JON BONSO, GEROME PAGATPATAN AND
ACE BATACANDULO

AZURE
CERTIFIED

AZ-104 Microsoft Azure Administrator

Tutorials Dojo Study Guide

TABLE OF CONTENTS

INTRODUCTION

With the rapid advancement of technology, enterprises are adopting newer technologies that will help their businesses transform and grow. Microsoft Azure is one of the emerging technologies that you can leverage in this age since a lot of companies are shifting their existing infrastructures in the cloud. Unlike the traditional setup, cloud computing allows you to obtain resources on-demand with just one click on their platform, including the servers, storage, databases, networking, analytics, artificial intelligence, and a lot more.

Microsoft Azure offers a range of cloud services, depending on your business needs. These services are continuously upgrading, and new features are being added every year to deliver customer satisfaction. Since Azure's resources and services are too vast, the **Microsoft Azure Certification** program offers different certification paths that will help aspiring candidates and IT professionals validate their skills and knowledge to maximize the solutions created in the cloud.

Microsoft Azure is the second biggest cloud service provider in the market next to AWS, and a lot of companies are now adopting a **multicloud** strategy, which makes it all the more beneficial for IT professionals like you to expand your skill set and learn multiple cloud technologies. Learning is a lot more fun if you merge it with various cloud services. It will be an exciting and enjoyable journey for you, and the first step is to become **AZ-104 Microsoft Azure Administrator** certified. This eBook will help familiarize yourself with the basic cloud concepts as well as the core services of Microsoft Azure, which are the building blocks that will help you pass the exam and make a successful career shift to cloud computing.

Note: We took extra care to come up with these study guides and cheat sheets, however, this is meant to be just a supplementary resource when preparing for the exam. We highly recommend working on **hands-on sessions** and **practice exams** to further expand your knowledge and improve your test-taking skills.

AZ-104 MICROSOFT AZURE ADMINISTRATOR EXAM OVERVIEW

The Microsoft Azure Certification Program validates the technical skills and knowledge for building secure and reliable cloud-based applications using the Azure platform. By successfully passing the Microsoft Azure exam, individuals can prove their expertise to their current and future employers.

Exam Details

The AZ-104 Microsoft Azure Administrator examination is intended for IT Professionals who implement, manage and monitor an organization's cloud infrastructure. You can take this exam from a local testing center or online from the comfort of your home. The exam is composed of different types of questions.

For multiple-choice types of questions, you will have to choose one correct response out of four options.

For Drag and Drop questions, match the items by dragging them to their correct descriptions.

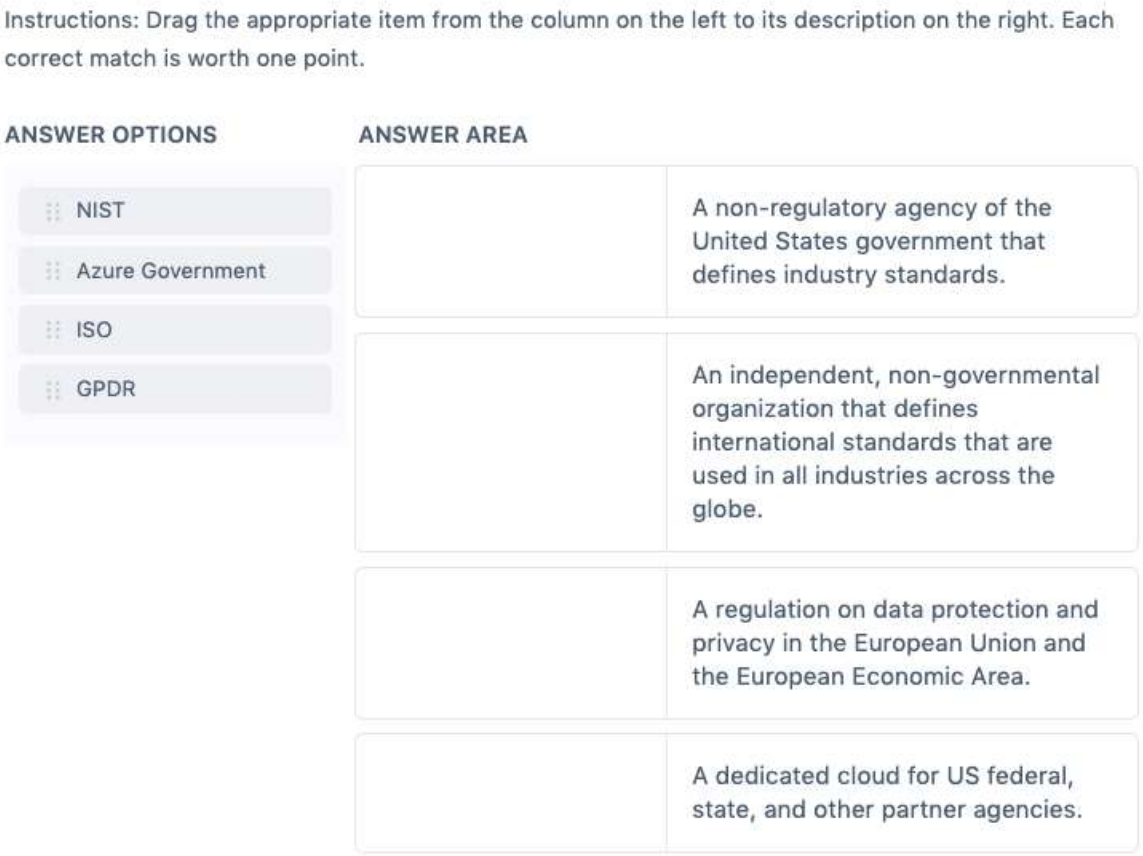

For Dropdown types of questions, select the correct answer from the drop-down list of options.

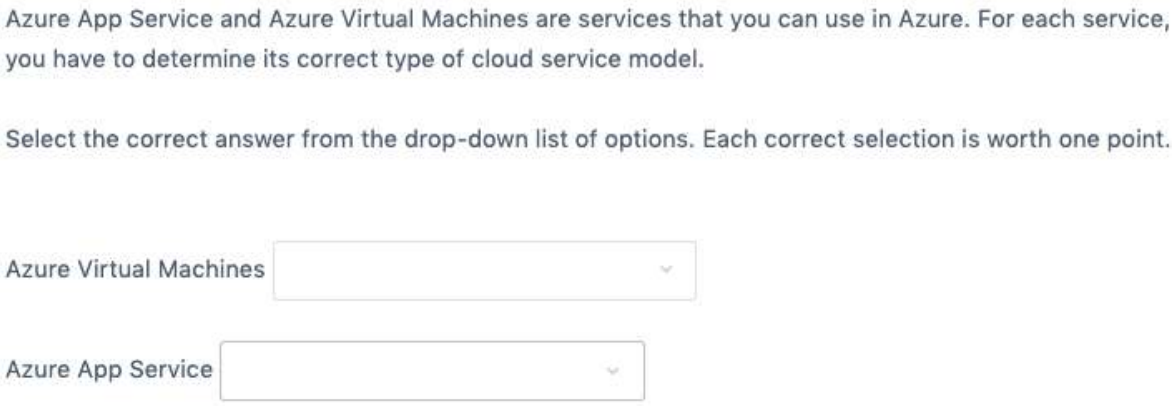

For Hotspot types of questions such as multiple Yes/No, evaluate whether the presented statements relating to a certain topic are correct/incorrect.

For each of the following items, choose **Yes** if the statement is true or choose **No** if the statement is false. Take note that each correct item is worth one point.

Questions	Yes	No
Azure virtual machines are billed on a per-hour basis.	○	○
When you delete a virtual machine in Azure, by default, any disks that are attached to the VM are deleted.	○	○
Disks attached to stopped virtual machines do not incur costs.	○	○

You can take the exam via online proctoring or from a testing center close to you.

Exam Code:	AZ-104
Prerequisites:	None
No. of Questions:	50-60
Score Range:	100-100
Cost:	165 USD
Passing Score:	700
Time Limit:	100 - 120 minutes

Exam Domains

The AZ-104 Microsoft Azure Administrator exam evaluates your skills across several domains, with the following updated weightage and topic coverage: Manage Azure identities and governance (20–25%), Implement and manage storage (15–20%), Deploy and manage Azure compute resources (20–25%), Implement and manage virtual networking (15–20%), and Monitor and maintain Azure resources (10–15%). These domains focus on your ability to manage and configure Azure services, including identity and governance tasks, storage solutions, compute resources, networking, and monitoring. The exam tests your proficiency in managing Azure environments and ensuring they are optimized for performance, security, and availability.

Manage Azure identities and governance

- Manage Microsoft Entra users and groups
- Manage access to Azure resources
- Manage Azure subscriptions and governance

Implement and manage storage

- Configure access to storage
- Configure and manage storage accounts
- Configure Azure Files and Azure Blob Storage

Deploy and manage Azure compute resources

- Automate deployment of resources by using Azure Resource Manager (ARM) templates or Bicep files
- Create and configure virtual machines
- Provision and manage containers in the Azure portal
- Create and configure an Azure App Service

Implement and manage virtual networking

- Configure and manage virtual networks in Azure
- Configure secure access to virtual networks
- Configure name resolution and load balancing

Monitor and backup Azure resources

- Monitor resources in Azure Monitor
- Implement backup and recovery

Exam Scoring System

You can get a score from 100 to 1,000 with a minimum passing score of 700 when you take the AZ-104 Microsoft Azure Administrator exam. Microsoft uses a scaled scoring model to associate scores across multiple exam types that may have different levels of difficulty. Your complete score report will be sent to you by email 1 - 5 business days after your exam. However, as soon as you finish your exam, you'll immediately see a pass or fail notification on the testing screen.

For individuals who unfortunately do not pass their exams, you must wait 24 hours before you are allowed to retake the exam. There is no hard limit on the number of attempts you can retake an exam.

Once you receive your score report via email, the result should also be saved in your Microsoft Certification account. The score report contains a table of your performance in each domain and it will indicate whether you have met the level of competency required for these. Take note that you do not need to achieve competency in all areas for you to pass the exam. In the first part of the report, there will

be a performance summary by exam section that highlights your strengths and weaknesses, which can help you determine the areas you need to improve on.

Exam Benefit

If you successfully pass any Microsoft Certification exam, you will receive a **Certified Digital Badge**. You can showcase your achievements to your colleagues and employers by adding these digital badges to your email signatures, Linkedin profile, or on your social media accounts. To view your badges, simply go to the "Dashboard" section of your Acclaim Account.

You can visit the official Microsoft Certification FAQ page to view the frequently asked questions about getting certified and other information about the Microsoft Certification: https://docs.microsoft.com/en-us/learn/certifications/certification-exam-policies.

AZ-104 MICROSOFT AZURE ADMINISTRATOR EXAM - STUDY GUIDE AND TIPS

The AZ-104 Microsoft Azure Administrator certification exam is designed for people who have experience in implementing, managing, and monitoring a Microsoft Azure environment. The exam will test your technical skills in implementing solutions based on different scenarios. Having prior experience in infrastructure management will help you understand the concepts and services easily.

The content of the exam will test your ability to perform the following:

- Manage Azure identities and governance
- Implement and manage storage
- Deploy and manage Azure compute resources
- Configure and manage virtual networking
- Monitor and backup Azure resources

For more information about the AZ-104 exam, you can check out this exam skills outline. This study guide will provide you with comprehensive review materials to help you pass the exam with flying colors.

Study Materials

For the Microsoft Azure Administrator exam, we recommend that you check out these study materials first before you take the actual exam. These resources will help you understand complex concepts and services that will be useful on your exam day.

1. **Microsoft Learn** – this website provides different learning paths for various Microsoft certifications. For the AZ-104 certification exam, you can focus on the following modules:
    - Prerequisites for Azure administrators
    - Manage identities and governance in Azure
    - Implement and manage storage in Azure
    - Deploy and manage Azure compute resources
    - Configure and manage virtual networks for Azure administrators

2. **Azure Documentation** – these documentations contain an overview, tutorials, examples, and how-to guides that will help broaden your knowledge on different Azure services.

3. **Azure Blog** – to get updated on new technologies and offerings of Microsoft Azure, you can subscribe to their newsletter.

4. **Azure FAQs** – you can find the FAQs section on the Azure documentation. The FAQs section is a compiled list of commonly asked questions, use cases, and a comparison of several Azure services.

5. **Azure free account** – the Azure portal will help you get hands-on experience with its 12-month trial. You'll also get free credits that you can spend for the first 30 days.

6. **Tutorials Dojo's Azure Cheat Sheets** – with the help of our cheat sheets, you can easily understand the information found in the Azure documentation. These are presented in bullet point format to highlight the important concepts.

7. **Tutorials Dojo's AZ-104 Microsoft Azure Administrator Practice Exams** – our practice exams have always been regarded as the best in the market. Each question in our practice tests contains detailed explanations at the end of each set to help you digest important concepts that will help you pass your Microsoft Azure certification exam on your first try.

Azure Services to Focus On

Your primary source of information when studying for the AZ-104 certification exam is the Azure documentation. To comprehend the different scenarios in the exam, you should have a thorough understanding of the following services:

1. **Azure Virtual Network** – you should know how to create a VNet peering, security rules, configuration of private/public IP addresses, network interface, subnets, and virtual networks.

2. **Azure DNS** – the configuration of custom DNS, private, and public DNS zone.

3. **Azure Application Gateway** – you should know when to use a load balancer and a web traffic load balancer, and how to create a web application firewall.

4. **Azure Load Balancer** – the types of load balancing rules, the difference between a public load balancer, and an internal load balancer.

5. **Azure VPN Gateway** – know how to configure VPN and VPN gateway.

6. **Azure ExpressRoute** – understand the concepts of ExpressRoute and how you would implement it in your environment.

7. **Azure Virtual Machines** – learn how to deploy and configure a VM, scale sets, highly available solutions, moving and redeploying of VM, creating a backup, backup policy, and recovery services vaults.

8. **Azure App Service** – learn how to create an app service plan and what run time can be put in the same app service plan.

9. **Azure Container Instances** – understand the concepts of containers and how to use ACI.

10. **Azure Blob** – you need to learn how to configure storage accounts, import/export of data, storage tiers, replication, and authentication.

11. **Azure Files** – learn how to create a file share, file sync, copy data using AZCopy.

12. **Microsoft Entra ID** – you should know how to manage a user, group, guest accounts, joined devices, device settings, and best practices.

13. **Azure RBAC** – learn how to create and assign a role and the types of built-in roles.

14. **Azure Policy** – you need to learn how to read and create a policy.

15. **Azure Monitor** – you should know how to interpret metrics, the configuration of monitor agent, query and analyze logs, set up alerts and actions, and other service features.

We suggest that you check out **Tutorials Dojo's Azure Cheat Sheets**, which provide bullet-point summaries of the most important concepts on different Azure services.

Validate Your Knowledge

If you're feeling confident because you've followed the recommended materials above, it's time to test your knowledge of various Azure concepts and services. For high-quality practice exams, you can use the Tutorials Dojo **AZ-104 Microsoft Azure Administrator Associate practice exams**.

These practice tests cover the relevant topics that you can expect from the real exam. It also contains different types of questions such as single choice, multiple responses, hotspot, yes/no, drag and drop, and case studies. Every question on these practice exams has a detailed explanation and adequate reference links that help you understand why the correct answer is the most suitable solution. After you've taken the exams, it will highlight the areas that you need to improve on. Together with our cheat sheets, we're confident that you'll be able to pass the exam and have a deeper understanding of how Azure works.

Sample Practice Test Questions:

Question 1

Your company has an Microsoft Entra tenant named tutorialsdojo.onmicrosoft.com and a public DNS zone for tutorialsdojo.com.

You added the custom domain name tutorialsdojo.com to Microsoft Entra ID. You need to verify that Azure can verify the domain name.

What DNS record type should you use?

1. SRV
2. NSEC
3. NSEC3
4. MX

Correct Answer: 4

Microsoft Entra ID is Microsoft's cloud-based identity and access management service, which helps your employees sign in and access resources in:

– External resources, such as Microsoft Office 365, the Azure portal, and thousands of other SaaS applications.

– Internal resources, such as apps on your corporate network and intranet, along with any cloud apps developed by your own organization.

Microsoft Online business services, such as Office 365 or Microsoft Azure, require Microsoft Entra ID for sign-in and to help with identity protection. If you subscribe to any Microsoft Online business service, you automatically get Microsoft Entra ID with access to all the free features.

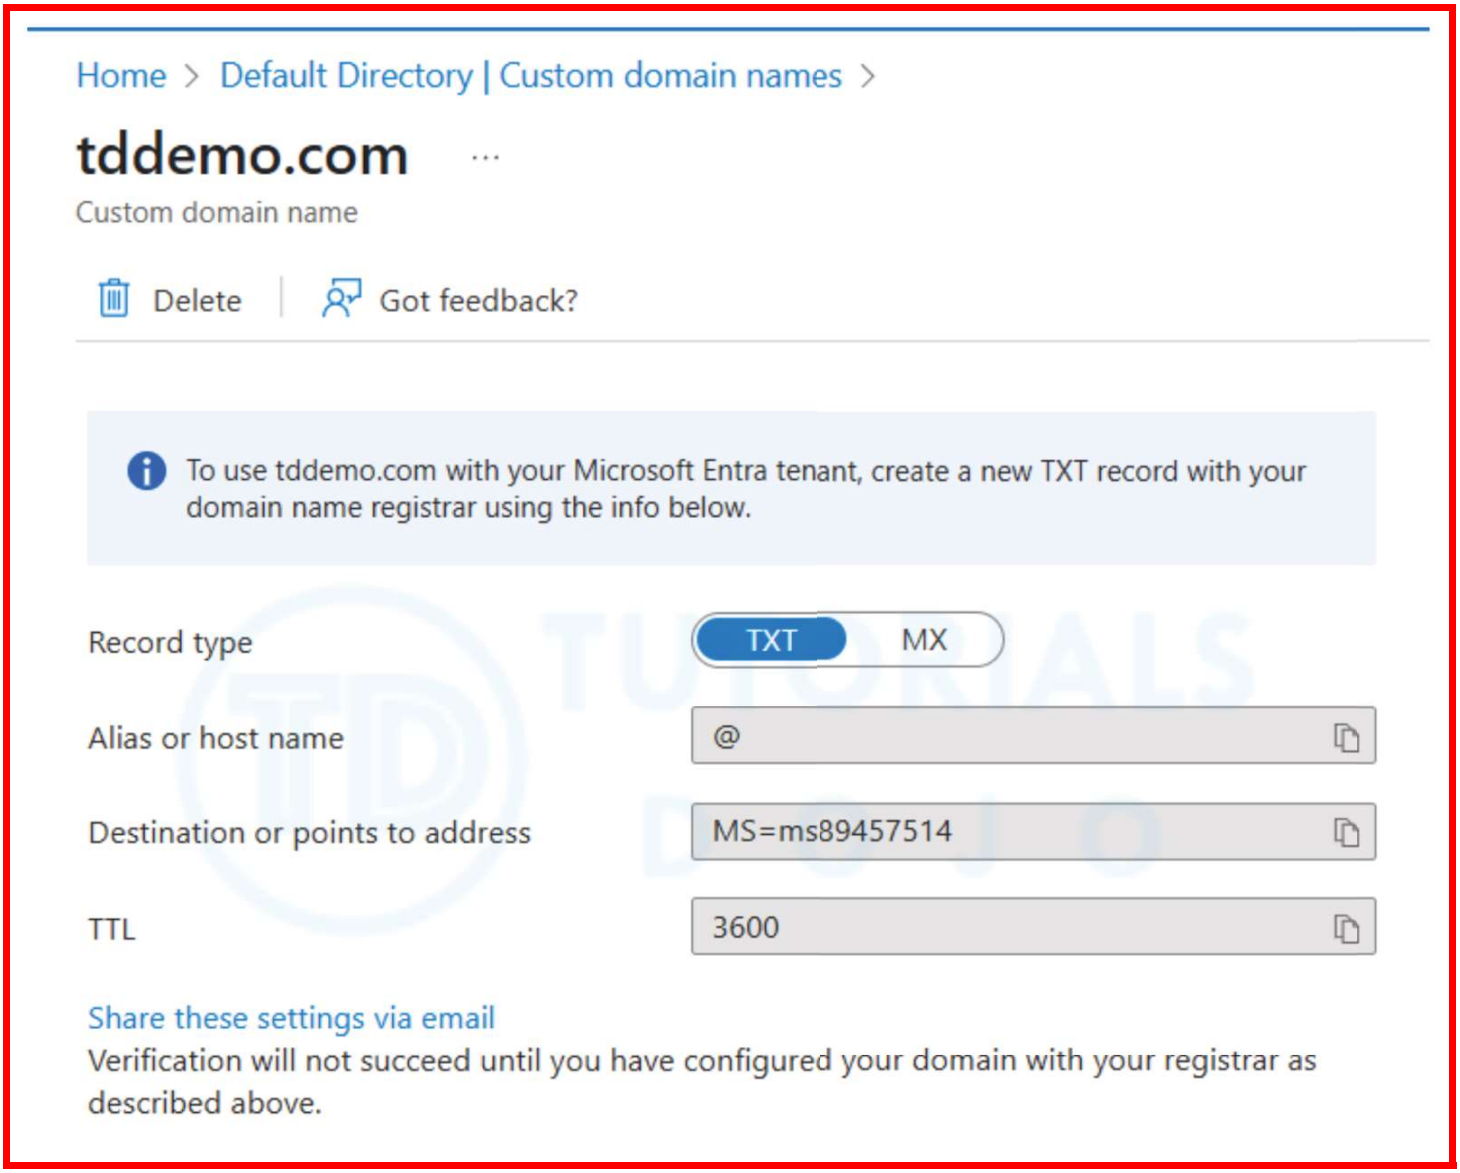

Every new Microsoft Entra ID tenant comes with an initial domain name, *<domainname>.onmicrosoft.com*. You can't change or delete the initial domain name, but you can add your organization's names. Adding custom domain names helps you to create user names that are familiar to your users, such as *azure@tutorialsdojo.com*.

You can verify your custom domain name by using TXT or MX record types.

Hence, the correct answer is: **MX.**

SRV, NSEC, and NSEC3 are incorrect because these record types are not supported by Microsoft Entra ID for verifying your custom domain. Only TXT and MX record types are supported.

References:

https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/add-custom-domain

https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/active-directory-whatis

Check out this Microsoft Entra ID Cheat Sheet:

https://tutorialsdojo.com/microsoft-entra-id/

Question 2

You plan to automate the deployment of Windows Servers using a virtual machine scale set.

You need to make sure that the web components are installed in the virtual machines.

Which two actions should you perform?

1. Create a configuration script.
2. Create an automation account.
3. Create a policy.
4. Configure the extensionProfile section of the ARM template.
5. Create a new scale set.

Correct Answer: 1, 4

Azure virtual machine scale sets let you create and manage a group of load-balanced VMs. The number of VM instances can automatically increase or decrease in response to demand or a defined schedule. Scale sets provide high availability to your applications and allow you to centrally manage, configure, and update a large number of VMs.

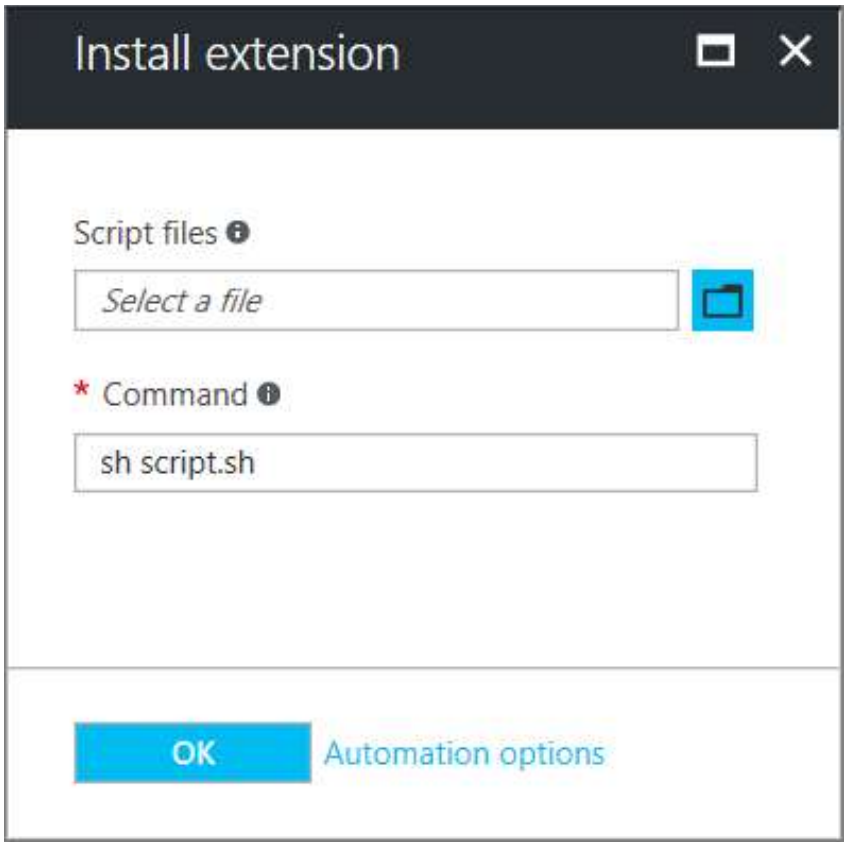

The Custom Script Extension downloads and executes scripts on Azure virtual machines. This extension is useful for post-deployment configuration, software installation, or any other configuration or management tasks.

Hence, the correct answers are:

– Create a configuration script.

– Configure the extensionProfile section of the ARM template.

The option that says: **Create an automation account** is incorrect because an automation account wouldn't help you automatically install web components. You still need to create a configuration script and extensionProfile in the ARM template.

The option that says: **Create a policy** is incorrect because this option only evaluates resources in Azure. Take note that you don't need to create a policy to install web components.

The option that says: **Create a new scale set** is incorrect because this wouldn't install the required web components. Instead of creating a new scale set, you should use a custom script extension to install the web components in the VMs.

References:

https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-deploy-app

https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/tutorial-install-apps-template#what-is-the-azure-custom-script-extension

https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-deploy-app#already-provisioned

Check out this Azure Virtual Machines Cheat Sheet:

https://tutorialsdojo.com/azure-virtual-machines/

For more **AZ-104 practice exam** questions with detailed explanations, check out the **Tutorials Dojo Portal**:

AZ-900 Microsoft Azure Fundamentals Practice Exams

AZ-104 Microsoft Azure Administrator Practice Exams

Final Remarks

It is not enough to understand the concepts at a high level. You also need to get hands-on experience by using the Microsoft Azure Portal. Simulate different scenarios that will help you deepen your understanding of various services. The combination of practical and theoretical knowledge will help you analyze difficult questions in the exam.

A few reminders that we can give is to always check the time and review your answers before proceeding to the next question (especially in the case study and yes/no questions). Before your scheduled exam day, don't forget to take a good rest. If you're not feeling confident yet, there's always an option to reschedule your exam. Good luck, and we wish you all the best.

CLOUD COMPUTING CONCEPTS

Cloud computing is the delivery of services over the Internet that helps you reduce your operating costs, run your infrastructure efficiently, and scale as business requirements change.

Cloud Service Models

- The three cloud computing service models are IaaS, PaaS, and SaaS.
- You can also use serverless computing to eliminate the need to manage infrastructure.
- The shared responsibility model determines the security tasks that are handled by the cloud provider and handled by the customer.
    - Azure is responsible for protecting the infrastructure such as hosts, network, and data center.
    - The customer is responsible for protecting their data, endpoints, account, and access management.
- IaaS, PaaS, and SaaS have different levels of managed services:

IaaS	PaaS	SaaS
Applications	Applications	Applications
Data	Data	Data
Runtime	Runtime	Runtime
MIddleware	MIddleware	MIddleware
O/S	O/S	O/S
Virtualization	Virtualization	Virtualization
Servers	Servers	Servers
Storage	Storage	Storage
Networking	Networking	Networking

You Manage Vendor Manages

Infrastructure as a service (IaaS)

- Most user management
- You are responsible for managing the **operating systems, data,** and **applications.**
- IaaS helps you to extend resources rapidly to meet the spikes required for your application.
- Used in the following scenarios:
    - **Migrating workloads –** move existing applications to the cloud.
    - **Test and development –** quickly set up and dismantle test and development environments. IaaS makes scaling development and testing environments fast and economical.
    - **Storage, backup, and recovery –** simplify the planning and management of backup and recovery systems.
    - **Website hosting** – less expensive than traditional web hosting.
    - **High-performance computing (HPC) –** clusters of computers that help solve complex problems involving millions of variables or calculations.
    - **Big data analysis –** for massive data sets that require a huge amount of processing power.

Platform as a service (PaaS)

- Less user management
- The operating systems are managed by the cloud provider, while the user is responsible for the applications and data they run and store.
- PaaS offers all the functionality you need to support the entire lifecycle of web applications: **building, testing the application**, **deploying the source code**, **managing**, and **updating** within the same integrated environment.
- Used in the following scenarios:
    - **Development framework –** a framework for creating or customizing cloud-based applications.
    - **Analytics or business intelligence –** find insights and patterns, and predict outcomes to improve business decisions.

Software as a service (SaaS)

- Least amount of management
- The cloud provider is responsible for managing everything, and the end-user just uses the software.

Serverless Computing

- Function as a Service (FaaS)
- You simply deploy the code with a serverless platform, and it runs at high availability.
- Dynamically scales up and down to meet the demands of each workload within seconds.

- A **pay-per-execution model** that charges sub-second billing only for the time and resources required to execute the code.

Cloud Architecture Models

- Three deployment methods of cloud computing: **Public vs Private vs Hybrid.**
- The model you choose for cloud deployment depends on your budget, security, scalability, and maintenance needs.

Public Cloud

- Focus on maintaining your applications without having to worry about purchasing, managing, or maintaining the hardware on which it runs.
- You can use multiple public cloud providers of varying scale.

Advantages	Disadvantages
High scalability/agility.	Specific security requirements.
Pay-as-you-go pricing.	Government policies, industry standards, or legal requirements.
You are not responsible for the updates and maintenance of the hardware.	You don't own the hardware or services and you also can't manage them as you may want to.
The required technical knowledge is minimal.	Maintaining a legacy application might be hard to meet.

Private Cloud

- A dedicated on-premises datacenter configured to be a cloud environment that provides users in your organization with self-service access to compute resources.
- You are responsible for the purchase and maintenance of the hardware and software services.
- You can use a private cloud when an organization has data that cannot be put in the public cloud, perhaps for legal reasons.

Advantages	Disadvantages
Any scenario or legacy application configuration is supported.	CapEx involved – principal cost is the procurement of the equipment.
You have control (and responsibility) over security.	To scale, you must buy, install, and set up new hardware.

Compliance or security requirements in your organization.	Private clouds require IT skills and expertise.

Hybrid Cloud

- Data and applications can move between **private** and **public clouds**.
- When there is a spike in demand in your private cloud, you can "burst through" to the public cloud for additional computing resources.

Advantages	Disadvantages
Maintain a private infrastructure for sensitive assets.	More expensive than selecting one deployment model since it involves some CapEx cost upfront
Take advantage of the resources in the public cloud when needed.	It can be more complicated to set up and manage
With the ability to scale to the public cloud, you pay for extra computing power only when needed.	
Allows you to use your own equipment to meet the security and compliance requirements in your organization.	

AZURE BASICS

Azure Overview

Azure is a cloud computing platform that was introduced by Microsoft in 2010. It enables you to create, manage, and deploy applications across a large global network. Microsoft Azure also provides a variety of services to assist your business in addressing current and potential business challenges in your infrastructure and applications.

Today, Microsoft Azure has the second-largest share in the cloud industry. It also has specialized regions for compliance or legal purposes.

Advantages of Azure Cloud Computing

- **Cost –** Eliminate the capital expense of buying hardware, software, and setting up of data centers. The principle of the cloud is, you will only pay for the computing resources you have consumed.
- **Global scale –** One of the benefits of cloud computing is the ability to scale elastically. This means that you can easily add resources such as compute and storage capacity in different regions with just a few clicks.
- **Performance –** Cloud computing services are hosted on a global network of secure data centers that are upgraded with the latest generation of computing hardware on a regular basis. Compared to a single corporate datacenter, this has several advantages, including lower application network latency and greater economies of scale.
- **Security –** Cloud service providers offer a broad set of policies, technologies, and controls to protect your data and infrastructure against potential threats.
- **Speed** – In a cloud computing environment, you can provision computing resources in minutes with just a few clicks. Providing businesses with a great deal of flexibility and relieving capacity planning pressure.
- **Productivity –** The cloud provides a lot of convenience to your IT teams since it reduces the time needed to obtain additional resources, allowing them to focus solely on achieving more important business goals.
- **Reliability –** With cloud computing, you can easily manage backup data, disaster recovery and business continuity since the data can be mirrored at multiple redundant sites.

Azure Global Infrastructure

Regions

- Each region has more than one data center, which is a physical location.
- A group of data centers deployed in a latency-defined perimeter and connected through a dedicated regional low-latency network.
- Criteria in choosing a Region:
    - **Location –** a region closest to your users minimizes the latency.
    - **Features –** some features are not available in all regions.
    - **Price –** the price of services varies from region to region.
- Each Region is paired within the same geographic area.
- If the primary region has an outage, you can **failover** to the secondary region.
- You can use paired regions for **replication.**
- Regions that are unique when it comes to compliance:
    - **Azure Government Cloud –** only US federal, state, local, and tribal governments and their partners have access to this dedicated instance.
    - **China Region –** data center is physically located within China and has no connection outside of China, including other Azure regions.

Availability Zones

- Each availability zone is a physical location within a region.
- A zone is composed of one or more data centers with independent power, cooling, and networking facilities.
- Azure services that support Availability Zones fall into two categories:
    - **Zonal services –** a resource is pinned to a specific zone.
    - **Zone-redundant services –** replicates automatically across zones.

Azure Security and Compliance

In the cloud, the responsibility of security is a shared one. Microsoft Azure secures what they can on their end, while you secure what you can on your end. Only this way can everyone protect their valuable data. Also as a customer, you inherit all the best practices of Azure policies, architecture, and operational processes built to satisfy the requirements of their most security-sensitive customers.

Microsoft Azure has also developed multiple tools and services to help you achieve your security objectives. You can also review the numerous audits and certifications that third-party auditors have conducted on Azure so that whenever you need to fulfill strict compliance with the use of a service, you can simply verify its status through the catalog.

Azure Pricing

- Azure offers pay-as-you-go and reserved instances for pricing.
- Azure Pricing Factors:
    - Resource size and resource type.
    - Different Azure locations have different prices for services.
    - The bandwidth of your services.
    - Any data transfer between two different billing zones is charged.
        - **Ingress (data in)** = free
        - **Egress (data out)** = charged based on data going out of Azure datacenters.
- Factors that can reduce costs:
    - By purchasing a **reserved instance** (one-year or three-year terms), you can significantly reduce costs by up to 72 percent compared to pay-as-you-go pricing.
    - A **reserved capacity** is a commitment for a period of one or three years for SQL Database and SQL Managed Instance.
    - **Hybrid Benefit** allows you to use your on-premises Software Assurance-enabled Windows Server and SQL Server licenses on Azure.
    - If you purchase an unused compute capacity, you can get deep discounts up to 90 percent compared to pay-as-you-go pricing. **A spot virtual machine** is for workloads that can tolerate interruptions.
- All resources belong to a **subscription**.
    - An Azure account can have multiple subscriptions.
    - Organize your resources and subscriptions using **Azure management groups**.
- **Azure Cost Management** gives you a detailed view of current and projected costs.
- For new accounts, the **Azure Free Tier** is available.
    - Free Tier offers limited usage of Azure products at no charge for 12 months.
    - You also get $200 credit that you can spend during the first 30 days.
    - More details at https://azure.microsoft.com/en-us/free/

  - Estimate your expected monthly costs using **Azure Pricing Calculator.**
- **Total Cost of Ownership (TCO) Calculator**
  - Estimate total savings over a period of time by using Azure.
  - Compares costs and savings against on-premises and co-location environments.
- **Azure Support Plans:**
  - **Basic –** included for all Azure customers.
  - **Developer –** recommended for non-production environments. Limited access to technical support during business hours by email only.
  - **Standard –** appropriate for production workload environments. Has 24/7 access to Azure's technical support engineers by phone or email.
  - **Professional Direct –** suitable for business-critical workloads. Has 24/7 access to Azure's technical support engineers by phone or email. Provides access to Operations Support, ProDirect delivery managers, and Support APIs.

Service Level Agreement (SLA)

- It is the commitment of Microsoft for the uptime and connectivity of a service.
- You could obtain a service credit if the service level agreement is not met by Microsoft.
- Composite SLAs include several resources (*with different availability levels*) to support an application.
- SLAs for multi-region deployments distribute the application in more than one region for high availability and use Azure Traffic Manager for failover if one region fails.

Service Lifecycle

- **Private Preview** is only available to a few customers for early access to new technologies and features.
- **Public Preview** makes the service in the public phase and can be used by any customers to evaluate the new features but SLA does not apply.
- **General Availability** is the release of service to the general public and is fully supported by SLAs.
- Azure updates allow you to get the latest updates on any Azure products and features.

Azure Well-Architected Framework - Five Pillars

- **Operational Excellence** - run and monitor systems to deliver business value and to continually improve supporting processes and procedures.
- **Reliability** - recover the system from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues.
- **Performance Efficiency** - use computing resources efficiently to meet system requirements and to maintain that efficiency as demand changes and technologies evolve.
- **Cost Optimization** - avoid or eliminate unneeded costs or suboptimal resources.
- **Security** - protect information, systems, and assets while delivering business value through risk assessments and mitigation strategies.

Best Practices when Architecting in the Cloud

- **Design for self healing** - Failures occur in a distributed system. Design your application to be self-healing in the event of failure.
- **Make all things redundant** - Design a resilient and highly available application to avoid single points of failure.
- **Minimize coordination** - To achieve scalability, you must minimize coordination between application services.
- **Design to scale out** - Design an application that can scale horizontally (adding or removing new instances) as needed.
- **Partition around limits** - Use partition for database, network, and compute limits
- **Design for operations** - The operations team must be able to access the tools they need for the application.
- **Use managed services** - When designing an application, use PaaS rather than IaaS.
- **Use the best data store for the job** - Select the storage technology that is most appropriate for your data and its intended use.
- **Design for evolution** - An evolutionary design is required for continuous innovation.
- **Build for the needs of business**- Always consider the business requirements when designing an application.

THE DIFFERENT AZURE SERVICES

Compute is the processing power required by applications and systems to carry out computational tasks.
Services: Virtual Machine, App Service, Functions, and Kubernetes Service

Storage in the cloud is used to store different types of data, such as objects, files, and backups.
Services: Blob, Disk, and Files

Database is a system to store and manage structured and unstructured information.
Services: SQL Database and Cosmos DB

Networking provides a global link to distribute the application all over the world.
Services: Virtual Network, Load Balancer, CDN, and DNS

Security allows you to authenticate and authorize users and services to access your applications.
Services: Microsoft Entra ID, RBAC, and Security Center

Management and Governance is a tool to control and monitor your infrastructure services.
Services: Monitor, Policy, and Advisor

DEEP DIVE

Azure Virtual Machines

Components of a Virtual Machine

1. When creating a virtual machine, you always start off by choosing a **subscription** and **resource group**. A subscription is a container where you can provision Azure resources. Before you can deploy resources, you also need to create a new resource group. This is a logical group to organize and manage all your resources in your subscription.

2. After you have chosen the resource group, you configure the **availability option** of your virtual machine. You can choose between the availability zone, availability set or no infrastructure redundancy option. The option you selected here would determine the availability and resiliency of your applications.

3. The **image** of your virtual machine contains the OS, settings, and other applications that you will use in your server. In the Azure Marketplace, you can choose between images provided by Microsoft or your own custom image

4. Once you have chosen the image of your virtual machine, select the **type and size** of your virtual machine. This will determine the physical properties of your instance, such as vCPUs, RAM, disks, and more.

5. During the creation of your virtual machine, you can also specify whether you'd like to launch it in a **spot instance** or use another instance billing type (pay as you go or reserved).

6. To access your virtual machine, you will need to use a **key pair**. It is generated after you launch your virtual machine. Make sure to secure your copy of your public key. Once you delete your public key, you wouldn't be able to directly access your instance.

7. After you have configured the basic settings, you need to add **storage** for your virtual machine. The disks that can be added are the operating system disk, data disk, and temporary disk. Encryption for your disks is automatically configured.

8. You also need to configure which **virtual network** the virtual machine should be launched in. And the **network security group** will serve as a firewall to your servers. It contains rules that allow or deny network traffic coming to or from your firewall.

9. When you have configured the network settings of your virtual machines, you can also enable **monitoring, auto-shutdown, and backup** in the management options.

10. In the advanced configuration option of your virtual machine, you can add **extensions** for post-deployment configuration, **custom and user data** to execute certain commands while the instance is being provisioned, and **proximity placement** group to enable you to group your resources closer in the same region.

11. Lastly, you can add **tags** to easily identify and classify your resources.

12. Once you have reviewed the configuration of your instance, proceed with the launch. Wait for your virtual machine to finish preparing itself, and you should be able to connect to it if there aren't any issues.

13. If you are having difficulties connecting to a virtual machine, you can try redeploying the VM to move it to a new node in the Azure infrastructure. Don't worry, all of the existing configurations in the resource will still be there after completing the redeployment.

Types of Virtual Machines

1. **General Purpose -** provides a balanced CPU-to-memory ratio. This instance is ideal for testing, development, and low to medium-traffic web servers. The B-series have burstable performance that allows the VM to use the build-up credits when the application requires higher CPU performance.

2. **Compute Optimized -** designed to have a high CPU-to-memory ratio. Instances belonging to this family are well suited for medium-traffic web servers, network appliances, batch processes, analytics, application, and gaming servers.

3. **Memory Optimized -** offers a high memory-to-CPU ratio. Ideal for relational database servers, medium to large caches, and in-memory analytics.

4. **Storage Optimized -** provides high disk throughput and IO. This VM size is ideal for SQL, NoSQL databases, big data, data warehousing, and large transactional databases.

5. **GPU -** designed for compute-intensive, graphics-intensive, and visualization workloads. It is available in single, multiple, or fractional GPUs.

6. **High-performance compute -** the HPC VM size is the most powerful and fastest CPU with high throughput network interfaces. It is optimized for fluid dynamics, explicit and implicit finite element analysis, weather modeling, seismic processing, reservoir simulation, and RTL simulation.

Virtual Machine Disks

The disks of a virtual machine are block-level storage volumes. This storage is managed by Azure and mainly used for Azure VMs. With managed disks, all you have to do is specify the type and size of the disk and provision it.

In Azure, there are three types of disk roles:

1. **Operating system (OS) disk** - in order for the virtual machine to operate, it must have an OS disk. There are a variety of images that you can choose from in the Azure Marketplace. An example of images that you can use are Windows Server, Ubuntu, Debian, RHEL, etc. There are two types of OS disk:
    a. Persistent OS disk - this type of disk supports all sizes of VM, and the data is preserved even if you upgrade your OS disk and VM size.
    b. Ephemeral OS disk - use ephemeral OS disks if you need lower read/write latency and faster VM reimage. This type of disk is ideal for stateless applications, and it can be stored on VM cache or VM temp/resource disk if sufficient space is available.

2. **Data disk** - this disk is also managed by Azure, and you can store your application data or any other data that you need to keep. Before you use a data disk, there are two options that you can select:
    a. Create and attach a new disk - you have the option to create the new disk from a snapshot, storage blob, or an empty disk.
    b. Attach an existing disk - allows you to add the disks you've already created. It's also important to know that the number of data disks that you can attach will depend on the size of your VM.

3. **Temporary disk** - provides you short-term storage to store pages and swap files. Take note that the data on this disk may be lost when you redeploy a VM or during a maintenance event. Also, to configure a server-side encryption on this disk, you need to enable encryption at host.

The available disks that you can choose from are:

1. **Ultra Disk** - ideal for IO-intensive workloads, top-tier databases, and other transaction-heavy workloads. This storage has the highest disk size, throughput, and IOPS.
2. **Premium SSD** - designed for production and performance-sensitive workloads.
3. **Standard SSD** - used for web servers and dev/test environments.
4. **Standard HDD** - ideal for backup, non-critical data, and infrequent access.

Detail	Standard HDD	Standard SSD	Premium SSD	Ultra Disk
Disk type	HDD	SSD	SSD	SSD
Scenario	Backup, non-critical, infrequent access	Web servers, and light applications of enterprise	Production and performance sensitive workloads	IO-intensive workloads, top tier databases, and other transaction-heavy workloads
Max Disk Size	32,767 GiB	32,767 GiB	32,767 GiB	65,536 GiB
Max Throughput	500 MB/s	750 MB/s	900 MB/s	2,000 MB/s
Max IOPS	2,000	6,000	20,000	160,000

It's also very important to understand how you can secure your data inside your virtual machine disks. Let's now take a look at disk encryption, Azure managed disks supports three types of encryption:

1. **Server Side Encryption (SSE)** - the data stored on managed disks are automatically encrypted at rest by default when persisting it to the cloud.
    a. Platform-managed keys - the keys are managed by Azure. The data, images, and snapshots written to an existing managed disks are automatically encrypted-at-rest.
    b. Customer-managed keys - since you are providing your own keys, you also manage the level of encryption on each managed disk. To manage your own keys, you can use Azure Key Vault. This service enables you to import your own RSA keys or generate a new ones.

2. **Azure Disk Encryption (ADE)** - provides volume encryption on both OS and data disks of Azure VMs. The encryption for Windows is done using BitLocker. On the other hand, the encryption for Linux is done using DM-Crypt.

3. **Encryption at host -** this type of encryption is different from SSE. The encryption of data is provided by the server hosting your virtual machine and the encrypted data flows into the Azure Storage service.

	Encryption at rest	Temp Disk Encryption	Encryption of Caches	Encrypted Data Flows	Customer Keys	Encryption Status
Encryption at rest with PMK	✓	-	-	-	-	Unhealthy
Encryption at rest with CMK	✓	-	-	-	✓	Unhealthy
Azure Disk Encryption	✓	✓	✓	✓	✓	Healthy
Encryption at Host	✓	✓	✓	✓	✓	Unhealthy

Note:
- The encrypted data flows are between Compute and Storage service.
- The disk encryption status is labeled by Azure Security Center.

When creating a copy of your managed disks, there are comparisons between images and snapshots. As discussed earlier in data disks, snapshots allow you to create a point in time recovery. But how is it different from images?

Let's look at the differences between the two:
1. **Snapshots -** a full, read-only copy of your virtual hard drive. It can be taken at any point in time. The existence of a managed disk snapshot is independent of the source disk. This means that it applies only to one disk. You can also use snapshots to create a new disk and attach it to a virtual machine.
2. **Images -** contain all the managed disks associated with the virtual machine. The created image can be used to launch hundreds of virtual machines without managing any storage accounts.

To conclude the comparison, a snapshot is only aware of the disk that it contains. For scenarios that require the coordination of multiple disks, like striping, snapshot wouldn't be able to meet this requirement. Therefore, this is where you would want to use custom images.

When talking about how the virtual machine handles unexpected disk traffic, Azure offers a feature called **bursting**. This will grant the virtual machine and disk the ability to boost the IOPS and MB/s performance for a period of time. In other words, it will allow you to get more use out of your disk and also helps you avoid upgrading the disk just to accommodate traffic spikes. The bursting on virtual machines and disks are independent from one another. So if you need to burst the disk performance, you don't need to burst the virtual machine. Bursting is enabled by default for both virtual machine and disk.

The following resources support bursting:

1. **Burstable Virtual Machines:**
    a. **General Purpose:** B, Dsv3, Dasv4, Ddsv4, and Dsv4 series
    b. **Compute Optimized:** Fsv2 series
    c. **Memory Optimized:** Esv3, Easv4, Edsv4, and Esv4 series
    d. **Storage Optimized:** Lsv2 series
2. **Burstable Disk:**
    a. Premium SSD
    b. Standard SSD

Payment options for Virtual Machines

Azure provides you with a variety of options to pay for compute capacity. Here are the following payment options:

1. **Pay as you go -** you are billed on a per-second basis. You can start or stop anytime, and you only pay for what you use. This payment option is ideal for users who prefer flexibility or have unpredictable workloads that cannot be interrupted.
2. **Reserved Instance -** you get up to 72 percent price savings compared to pay-as-you-go, but in return, you need to pay the upfront cost and be committed for one or three years in a specified region. There are three options to scope a reservation:
    a. **Single resource group -** the reservation discount applies solely to the corresponding resources in the resource group you've chosen. Keep in mind that discounts will not be applied if the resource group is moved or deleted.
    b. **Single subscription -** the reservation discount applies only to the corresponding resources in the subscription you've selected.
    c. **Shared -** the reservation discount is applied to the corresponding resources in eligible subscriptions within the billing context. If the subscription is moved to a different billing

context, the discounts no longer apply to that subscription but will continue to apply to the remaining subscriptions in the billing context.
    i. The billing context for Enterprise Agreement customers is enrollment. In an enrollment, the reservation shared scope contains multiple Microsoft Entra ID tenants.
    ii. The billing scope for Microsoft Customer Agreement customers is billing profile.
    iii. The billing scope for individual subscriptions with pay-as-you-go rates is all eligible subscriptions.

After purchasing a reservation, you can always update the scope. Go to the reservation, click **Configuration**, and then rescope the reservation. Rescoping a reservation won't change the reservation term.

3. **Spot** - save up to 90 percent when you purchase unused compute capacity. This is only ideal for workloads that can tolerate interruptions. Discounts may vary based on:
    a. Region
    b. Virtual machine type
    c. Compute capacity

Since Azure Spot Virtual Machines are unused capacity, at any point in time, Azure infrastructure can evict Spot VMs with 30 seconds notice. Eviction is based on the capacity or the max price you've set. When creating a Spot VMs, you can set the eviction policy to Deallocate (default) or Delete.

The Deallocate policy moves your virtual machine to the stopped-deallocated state, allowing you to redeploy it later. However, there is no assurance that the allocation will be successful. Your quota will be depleted by the deallocated VMs, and you will be charged for the underlying disks.

If you want your virtual machines to be deleted when it is evicted, you can set the eviction policy to Delete. The underlying disks are also deleted, so you won't be charged for the storage. In the portal, you can look up the eviction rates by size in a certain region. Go to **View pricing history and compare prices in nearby regions** to see a table or graph of pricing for a specific size.

Availability Options for Virtual Machines

There are two ways to manage the availability and resiliency of your applications in a virtual machine:

1. **Availability zones** - to protect your resources from an entire data center failure, you need to deploy the VMs to a minimum of three Availability Zones to ensure resiliency. Azure services that support Availability Zones are classified into two types:
    a. Zonal services - resources are pinned to a specific Availability Zone.
    Examples: Virtual machines, Managed disks, Standard IP addresses

   b. Zone-redundant services - replicate resources automatically across Availability Zones to protect from single points of failure.
      Examples: Zone-redundant storage, SQL Database
2. **Availability sets -** to protect from hardware failures within a data center, you can deploy the virtual machine to an availability set. Each VM in an availability set is assigned to an update domain and fault domain. This option ensures that at least one is available during planned or unplanned maintenance events.
   a. Update domains (planned maintenance)
      i. A logical group of virtual machines that can undergo maintenance at the same time. By default, it has five non-user-configurable update domains. It can be increased up to 20 update domains and given 30 minutes to recover before maintenance is initiated on a different update domain.
   b. Fault domains (unplanned maintenance)
      i. A logical group of virtual machines that share a common power source and network switch. By default, VMs within an availability set are separated up to three fault domains.

Virtual Machine Scale Sets

When you need to improve the performance of your applications and also provide redundancy, you should scale your resources horizontally. Horizontal scaling means you are adding more servers to the system. By doing this, the workload will be distributed across multiple resources and accommodate the increasing demand. Take note that this type of scaling is different from vertical scaling. When you say scale vertically, you are increasing or decreasing the resources of a single server instead of adding new servers to the system.

The horizontal scaling service in Azure is called **virtual machine scale sets**. A VM scale set allows you to create and manage a group of load-balanced VMs. Since the workload is distributed, if one VM fails, you can still continue to access your application through other VMs with minimal interruption. You can also distribute VMs in a scale set within a single data center or across various data centers. This service supports both layer 4 basic traffic distribution and layer 7 advanced traffic distribution and TLS termination.

Virtual Machine Scale Sets provide the following key benefits:

- By creating scaling policies, you can automatically add or remove virtual machines based on host metrics. A host metric provides you visibility into the performance of the virtual machines in a scale set without the need to install and configure agents. An example of host metrics can be CPU Utilization, Network In, and many more.

- You can create health checks and set a repair policy to automatically replace unhealthy virtual machines. Unhealthy instances are reported by Application Health extension or Load Balancer health probes.
- You can associate virtual machine scale sets with a load balancer. This will allow you to distribute virtual machines across Availability Zones. By implementing this practice, you can make your application redundant and highly available.
- Lastly, virtual machine scale sets allow you to scale hundreds or even thousands of virtual machines.

Now that we know scale sets can be associated with load balancers, this will help us implement one of the best practices on architecting in the cloud by evenly distributing the virtual machines across different Availability Zones. The main reason why you need to configure it with a load balancer is to give you high availability. An application that can run continuously even if one of the virtual machines fails. Aside from distributing the load across AZs, one of the added benefits is you can use Load Balancer health probes for more robust health checks.

When associating scale sets with a load balancer, you have two options:

1. **Azure Application Gateway -** is an HTTP/HTTPs web traffic load balancer that has the capability to do the following: URL-based routing, SSL termination, session persistence, and web application firewall.
2. **Azure Load Balancer -** a TCP/UDP network traffic load balancer that supports port forwarding and outbound flows.

After going through load balancing, let's now talk about the scaling policy and how it works. A **scaling policy** can determine when a virtual machine should be added or removed to meet the current capacity requirements of your application. When you create a virtual machine scale set, you would see this configuration in the portal.

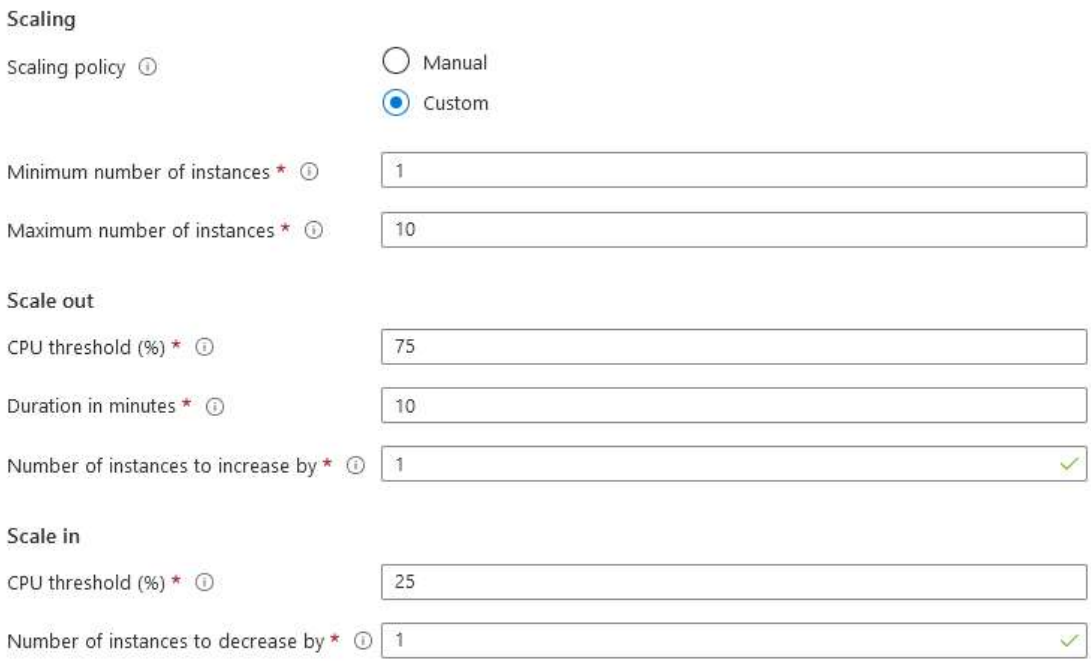

These configurations can only be seen if you select the custom scaling policy option. The first thing that you can set is the number of instances. But let's focus on the remaining two options, the scaling out and the scaling in. **Scale out** is when you need to add virtual machines to the scale set to increase the current capacity. In order to scale out, you should input certain values on the following fields:

- CPU threshold - is the CPU usage percentage threshold on when to trigger the scale out rule.
- Duration in minutes - is the amount of time that the autoscale will check the threshold again.
- Number of instances to increase by - this will determine how many virtual machines should be added when the scale out rule is triggered.

On the other hand, the **scale in** rule is when should the scale sets remove a virtual machine in order to decrease the capacity. Unlike scale out, you only need to input two values in the scale in fields. After you create a virtual machine scale set, you will see a lot of options available that you can configure in the scaling policy.

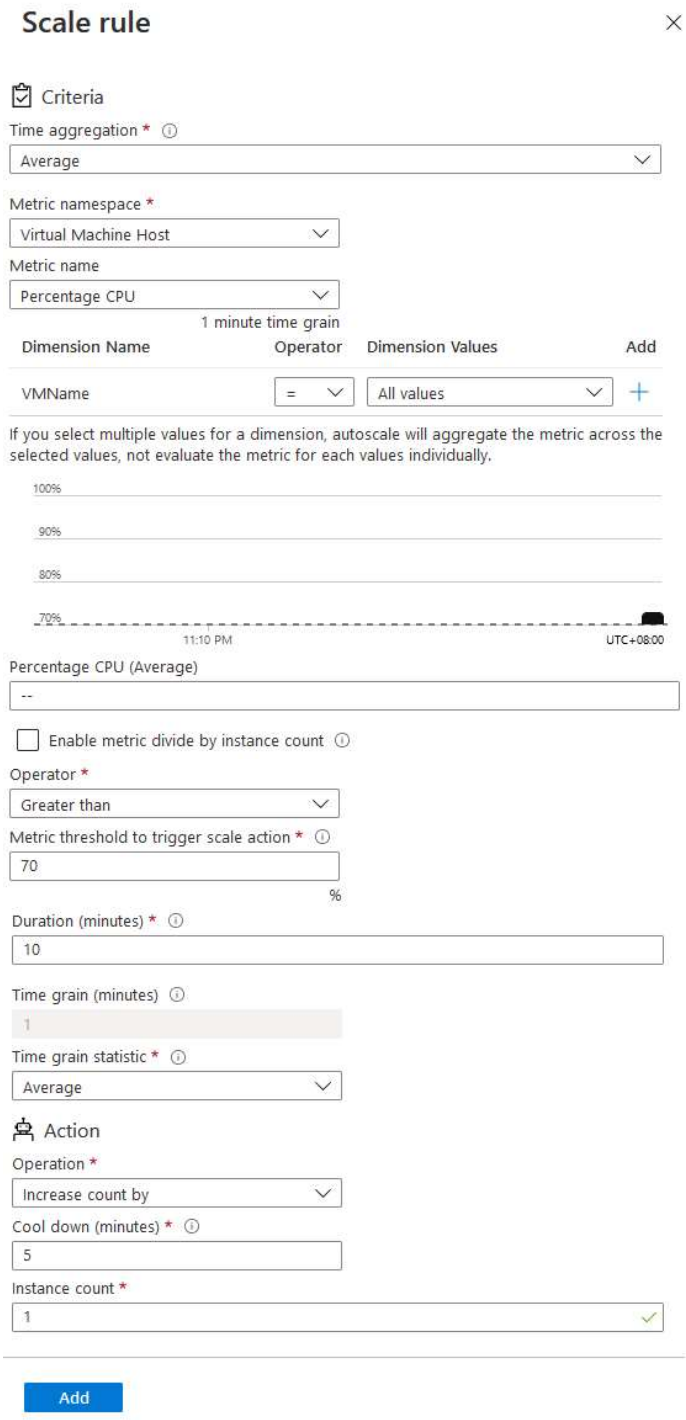

As seen in the image above, you can still configure other options in order to meet certain requirements on when to scale your virtual machines. Here are the options that you can customize:

1. **Metric Name -** allows you to set the metric that will be collected to your virtual machine. Some of the metrics that you can choose from are:
    - Percentage CPU

- Network In or Out
- Disk Read or Write Bytes
- Disk Read or Write Operations/Sec
- CPU Credits Consumed or Remaining

2. **Aggregates** - it is how you want to collect the data. For example, TimeAggregation = “Sum” will aggregate the sampled metrics by taking the sum. The methods that you can select from are:
   - Average
   - Minimum
   - Maximum
   - Sum
   - Last
   - Count
3. **Operators** - this will determine when to trigger scale action.
   - Greater than
   - Greater than or equal to
   - Less than
   - Less than or equal to
   - Equal to
   - Not equal to
4. **Actions** - what should the scaling policy do after it is triggered.
   - Increase count by
   - Increase percent by
   - Increase count to
   - Decrease count by
   - Decrease percent by
   - Decrease count to

If you want to collect more information based on different metrics, you need to install the following:

- **App Insights** - when you want to collect application metrics such as page load performance and session counts, you can install app insights in your application, and it will monitor your app and send telemetry to Azure.
- **Azure Diagnostic Extension** - when you want detailed Host-based metrics, you can install this extension. This agent will run inside your virtual machine. It will monitor and save performance metrics to an Azure storage service to collect more detailed information.