which of the following statements is true about Azure's scalability?**
- A. It requires manual intervention for scaling operations**
 - B. It can automatically adjust resources based on demand**
 - C. It is limited to virtual machines**
 - D. It can only scale down resources**
- 9. How does Azure Traffic Manager reduce latency?**
- A. It chooses only the fastest networks between endpoints.**
 - B. It selects the endpoint that's closest to the user's DNS server.**
 - C. It caches content, similar to how content delivery networks work.**
 - D. It prioritizes traffic based on bandwidth.**
- 10. What term describes the ability to scale applications and services in Azure?**
- A. Elasticity**
 - B. Scalability**
 - C. Capacity Management**
 - D. Performance Optimization**

Answers

SAMPLE

1. B
2. B
3. C
4. C
5. D
6. B
7. A
8. B
9. B
10. A

SAMPLE

Explanations

SAMPLE

1. Which Azure service allows you to build and host web apps, mobile back ends, and RESTful APIs without managing infrastructure?

A. Web Apps

B. Azure App Service

C. JSON

D. Azure Functions

The correct answer is Azure App Service because it is a fully managed platform designed for building, deploying, and scaling web applications and APIs. This service abstracts the underlying infrastructure, allowing developers to focus on their application code without the need to manage virtual machines, servers, or storage. Azure App Service supports multiple programming languages, integrates seamlessly with other Azure services, and provides features like auto-scaling, custom domains, and SSL certificates out of the box. While Web Apps is often used informally to refer to web applications on the Azure platform, Azure App Service is the official name of the service that encompasses not only web apps but also mobile back ends and RESTful APIs. This distinction confirms that Azure App Service is tailored for developers looking for a streamlined development experience without the burden of infrastructure management. JSON is a data format used for data interchange and is not related to hosting applications, so it does not apply to the context of the question. Azure Functions is a serverless compute service that allows you to run code in response to events; however, it is better suited for executing smaller code snippets rather than serving as a comprehensive platform for fully fledged web applications.

2. A resource group can contain resources from the same region as the resource group only. True or False?

A. True

B. False

C. Only if specified by the user

D. Not applicable

A resource group in Azure is a logical container that can hold various types of resources, such as virtual machines, storage accounts, and web apps. The key point is that a resource group can contain resources from different regions. This flexibility allows organizations to design and manage their resources across multiple geographical locations as needed. While a resource group is created in a specific region, it can contain resources deployed in other regions as well. For instance, you could have a resource group located in the East US region that includes virtual machines in both the East US and West Europe regions. This capability provides the advantage of organizing resources logically while supporting a globally distributed architecture. In summary, a resource group is not limited to containing only resources from the same region, making the statement false.

3. A company is deploying a critical business application on two virtual machines (VMs). Which configuration should the company use to deploy the solution for high availability and low latency?

A. Separate regions in a regional pair.

B. Separate availability sets.

C. Separate availability zones.

D. Separate resource groups in the same region.

When deploying critical business applications that require high availability and low latency, using separate availability zones is the best practice. Availability zones are physically separated locations within a region that provide redundancy and fault isolation. By deploying the virtual machines in different availability zones, the company ensures that if one zone experiences a failure, the other zone remains operational, thereby enhancing the application's availability. In addition, the proximity of availability zones within the same region optimizes latency, ensuring that the latency between VMs remains low while still providing resilience. This is crucial for applications that require real-time processing and minimal downtime. Using separate availability sets would offer some redundancy but would not provide the same level of physical separation as availability zones, making them less effective for protecting against certain types of failures. Deploying in separate regions would increase latency due to the distance between the regions and could complicate the application architecture. Lastly, using separate resource groups in the same region does not provide any fault tolerance benefits, as it lacks physical separation.

4. You want to deploy your own web application quickly without taking any responsibility for the Operating System, Web server, or system updates. What is the right solution for you?

A. DaaS

B. SaaS

C. PaaS

D. IaaS

The chosen answer is correct because Platform as a Service (PaaS) provides a cloud-based platform that enables developers to build, deploy, and manage applications without worrying about the underlying infrastructure, including the operating system and web server. PaaS abstracts these responsibilities, allowing you to focus on coding and deploying your application quickly. This model is particularly beneficial for scenarios where speed and efficiency are paramount, as it handles routine maintenance tasks like system updates and patches, which can greatly reduce the administrative burden on developers. In contrast, other options like DaaS (Desktop as a Service) and SaaS (Software as a Service) do not fit the scenario perfectly. DaaS focuses on providing virtual desktops, while SaaS delivers software applications over the internet without user control over the underlying infrastructure. IaaS (Infrastructure as a Service), meanwhile, requires users to manage the operating system and the underlying hardware, which is not aligned with your requirement to avoid responsibility for those components. Therefore, PaaS is indeed the most appropriate choice for quick deployment while minimizing operational responsibilities.

5. What does the Azure Resource Manager provide?

- A. Access Control
- B. Resource Provisioning
- C. Monitoring and Diagnostics
- D. All of the above**

The Azure Resource Manager (ARM) serves as a crucial management layer in Microsoft Azure that provides a wide range of capabilities for handling resources in a consistent and organized manner. One of its primary functions is to facilitate access control, allowing users to assign roles and permissions to manage who can take action on resources. By using role-based access control (RBAC), users can ensure that only authorized individuals can perform specific actions on Azure resources. In addition to access control, Azure Resource Manager plays a vital role in resource provisioning. This means that it allows users to create, update, and delete resources in Azure through templates, ensuring that resources are deployed in a standardized manner. This capability is particularly important for establishing infrastructure as code, enabling automated and repeatable deployments. Furthermore, ARM also supports monitoring and diagnostics of resources. It provides integrated monitoring capabilities, allowing users to track the performance and health of resources. This helps organizations ensure that their applications run smoothly and enables proactive troubleshooting. Given these functionalities—access control, resource provisioning, and monitoring and diagnostics—selecting all of the above is appropriate because the Azure Resource Manager encompasses these three critical aspects of resource management within Azure, making it an indispensable component of the Azure ecosystem.

6. Can machine learning models from Machine Learning Studio be managed by Azure Machine Learning Service?

- A. True
- B. False**

The correct answer is based on the relationship between Azure Machine Learning Studio and Azure Machine Learning Service. Azure Machine Learning Studio is an integrated development environment for building, training, and deploying machine learning models. However, it operates independently of Azure Machine Learning Service, which provides comprehensive capabilities for managing the entire machine learning lifecycle, including advanced experimentation, model management, and deployment functionalities. While you can develop models in Azure Machine Learning Studio, the specific management capabilities provided by Azure Machine Learning Service, such as version control, tracking, and fine-tuning deployed models, are not directly available for those models created in the studio. Therefore, machine learning models created and trained in Azure Machine Learning Studio cannot be managed effectively using the Azure Machine Learning Service. This distinction helps to clarify that while both tools serve related purposes within the Azure ecosystem, they function in separate contexts regarding model management.

7. Which service is designed to monitor the performance of container workloads?

A. Azure Monitor for containers

B. Azure Service Health

C. Application Insights

D. Azure Monitor for VMs

Azure Monitor for containers is designed specifically to monitor the performance and health of container workloads. It provides visibility into the performance of your containerized applications deployed in Azure Kubernetes Service (AKS) and other container orchestration platforms. By using this service, you can collect and analyze metrics and logs from your containers, allowing you to understand resource utilization, diagnose issues, and ensure optimal performance. This service is tailored to handle the complexities associated with containers, making it an essential tool for developers and IT professionals managing container-based environments. It offers capabilities such as live data, detailed insights, and integration with other Azure services for a comprehensive monitoring solution. The other choices serve different purposes: Azure Service Health focuses on the overall health of Azure services and not on specific workloads, Application Insights is more geared toward monitoring application performance and user interaction in a broader context rather than containers specifically, and Azure Monitor for VMs is dedicated to monitoring virtual machines rather than containers. Therefore, Azure Monitor for containers is the most fitting choice for monitoring container workloads.

8. Which of the following statements is true about Azure's scalability?

A. It requires manual intervention for scaling operations

B. It can automatically adjust resources based on demand

C. It is limited to virtual machines

D. It can only scale down resources

The statement that Azure can automatically adjust resources based on demand is correct. This feature of Azure is known as "auto-scaling," which allows Azure services to dynamically scale up or down based on the workload requirements. This capability is particularly beneficial in managing fluctuating workloads without requiring manual intervention, ensuring that applications remain efficient and cost-effective. Azure's scalability includes not just virtual machines, but encompasses a wide range of services like Azure App Service, Azure Kubernetes Service, and databases. This means that scaling can apply across various components and not just to a single type of resource. The ability to automatically manage resources helps businesses respond quickly to changing demands, enhancing performance while minimizing waste. This flexibility is a significant advantage of using cloud services like Azure.

9. How does Azure Traffic Manager reduce latency?

- A. It chooses only the fastest networks between endpoints.
- B. It selects the endpoint that's closest to the user's DNS server.**
- C. It caches content, similar to how content delivery networks work.
- D. It prioritizes traffic based on bandwidth.

Azure Traffic Manager reduces latency primarily by selecting the endpoint that is closest to the user's DNS server. When a user makes a request, Traffic Manager utilizes the DNS resolution process to determine which of the available endpoints can respond the quickest based on geographic proximity. The principle behind this is that shorter distances typically result in lower latency, helping to ensure that users connect to the most responsive service region. This approach optimally directs the traffic to the location that can fulfill the request faster, thereby enhancing the overall user experience. By optimizing the selection process through DNS queries, Traffic Manager effectively minimizes the time it takes for users to access resources. Other strategies such as choosing the fastest networks or caching content have differing functions and are more aligned with other Azure services like Azure CDN, which specifically focuses on content delivery. Prioritizing traffic based on bandwidth is another different management technique that doesn't directly relate to reducing latency through intelligent endpoint selection in the context of Traffic Manager.

10. What term describes the ability to scale applications and services in Azure?

- A. Elasticity**
- B. Scalability
- C. Capacity Management
- D. Performance Optimization

The correct term that describes the ability to scale applications and services in Azure is "Elasticity." Elasticity refers to the dynamic allocation of resources as workload demands change, allowing applications to automatically scale up or down in response to varying levels of traffic or usage. This capability enables organizations to effectively manage costs and maintain performance without manual intervention. While "Scalability" is often used in similar contexts, it generally refers to the capability of a system to handle a growing amount of work or its potential to accommodate growth. In contrast, elasticity emphasizes the automatic and real-time response to current resource demands, making it a more precise term in the context of cloud computing. "Capacity Management" relates to planning and managing capacity resources to meet future demands, while "Performance Optimization" focuses on enhancing the performance of applications. Neither of these terms encompasses the dynamic scaling feature that elasticity provides in Azure environments.