PassQueen
a.k.a dumpsqueen.com

Designing Microsoft Azure Infrastructure Solutions

Microsoft AZ-305

Version Demo

Total Demo Questions: 39

Total Premium Questions: 396

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design identity, governance, and monitoring solutions	126
Topic 2, Design data storage solutions	115
Topic 3, Design business continuity solutions	35
Topic 4, Design infrastructure solutions	113
Topic 5, Mix Questions	1
Topic 6, Case Study 1	1
Topic 7, Case Study 2	2
Topic 8, Case Study 3	3
Total	396

QUESTION NO: 1

You have data files in Azure Blob Storage.

You plan to transform the files and move them to Azure Data Lake Storage.

You need to transform the data by using mapping data flow.

Which service should you use?

A. Azure Data Box Gateway

B. Azure Databricks

C. Azure Data Factory

D. Azure Storage Sync

ANSWER: C

Explanation:

Azure Data Factory is the appropriate service because mapping data flows are a native Azure Data Factory capability for designing and running visually authored, scalable data transformations without requiring Spark code. A pipeline can orchestrate the complete process: it can read the source files from Azure Blob Storage, execute the mapping data flow to apply the required transformations, and write the resulting files to Azure Data Lake Storage Gen2. Mapping data flows run on managed Azure integration runtime compute and support common transformation tasks such as selecting and deriving columns, filtering rows, joining datasets, aggregating data, and altering schemas. Azure Data Factory supports Azure Blob Storage as a source and Azure Data Lake Storage Gen2 as a sink through its storage connectors, enabling the required movement between the two services. This combination provides a managed, repeatable, and monitorable ETL or ELT workflow, with pipeline scheduling and operational monitoring available in the Azure Data Factory service. Microsoft documents mapping data flows as the visual transformation component of Azure Data Factory and documents the relevant storage connectors for reading and writing file-based data. See Mapping data flow overview and Azure Data Lake Storage Gen2 connector.

QUESTION NO: 2

You have an Azure subscription that contains the resources shown in the following table.

You need to recommend a load balancing solution that will distribute incoming traffic for VMSS1 across NVA1 and NVA2. The solution must minimize administrative effort.

What should you include in the recommendation?

A. Gateway Load Balancer

B. Azure Front Door

C. Azure Application Gateway

D. Azure Traffic Manager

ANSWER: A

Explanation:

Gateway Load Balancer is designed for transparent service insertion of network virtual appliances (NVAs), such as firewalls, intrusion-detection systems, and packet-inspection appliances. It can place NVA1 and NVA2 in its backend pool and distribute traffic flows across them while preserving flow symmetry. This enables incoming traffic intended for VMSS1 to be sent through an available NVA without requiring application-level proxy configuration or custom routing on the workload instances.

A Gateway Load Balancer is chained to a standard Azure Load Balancer frontend through a tunnel interface. The workload-facing load balancer continues to provide the application or network load-balancing entry point, while the Gateway Load Balancer transparently forwards traffic to the NVA fleet for inspection and then returns it to the original path. This architecture scales the NVA layer independently of VMSS1 and uses Azure-managed health probes and flow distribution, reducing the operational work required to steer traffic through multiple appliances.

Microsoft describes Gateway Load Balancer as a high-performance, highly available solution for deploying and scaling third-party NVAs transparently: Azure Gateway Load Balancer overview. For the supported chaining design and traffic flow, see Deploy Gateway Load Balancer.

QUESTION NO: 3 - (SIMULATION)

You have an Azure subscription that contains the resources shown in the following table.

Log files from App1 are ingested to App 1 Logs. An average of 120 GB of log data is ingested per day.

You configure an Azure Monitor alert that will be triggered if the App1 logs contain error messages.

You need to minimize the Log Analytics costs associated with App1. The solution must meet the following requirements:

• Ensure that all the log files from App1 are ingested to App 1 Logs.

• Minimize the impact on the Azure Monitor alert.

Which resource should you modify, and which modification should you perform? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

ANSWER: See the explanation for the answer

Explanation:

Modify: App1 Logs (the Log Analytics workspace).

Modification: Configure a **commitment tier / capacity reservation of 100 GB per day**.

The workspace ingests approximately 120 GB per day. A 100-GB/day commitment tier reduces ingestion cost compared to pay-as-you-go, while the additional approximately 20 GB/day is billed as overage. This retains the logs in the Analytics table plan, so all logs continue to be ingested and the existing Azure Monitor log alert continues to work with minimal impact.

Do not move the data to the Basic log plan, because Basic logs have query and alerting limitations that can affect the Azure Monitor alert.

QUESTION NO: 4

You are designing an internet-facing application that will run on virtual machines in a single Azure region.

The solution must meet the following requirements:

• Load balance HTTP and HTTPS traffic.

• Route requests to different backend pools based on the URL path.

• Provide Web Application Firewall protection.

• Scale automatically as request volume changes.

What should you include in the recommendation?

A. Azure Application Gateway WAF v2

B. Azure Load Balancer Standard

C. Azure Traffic Manager

D. Azure VPN Gateway

ANSWER: A

Explanation:

Azure Application Gateway WAF v2 is the appropriate recommendation. Application Gateway is a regional Layer 7 load-balancing service that supports HTTP and HTTPS traffic. Its URL path-based routing capability can send requests for different paths, such as `/images` and `/api`, to separate backend pools.

The WAF v2 SKU includes Azure Web Application Firewall capabilities that protect web applications against common web exploits and vulnerabilities. The v2 SKU also supports autoscaling, allowing the gateway capacity to adjust based on traffic demand. Azure Load Balancer operates at Layer 4 and does not provide URL path routing or Web Application Firewall capabilities, while Azure Front Door is primarily selected for global application delivery across regions.

See Azure Application Gateway v2 overview and Web Application Firewall on Azure Application Gateway.

QUESTION NO: 5

You are designing a large Azure environment that will contain many subscriptions.

You plan to use Azure Policy as part of a governance solution.

To which three scopes can you assign Azure Policy definitions? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. management groups

B. subscriptions

C. resource groups

D. Azure Active Directory (Azure AD) administrative units

E. compute resources

ANSWER: A B C

Explanation:

Azure Policy definitions and initiatives are applied by creating a policy assignment at an Azure Resource Manager scope. **management groups** are a suitable scope for a large, multi-subscription environment because an assignment made there is inherited by the subscriptions, resource groups, and resources below that management group. This lets central governance teams consistently enforce organizational requirements, such as allowed locations, required tags, or security configuration baselines.

subscriptions are also valid assignment scopes. A subscription-level assignment applies to resources in that subscription and is useful where a policy must apply across all workloads in one billing or isolation boundary.

resource groups are valid for assigning policy to a defined set of related resources. This supports workload-specific governance when requirements differ among resource groups in the same subscription. Azure Policy evaluates resources within the assignment scope and can enforce or audit compliance according to the assigned definition or initiative. For the supported assignment scope hierarchy and inheritance behavior, see Azure Policy overview and Assign Azure Policy definitions through the Azure portal.

QUESTION NO: 6

You have an Azure subscription that contains a Windows Virtual Desktop tenant.

You need to recommend a solution to meet the following requirements:

Start and stop Windows Virtual Desktop session hosts based on business hours.

Scale out Windows Virtual Desktop session hosts when required.

Minimize compute costs.

What should you include in the recommendation?

A. Microsoft Intune

B. a Windows Virtual Desktop automation task

C. Azure Automation

D. Azure Service Health

ANSWER: C

Explanation:

Azure Automation is the appropriate recommendation because it provides runbooks that can automate the operational lifecycle of session-host virtual machines. A runbook can be scheduled to start the required session hosts before business hours and stop or deallocate them after business hours, avoiding charges for VM compute when capacity is not needed. Azure Automation also supports PowerShell-based logic that can query host-pool usage and session information, calculate the required capacity, and start additional session hosts as demand rises. This allows the organization to retain enough available hosts for users while reducing the number of running VMs during quiet periods. Runbooks can run on schedules and can be triggered or integrated with monitoring and automation workflows, making the approach suitable for both planned business-hour operations and demand-based scaling. Deallocating unused session-host VMs is especially important for minimizing compute cost, because compute billing stops when a VM is deallocated, although associated resources such as managed disks can still incur charges. Microsoft documents Azure Automation as a service for process automation through runbooks and schedules, and documents autoscaling concepts for Azure Virtual Desktop host pools at Azure Automation overview and Azure Virtual Desktop autoscale scaling plans.

QUESTION NO: 7

Your company has the infrastructure shown in the following table.

Location	Resource
Azure	• Azure subscription named Subscription1 • 20 Azure web apps
On-premises datacenter	• Active Directory domain • Server running Azure AD Connect • Linux computer named Server1

The on-premises Active Directory domain syncs with Azure Active Directory (Azure AD).

Server1 runs an application named App1 that uses LDAP queries to verify user identities in the on-premises Active Directory domain.

You plan to migrate Server1 to a virtual machine in Subscription1.

A company security policy states that the virtual machines and services deployed to Subscription1 must be prevented from accessing the on-premises network.

You need to recommend a solution to ensure that App1 continues to function after the migration. The solution must meet the security policy.

What should you include in the recommendation?

A. Azure AD Application Proxy

B. the Active Directory Domain Services role on a virtual machine

C. an Azure VPN gateway

D. Azure AD Domain Services (Azure AD DS)

ANSWER: D

Explanation:

Azure AD Domain Services (Azure AD DS) is the appropriate recommendation because it supplies managed domain services within Azure without requiring the migrated virtual machine to contact the on-premises network. Azure AD DS provides a managed domain containing synchronized identity information from Microsoft Entra ID (formerly Azure AD) and supports the directory capabilities that a legacy LDAP-dependent application requires, including LDAP queries, domain join, Kerberos authentication, NTLM authentication, and Group Policy.

App1 can therefore query LDAP against the managed domain endpoints in the Azure virtual network after Server1 is migrated. The organization can deploy Azure AD DS in a virtual network available to the workload in Subscription1 while maintaining network isolation from the on-premises Active Directory environment. This preserves the application's directory-based identity validation model while satisfying the requirement that Subscription1 resources not access the on-premises network. Azure AD DS is a managed service, so Microsoft operates the domain controllers and the organization does not need to deploy, administer, or replicate its own Active Directory Domain Services domain controllers for this scenario.

For service capabilities and deployment guidance, see Microsoft Entra Domain Services overview and Network considerations for Microsoft Entra Domain Services.

QUESTION NO: 8

You have an Azure subscription that contains production resources.

You need to provide members of the operations team with the Contributor role for production resources. The solution must meet the following requirements:

• The users must not have permanent active Contributor permissions.

• Users must activate the role only when required.

• Role activation must require multifactor authentication and justification.

Which two features should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create eligible Contributor role assignments by using Microsoft Entra Privileged Identity Management.

B. Configure Microsoft Entra Privileged Identity Management activation settings to require multifactor authentication and justification.

C. Create permanent active Contributor role assignments.

D. Create an Azure Policy assignment that has a Deny effect.

E. Configure a management lock on the production resource groups.

ANSWER: A B

Explanation:

Microsoft Entra Privileged Identity Management (PIM) supports **eligible Azure resource role assignments**. An eligible assignment does not grant active Contributor access until a user activates the role. This reduces standing privilege while allowing operations personnel to obtain access for a limited activation period when required.

PIM role settings can require users to provide a justification and perform multifactor authentication during activation. These activation controls create an auditable, just-in-time elevation process for the Contributor role. Permanent active role assignments do not satisfy the requirement because they provide continuous permissions and do not require activation.

See Assign Azure resource roles in Privileged Identity Management and Configure Azure resource role settings in Privileged Identity Management.

QUESTION NO: 9

You have an Azure subscription that contains a Basic Azure virtual WAN named Virtual/WAN1 and the virtual hubs shown in the following table.

You have an ExpressRoute circuit in the US East region.

You need to create an ExpressRoute association to VirtualWAN1.

What should you do first?

A. Upgrade VirtualWAN1 to Standard.

B. Create a gateway on Hub1.

C. Create a hub virtual network in US East.

D. Enable the ExpressRoute premium add-on.

ANSWER: A

Explanation:

Upgrade VirtualWAN1 to Standard. A Basic Azure Virtual WAN supports only site-to-site VPN connectivity; it does not support ExpressRoute connectivity. An ExpressRoute circuit can be associated with a Virtual WAN only when the Virtual WAN uses the Standard SKU. Therefore, upgrading the existing Basic virtual WAN is the required prerequisite before an ExpressRoute gateway can be deployed in an appropriate virtual hub and the circuit connection can be created.

The Standard virtual WAN SKU provides the capabilities required for ExpressRoute, including ExpressRoute gateway deployment within a virtual hub and connectivity between the ExpressRoute circuit and resources connected to the Virtual WAN. The upgrade is performed at the Virtual WAN resource level, and Microsoft documents that a Basic virtual WAN can be upgraded to Standard. After the upgrade, the ExpressRoute configuration can proceed using the supported gateway and circuit-connection workflow. For the relevant SKU capabilities and upgrade guidance, see About Azure Virtual WAN and About ExpressRoute in Azure Virtual WAN.

QUESTION NO: 10

You have an Azure virtual machine named VM1 that runs Windows Server 2019 and contains 500 GB of data files.

You are designing a solution that will use Azure Data Factory to transform the data files, and then load the files to Azure Data Lake Storage

What should you deploy on VM1 to support the design?

A. the self-hosted integration runtime

B. the Azure Pipelines agent

C. the On-premises data gateway

D. the Azure File Sync agent

ANSWER: A

Explanation:

The self-hosted integration runtime is the required component to deploy on VM1. Azure Data Factory uses an integration runtime as the compute and connectivity infrastructure for moving data between a source and a destination. A self-hosted integration runtime runs on a customer-managed Windows computer, including an Azure virtual machine, and enables Data Factory to access data stored locally on that machine or within a private network environment. In this scenario, the 500 GB of source files reside on VM1. Installing and registering a self-hosted integration runtime on VM1 lets a Data Factory pipeline securely read those files and copy them to Azure Data Lake Storage. It can also support pipeline activities that require connectivity to resources accessible from VM1. The runtime initiates outbound connections to Azure, so it avoids requiring inbound access from the Data Factory service to the virtual machine. For resilience or higher throughput, a self-hosted integration runtime can be configured with multiple nodes, but one runtime installed on VM1 is sufficient to support the stated design. See Create and configure a self-hosted integration runtime and Integration runtime in Azure Data Factory.

QUESTION NO: 11

You are designing a large Azure environment that will contain many subscriptions.

You plan to use Azure Policy as part of a governance solution.

To which three scopes can you assign Azure Policy definitions? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. management groups

B. subscriptions

C. C. Azure Active Directory (Azure AD) tenants

D. resource groups

E. Azure Active Directory (Azure AD) administrative units

F. compute resources

ANSWER: A B D

Explanation:

Azure Policy definitions, including initiative definitions, are applied by creating a policy assignment at a supported Azure scope. **management groups** are appropriate for governance that must be inherited by multiple subscriptions, which is especially useful in a large environment. A policy assignment at a management group applies to subscriptions and resource groups beneath it, subject to any exclusions configured in the assignment.

subscriptions are also a supported assignment scope. This scope is useful when a policy must govern all applicable resources in one subscription without applying the same rule across every subscription in the management-group hierarchy.

resource groups are a supported scope for assignments that apply to resources deployed within a particular resource group. This enables more targeted governance when different workloads or environments within a subscription have distinct requirements. Azure Policy also supports assignment at an individual resource scope, but the supported scope is the specific resource rather than a broad resource category. For the three requested assignment scopes, management groups, subscriptions, and resource groups are the applicable selections. See Azure Policy overview and Assign Azure Policy definitions.

QUESTION NO: 12

You have an Azure Functions microservice app named Appl that is hosted in the Consumption plan. App1 uses an Azure Queue Storage trigger.

You plan to migrate App1 to an Azure Kubernetes Service (AKS) cluster.

You need to prepare the AKS cluster to support Appl. The solution must meet the following requirements:

• Use the same scaling mechanism as the current deployment.

• Support kubenet and Azure Container Netwoking Interface (CNI) networking.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Configure the horizontal pod autoscaler.

B. Install Virtual Kubelet.

C. Configure the AKS cluster autoscaler.

D. Install Kubernetes-based Event Driven Autoscaling (KEDA).

ANSWER: C D

Explanation:

Install Kubernetes-based Event Driven Autoscaling (KEDA) to retain the event-driven scaling behavior used by an Azure Functions app running in the Consumption plan. KEDA monitors the configured event source and adjusts the number of workload replicas in response to demand. For an Azure Queue Storage-triggered workload, the Azure Queue scaler can use queue length to scale the application, including scaling replicas down to zero when no messages require processing. KEDA is designed for Kubernetes event-driven workloads and supports AKS clusters that use either kubenet or Azure CNI networking. See KEDA on AKS and the KEDA Azure Queue Storage scaler documentation.

Configure the AKS cluster autoscaler as the complementary infrastructure-scale component. KEDA increases or decreases the number of application pods according to queue activity, while the cluster autoscaler adjusts the AKS node count when the scheduled pod capacity requires more or fewer nodes. Together, these mechanisms provide event-driven application scaling and the cluster capacity needed to host the resulting replicas. The AKS cluster autoscaler can be enabled for node pools and automatically adds nodes for unschedulable pods or removes underutilized nodes when appropriate.

QUESTION NO: 13

You need to recommend a notification solution for the IT Support distribution group.

What should you include in the recommendation?

A. Azure Network Watcher

B. an action group

C. a SendGrid account with advanced reporting

D. Azure AD Connect Health

ANSWER: D

Explanation:

Azure AD Connect Health is the appropriate recommendation because it provides health monitoring and alerting for hybrid identity components, including Azure AD Connect synchronization. Its alerting configuration supports adding email recipients, including an IT Support distribution group, so that the operational team receives notifications when a monitored issue requires attention. This directly supports a centralized support-notification process for the on-premises Active Directory and hybrid identity environment. Azure AD Connect Health presents alerts in the Microsoft Entra admin center and can send alert notifications to the configured recipients. Administrators can configure the notification email addresses from the

service's health and alerts settings, allowing distribution lists to be used rather than requiring individual support staff to be configured separately. The service also supplies actionable health insights for synchronization and related identity infrastructure, helping the support team identify and remediate issues promptly. Microsoft documents Azure AD Connect Health monitoring and alert operations at Microsoft Entra Connect Health operations. For an overview of the health monitoring capabilities available for hybrid identity, see What is Microsoft Entra Connect Health?.

QUESTION NO: 14

You use Azure virtual machines to run a custom application that uses an Azure SQL database on the back end.

The IT apartment at your company recently enabled forced tunneling,

Since the configuration change, developers have noticed degraded performance when they access the database

You need to recommend a solution to minimize latency when accessing the database. The solution must minimize costs

What should you include in the recommendation?

A. Azure SQL Database Managed instance

B. Azure virtual machines that run Microsoft SQL Server servers

C. Always On availability groups

D. virtual network (VNET) service endpoint

ANSWER: D

Explanation:

A virtual network (VNET) service endpoint is the appropriate recommendation because it provides an optimized route from the virtual machines' subnet to the Azure SQL Database service. Forced tunneling sends Internet-bound traffic to an on-premises location or network virtual appliance through a user-defined route. Without a service endpoint, connections to the public endpoint for Azure SQL Database can follow that forced-tunnel path, adding network hops and latency.

Enabling the `Microsoft.Sql` service endpoint on the subnet installs service-specific routing that directs traffic to Azure SQL Database over the Microsoft backbone network. This route takes precedence over the forced-tunneling default route for the Azure SQL service address prefixes, avoiding the unnecessary on-premises traversal while retaining forced tunneling for other applicable traffic. The SQL database firewall can then be configured to allow access from the selected virtual network subnet, further limiting access to the service endpoint traffic.

Virtual network service endpoints do not have an additional charge, so this approach minimizes cost while addressing the latency introduced by forced tunneling. See Virtual Network service endpoints and Azure user-defined routes and forced tunneling.

QUESTION NO: 15

You are designing a solution that calculates 3D geometry from height-map data. You need to recommend a solution that meets the following requirements:

• Performs calculations in Azure.

• Ensures that each node can communicate data to every other node.

• Maximizes the number of nodes to calculate multiple scenes as fast as possible.

• Minimizes the amount of effort to implement the solution.

Which two actions should you include in the recommendation? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create a render farm that uses virtual machine scale sets.

B. Enable parallel file systems on Azure.

C. Create a render farm that uses virtual machines.

D. Enable parallel task execution on compute nodes.

E. Create a render farm that uses Azure Batch.

ANSWER: D E

Explanation:

Enable parallel task execution on compute nodes and **Create a render farm that uses Azure Batch** satisfy the requirements together. Azure Batch is a managed service for running large-scale parallel and high-performance computing workloads. It provisions and manages pools of compute nodes, schedules jobs and tasks, and can scale pool capacity to provide a large number of workers for rendering or calculating multiple scenes concurrently. This avoids the operational work of manually building, managing, and scheduling a fleet of individual virtual machines.

Parallel task execution in Azure Batch supports workloads whose participating tasks or nodes must communicate during execution. Batch provides multi-instance task capabilities, including coordination across the allocated compute nodes, and supports task communication patterns needed by tightly coupled distributed processing. This enables geometry-calculation workloads to distribute work among nodes while allowing the nodes to exchange data as required.

Using Azure Batch with parallel execution therefore combines managed elastic compute capacity with coordinated multi-node processing, maximizing throughput while minimizing implementation and infrastructure-management effort. See Azure Batch service overview and Use multi-instance tasks to run MPI applications in Azure Batch.

QUESTION NO: 16

A company named Contoso, Ltd. has an Azure Active Directory (Azure AD) tenant that is integrated with Microsoft Office 365 and an Azure subscription.

Contoso has an on-premises identity infrastructure. The infrastructure includes servers that run Active Directory Domain Services (AD DS), and Azure AD Connect

Contoso has a partnership with a company named Fabrikam, Inc. Fabrikam has an Active Directory forest and an Office 365 tenant. Fabrikam has the same on-premises identity infrastructure as Contoso.

A team of 10 developers from Fabrikam will work on an Azure solution that will be hosted in the Azure subscription of Contoso. The developers must be added to the Contributor role for a resource in the Contoso subscription.

You need to recommend a solution to ensure that Contoso can assign the role to the 10 Fabrikam developers. The solution must ensure that the Fabrikam developers use their existing credentials to access resources.

What should you recommend?

A. Configure a forest trust between the on-premises Active Directory forests of Contoso and Fabrikam.

B. Configure an organization relationship between the Office 365 tenants of Fabrikam and Contoso.

C. In the Azure AD tenant of Contoso, use MIM to create guest accounts for the Fabrikam developers.

D. Configure an AD FS relying party trust between the fabrikam and Contoso AD FS infrastructures.

ANSWER: C

Explanation:

In the Azure AD tenant of Contoso, use MIM to create guest accounts for the Fabrikam developers. This approach uses Azure AD B2B collaboration, now called Microsoft Entra B2B collaboration, to represent each Fabrikam developer in Contoso's tenant as a guest user. The guest object is a security principal in Contoso's directory, which makes it available for

Azure role-based access control assignments at the required scope, such as an individual resource. Contoso can therefore assign the Contributor role to each guest account without creating separate Contoso-managed identities for the developers.

Guest users authenticate through their home organization, Fabrikam, using the identities and credentials that they already use. This meets the requirement to preserve the developers' existing sign-in experience while allowing Contoso to control authorization to its Azure resources. Microsoft Identity Manager can be used to automate the creation and lifecycle management of these guest identities where appropriate. After the users accept the B2B invitation and are present in the Contoso tenant, Azure RBAC evaluates their assigned Contributor permissions when they access the resource. See Microsoft Entra B2B collaboration overview and Assign Azure roles using the Azure portal.

QUESTION NO: 17 - (DRAG DROP)

DRAG DROP

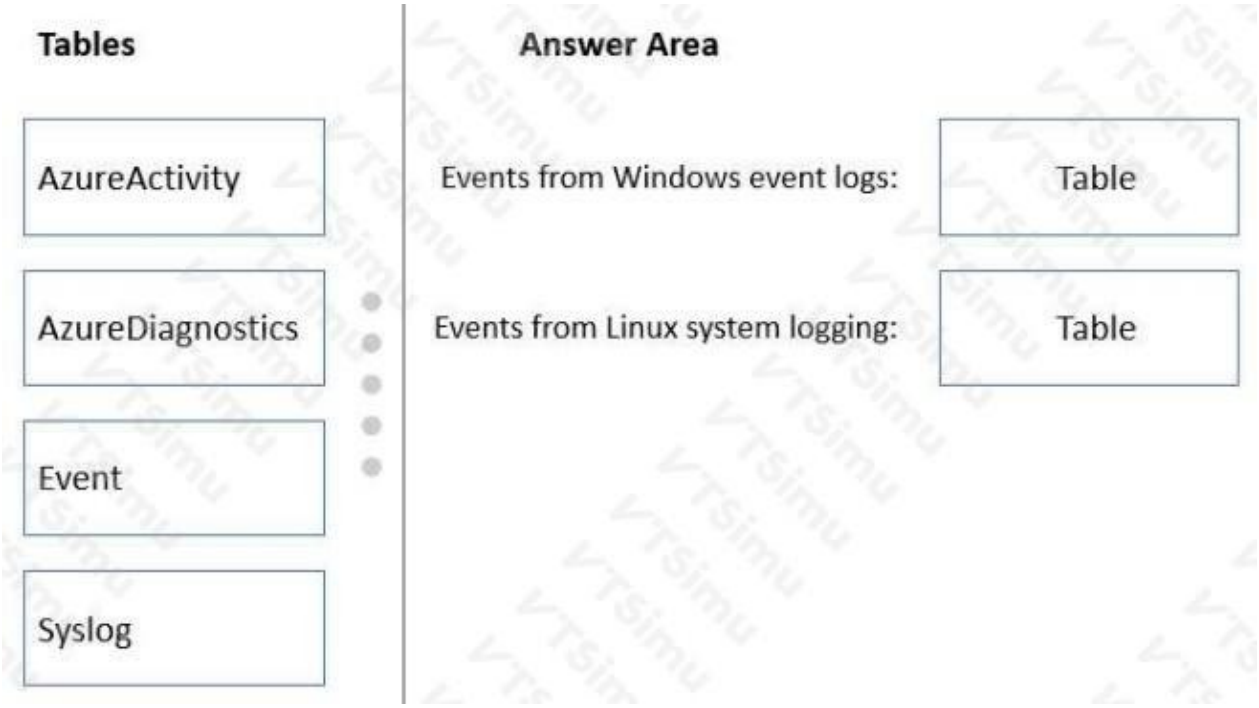
You have an Azure subscription. The subscription contains Azure virtual machines that run Windows Server 2016 and Linux.

You need to use Azure Monitor to design an alerting strategy for security-related events.

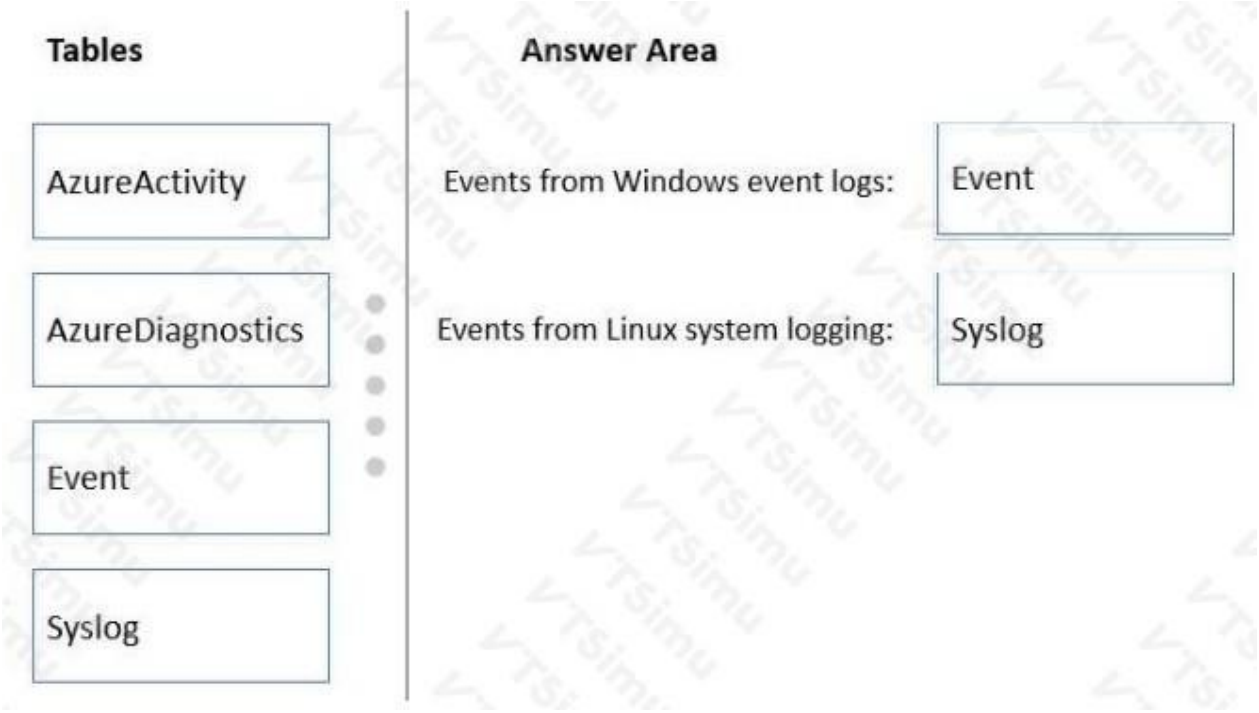
Which Azure Monitor Logs tables should you query? To answer, drag the appropriate tables to the correct log types. Each table may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:



ANSWER:



Explanation:

For security-related alerting from these virtual machines, Azure Monitor Logs should be queried through the tables that store the operating system event streams collected from each platform. Windows Server event records are written to the **Event** table. This table contains Windows Event Log entries collected from monitored computers, including the event log name, event ID, provider, level, computer name, and the event message. Those fields allow alert rules to identify important security signals, such as authentication activity, account-management events, audit-policy events, or other Windows security events selected for collection. Azure Monitor log alert rules can run Kusto queries against this table on a schedule and trigger an action when the query produces matching results. Microsoft documents the schema and purpose of this table in the Event table reference.

Linux system logging records are written to the **Syslog** table. Linux hosts use Syslog facilities and severities to categorize operating system and service messages, and the table preserves useful alerting fields such as the computer, process, facility, severity level, and message text. This supports security monitoring queries for events emitted by Linux authentication and system logging components, including messages associated with sign-in attempts, privilege-related activity, and system service events. A scheduled log search alert can filter the collected Syslog records by relevant facility, severity, host, process, or message content and notify the security operations team when defined conditions occur. The table's fields and collection model are described in Microsoft's Syslog table reference. Therefore, the correct design maps Windows event-log events to Event and Linux system-logging events to Syslog.

QUESTION NO: 18

You have an Azure subscription that contains a Windows Virtual Desktop tenant.

You need to recommend a solution to meet the following requirements:

What should you include in the recommendation?

A. Microsoft Intune

B. a Windows Virtual Desktop automation task

C. Azure Automation

D. Azure Service Health

ANSWER: C

Explanation:

Azure Automation is the appropriate recommendation when the solution requires automated operational actions for Windows Virtual Desktop, now called Azure Virtual Desktop. Azure Automation can host PowerShell runbooks that authenticate to Azure and perform repeatable management tasks against session-host virtual machines. For example, a runbook can start session hosts before expected business hours, stop or deallocate unused hosts after hours, and apply a schedule so that the process runs without an administrator manually operating each VM. This approach supports cost management because deallocated session-host VMs do not continue to incur compute charges, while still allowing capacity to be prepared for users when required.

Automation accounts provide scheduling, managed identities, job execution history, logging, and error handling capabilities that make the automation maintainable at scale. Runbooks can use Azure Resource Manager cmdlets or REST APIs to assess and manage virtual-machine state, and schedules can invoke them at defined times. In modern Azure Virtual Desktop designs, native scaling plans are often preferred where their capabilities meet the requirement; however, Azure Automation remains the applicable solution for the automated runbook-based management scenario indicated here. See Azure Automation overview and Azure Virtual Desktop autoscale and scaling plans.

QUESTION NO: 19

You plan to use an Azure Storage account to store data assets.

You need to recommend a solution that meets the following requirements:

• Supports immutable storage

• Disables anonymous access to the storage account

• Supports access control list (ACL)-based Azure AD permissions

What should you include in the recommendation?

A. Azure Blob Storage

B. Azure Data Lake Storage

C. Azure NetApp Files

D. Azure Files

ANSWER: B

Explanation:

Azure Data Lake Storage is the appropriate recommendation because it combines Blob storage durability and governance capabilities with a hierarchical namespace designed for analytics-oriented data lakes. When hierarchical namespace is enabled on a storage account, Azure Data Lake Storage Gen2 supports fine-grained, POSIX-like access control lists on directories and files. These ACLs can grant permissions to Microsoft Entra users, service principals, managed identities, and groups, allowing access to be controlled at the filesystem, directory, and file levels rather than only through broad container-level roles.

The underlying Azure Storage account can also be configured to prevent anonymous blob access by disabling anonymous access at the account level. In addition, Azure Data Lake Storage Gen2 workloads can use Azure Blob immutable storage policies to protect data using write-once, read-many retention controls. A time-based retention policy or legal hold helps preserve regulated or audit-sensitive data from modification or deletion for the required period. This combination satisfies all three requirements in a single storage-account-based design. See Access control in Azure Data Lake Storage Gen2 and Immutable storage for Azure Blob Storage.

QUESTION NO: 20

You plan to deploy an application to an Azure Kubernetes Service (AKS) cluster.

The application images will be stored in an Azure Container Registry named ACR1.

You need to ensure that AKS can pull images from ACR1. The solution must avoid storing registry credentials in Kubernetes secrets.

Which two actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Configure the AKS cluster to use a managed identity.

B. Assign the AcrPull role on ACR1 to the AKS kubelet identity.

C. Enable the admin user for ACR1.

D. Assign the Contributor role on ACR1 to the AKS kubelet identity.

E. Create a Kubernetes secret that contains the ACR1 access key.

ANSWER: A B

Explanation:

AKS can authenticate to Azure Container Registry by using a managed identity instead of an administrative registry account and stored credentials. The AKS kubelet identity is used by the cluster to pull images from Azure Container Registry. Using this identity eliminates the need to create image-pull secrets that contain a registry username and password.

The kubelet identity must be granted the **AcrPull** Azure role on ACR1. This role provides the required pull permission while following least-privilege principles. The AKS cluster must also use managed identity authentication so that Azure can obtain and manage identity tokens for the cluster. Enabling the ACR admin user would expose long-lived registry credentials, and the Contributor role grants permissions beyond those needed to pull images. See Authenticate with Azure Container Registry from Azure Kubernetes Service and Azure Container Registry roles and permissions.

QUESTION NO: 21

You have an Azure subscription.

Your on-premises network contains a file server named Server1. Server 1 stores 5 TB of company files that are accessed rarely.

You plan to copy the files to Azure Storage.

You need to implement a storage solution for the files that meets the following requirements:

• The files must be available within 24 hours of being requested.

• Storage costs must be minimized.

Which two possible storage solutions achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Create a general-purpose v1 storage account. Create a blob container and copy the files to the blob container.

B. Create a general-purpose v2 storage account that is configured for the Hot default access tier. Create a blob container, copy the files to the blob container, and set each file to the Archive access tier.

C. Create a general-purpose v1 storage account. Create a file share in the storage account and copy the files to the file share.

D. Create a general-purpose v2 storage account that is configured for the Cool default access tier. Create a file share in the storage account and copy the files to the file share.

E. Create an Azure Blob storage account that is configured for the Cool default access tier. Create a blob container, copy the files to the blob container, and set each file to the Archive access tier.

ANSWER: B E

Explanation:

Both solutions store the data as block blobs and explicitly assign every blob to the Archive access tier. Archive is Azure Blob Storage's lowest-cost online storage tier for data that is rarely accessed, making it appropriate for 5 TB of infrequently requested company files. Archive-tier blobs remain stored in Azure but are offline and must be rehydrated before they can be read. Microsoft documents that standard-priority rehydration can take up to 15 hours, which satisfies the requirement for availability within 24 hours. Where a faster response is needed, high-priority rehydration is also available.

A general-purpose v2 storage account supports blob access tiers, including Archive, and its account-level Hot setting does not prevent individual blobs from being moved to Archive. Likewise, an Azure Blob storage account can contain blobs in the Archive tier; setting the account's default tier to Cool affects newly uploaded blobs before the explicit per-blob Archive assignment. In each case, the explicit Archive tier is the key configuration that minimizes capacity cost while retaining a documented rehydration path within the stated time requirement. See Azure Blob Storage access tiers and Rehydrate blobs from the archive tier.

QUESTION NO: 22

You are designing a configuration-management solution for several Azure-hosted applications.

The solution must meet the following requirements:

• Store application settings centrally.

• Allow configuration values to be updated without redeploying the applications.

• Store passwords and connection strings separately from nonsecret configuration values.

• Restrict access to secret values by using Microsoft Entra ID.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure App Configuration

B. Azure Key Vault

C. Azure Automation

D. Azure Cache for Redis

E. Azure Service Bus

ANSWER: A B

Explanation:

Azure App Configuration provides a centralized service for application settings and feature flags. Applications can load configuration dynamically, allowing settings to be changed independently from application deployments. App Configuration is intended for nonsecret configuration values such as endpoint names, application behavior settings, and feature flags.

Azure Key Vault is appropriate for passwords, connection strings, certificates, and other secrets. Key Vault integrates with Microsoft Entra ID and Azure role-based access control, allowing applications to retrieve secret values by using managed identities. App Configuration can also store Key Vault references, enabling applications to retrieve configuration from App Configuration while keeping secret values in Key Vault.

See Azure App Configuration overview and Azure Key Vault overview.

QUESTION NO: 23

You architect a solution that calculates 3D geometry from height-map data.

You have the following requirements:

Perform calculations in Azure.

Each node must communicate data to every other node.

Maximize the number of nodes to calculate multiple scenes as fast as possible.

Require the least amount of effort to implement.

You need to recommend a solution.

Which two actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create a render farm that uses Azure Batch.

B. Enable parallel file systems on Azure.

C. Enable parallel task execution on compute nodes.

D. Create a render farm that uses virtual machine (VM) scale sets.

E. Create a render farm that uses virtual machines (VMs).

ANSWER: A C

Explanation:

Creating a render farm that uses Azure Batch is appropriate because Azure Batch is a managed service for running large-scale parallel and high-performance computing workloads. It provisions and manages compute pools, schedules jobs and tasks across many nodes, and can scale the pool to support a large number of scene calculations without requiring the solution architect to build the underlying orchestration platform. This directly reduces implementation and operational effort while providing the compute scale needed for rendering or geometry-processing workloads.

Enabling parallel task execution on compute nodes allows the workload to use Azure Batch capabilities for parallel and tightly coupled computation. In particular, Batch supports multi-instance tasks, which coordinate a group of compute nodes to run an application such as an MPI-based workload. MPI enables processes on participating nodes to exchange data during processing, satisfying the requirement for inter-node communication. The task is scheduled as one coordinated unit while Batch handles node allocation, task execution, and lifecycle management. Together, Azure Batch and parallel execution provide a managed, scalable approach for processing many scenes quickly while supporting communication across the nodes involved in an individual calculation.

For implementation guidance, see Azure Batch technical overview and Use multi-instance tasks to run MPI applications in Azure Batch.

QUESTION NO: 24

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company has deployed several virtual machines (VMs) on-premises and to Azure. Azure ExpressRoute has been deployed and configured for on-premises to Azure connectivity.

Several VMs are exhibiting network connectivity issues.

You need to analyze the network traffic to determine whether packets are being allowed or denied to the VMs.

Solution: Install and configure the Microsoft Monitoring Agent and the Dependency Agent on all VMs. Use the Wire Data solution in Azure Monitor to analyze the network traffic.

Does the solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because the stated solution does not provide the required determination of whether a specific packet flow is allowed or denied for a VM. The requirement is a network policy and routing diagnostic: an administrator must be able to test a flow defined by its direction, protocol, local and remote IP addresses, and ports, then identify the effective decision that applies to that flow. Azure Network Watcher IP flow verify is designed for this purpose. It evaluates the applicable network security group rules for an Azure VM network interface and reports whether the tested flow is allowed or denied, including the rule responsible for the result. This gives an actionable result for troubleshooting VM connectivity in an Azure environment connected through ExpressRoute.

The proposed monitoring approach collects traffic and dependency telemetry, but it does not satisfy the specific allow-or-deny packet-flow verification goal. Microsoft documents IP flow verify as a Network Watcher capability for checking whether a packet is permitted or blocked to or from a virtual machine. See IP flow verify overview and Azure Network Watcher monitoring and diagnostics.

QUESTION NO: 25

You have an Azure subscription that contains an Azure Blob storage account named store1.

You have an on-premises file server named Setver1 that runs Windows Sewer 2016. Server1 stores 500 GB of company files.

You need to store a copy of the company files from Server 1 in store1.

Which two possible Azure services achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point

A. an Azure Batch account

B. an integration account

C. an On-premises data gateway

D. an Azure Import/Export job

E. Azure Data factory

ANSWER: D E

Explanation:

An Azure Import/Export job is a complete solution for transferring the files to the Azure Blob storage account. The service supports an import workflow in which the organization copies data to compatible, encrypted disk drives, ships those drives to an Azure datacenter, and Azure transfers the data into the specified storage account. This approach is especially useful when transferring data over the network is impractical or when a physical data-transfer process is preferred. The import job supports importing files into Azure Blob storage. See Azure Import/Export service documentation.

Azure Data factory is also a complete solution because its Copy activity can copy files from an on-premises file system to Azure Blob storage. To reach Server1, the organization installs and configures a self-hosted integration runtime on-premises or on a machine that can access the file share. A pipeline can then use a File System connector as its source and an Azure Blob Storage connector as its destination. The pipeline can be run once to create the required copy and can also be

scheduled for recurring copies if needed. This provides a managed, network-based transfer path for the 500 GB of company files. See File System connector in Azure Data Factory.

QUESTION NO: 26 - (HOTSPOT)

You have an Azure subscription that is linked to an Azure Active Directory Premium Plan 2 tenant The tenant has multi-factor authentication (MFA) enabled for all users.

You have the named locations shown in the following table.

Name	IP address range	Trusted
NY	192.168.2.0/27	Yes
DC	192.168.1.0/27	No
LA	192.168.3.0/27	No

You have the users shown in the following table.

Name	Device operating system	User-risk level	Matching compliance policies
User1	Windows 10	High	None
User2	Windows 10	Medium	None
User3	macOS	Low	None

You plan to deploy the Conditional Access policies shown in the following table.

Name	Assignment	Conditions: Locations	Conditions: User risk	Conditions: Sign-in risk	Access Control: Grant
CA1	All users	Trusted locations	High, Medium	None	Block access
CA2	All users	NY	None	High, Medium	Block access
CA3	All users	LA	None	None	Grant access: Require device to marked as compliant

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
To ensure that the conditions in CA1 can be evaluated, you must enforce an Azure Active Directory (Azure AD) Identity Protection user risk policy .	○	○
To ensure that the conditions in CA2 can be evaluated, you must enforce an Azure Active Directory (Azure AD) Identity Protection sign-in risk policy.	○	○
To ensure that the conditions in CA3 can be evaluated, you must deploy Microsoft Endpoint Manager.	○	○

ANSWER:

Statements	Yes	No
To ensure that the conditions in CA1 can be evaluated, you must enforce an Azure Active Directory (Azure AD) Identity Protection user risk policy .	○	○
To ensure that the conditions in CA2 can be evaluated, you must enforce an Azure Active Directory (Azure AD) Identity Protection sign-in risk policy.	○	○
To ensure that the conditions in CA3 can be evaluated, you must deploy Microsoft Endpoint Manager.	○	○

Explanation:

The correct selections are No for the statement concerning enforcement of an Identity Protection user risk policy, No for the statement concerning enforcement of an Identity Protection sign-in risk policy, and Yes for the statement concerning deployment of Microsoft Endpoint Manager. Azure AD Premium P2, now Microsoft Entra ID P2, provides the Identity Protection risk detections and the Conditional Access risk conditions. Conditional Access evaluates the applicable user-risk or sign-in-risk level as part of processing a policy assignment. The separately configured Identity Protection user-risk and sign-in-risk policies are remediation policies: they can require actions such as password changes or MFA when their configured risk threshold is reached, but they are not prerequisites for Conditional Access to read and evaluate risk signals. Microsoft documents the relationship between these risk signals and Conditional Access in What are risk detections? and Conditional Access conditions.

Where a Conditional Access policy needs device-compliance information, that information must be produced by an endpoint management service. Microsoft Intune, administered through the Microsoft Intune admin center and commonly referred to in this context as Microsoft Endpoint Manager, enrolls devices, evaluates compliance policies, and reports the compliant-device state to Microsoft Entra ID. Conditional Access can then use that state while deciding whether to grant access. This integration is described in Conditional Access grant controls and Device compliance policies in Intune. Therefore, deploying the endpoint-management capability is necessary for the device-compliance condition involved in the third policy to be evaluated.

QUESTION NO: 27

You plan to migrate 200 VMware virtual machines from an on-premises datacenter to Azure.

You need to use Azure Migrate to discover the servers, assess their Azure readiness and costs, and migrate the servers to Azure virtual machines.

Which two Azure Migrate tools should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Migrate: Discovery and assessment

B. Azure Migrate: Server Migration

C. Azure Database Migration Service

D. Azure Data Box

E. Azure Site Recovery

ANSWER: A B

Explanation:

Azure Migrate: Discovery and assessment discovers on-premises servers and collects configuration and performance data for assessment. Assessments can evaluate Azure readiness, recommend Azure VM sizes, estimate costs, and identify dependencies when the applicable discovery configuration is used.

Azure Migrate: Server Migration performs the migration of supported servers to Azure virtual machines. For VMware workloads, Server Migration can replicate virtual machines to Azure and perform test migrations and production cutovers. Using Discovery and assessment before Server Migration helps ensure that the migration design is based on measured workload requirements and readiness information.

See Azure Migrate services overview and Azure Migrate: Server Migration.

QUESTION NO: 28 - (DRAG DROP)

Your company identifies the following business continuity and disaster recovery objectives for virtual machines that host sales, finance, and reporting application in the company's on-premises data center.

•The finance application requires that data be retained for seven years. In the event of a disaster, the application must be able to run from Azure. The recovery in objective (RTO) is 10 minutes,

• The reporting application must be able to recover point in-time data al a daily granularity. The RTO is eight hours.

•The sales application must be able to fail over to second on-premises data center.

You need to recommend which Azure services meet the business community and disaster recovery objectives. The solution must minimize costs.

What should you recommend for each application? To answer, drag the appropriate services to the correct application. Each service may be used owe. More than once not at an You may need to drag the spin bar between panes or scroll 10 view content.



ANSWER:



Explanation:

The recommended mapping uses Azure Site Recovery for the workloads that require disaster-recovery failover and Azure Backup for the workload that requires recoverable historical data. For the sales application, Azure Site Recovery provides replication of on-premises virtual machines to a secondary on-premises datacenter and supports orchestrated failover. This directly supports the requirement to continue the workload from the second datacenter when the primary site is unavailable. Microsoft documents this scenario under VMware/physical-machine disaster recovery and failback with Azure Site Recovery: Set up disaster recovery to a secondary on-premises site.

The finance application needs two complementary capabilities. Azure Site Recovery replicates its virtual machines to Azure so that the application can be brought online in Azure during a datacenter disaster. Its replication, recovery-plan, and failover capabilities support the short recovery-time objective. Azure Backup provides the long-term retention required for the finance data, including retention policies that can retain recovery points for years. Together, the services provide both operational disaster recovery and seven-year data retention. Details on Azure Backup retention are available at Back up Azure virtual machines, while Site Recovery failover concepts are described at Run a failover in Azure Site Recovery.

The reporting application requires the ability to restore data to a daily point in time and has an eight-hour recovery-time objective. Azure Backup is appropriate because backup policies create scheduled recovery points and enable restoration from a selected retained recovery point. A daily backup schedule supplies the required daily recovery granularity, and the relatively long recovery window permits recovery through the backup-and-restore process. This is the lowest-cost service combination that meets that stated recovery need because it uses protected recovery points rather than continuous disaster-recovery replication. Azure Backup policy scheduling and retention are documented at Back up Azure VMs with a backup policy.

QUESTION NO: 29

You have an Azure Active Directory (Azure AD) tenant that syncs with an on-premises Active Directory domain.

Your company has a line-of-business (LOB) application that was developed internally.

You need to implement. SAML single sign-on (SSO) and enforce multi-factor authentication (MFA) when users attempt to access the application from an unknown location.

Which two features should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Azure AD enterprise applications

B. Azure AD Identity Protection

C. Azure Application Gateway

D. Conditional Access policies

E. Azure AD Privileged Identity Management (PIM)

ANSWER: A D

Explanation:

Azure AD enterprise applications and Conditional Access policies provide the required application integration and sign-in control. An internally developed LOB application can be registered and configured as an enterprise application in Microsoft Entra ID (formerly Azure AD). The enterprise application represents the service principal in the tenant and supports configuring single sign-on using SAML. SAML settings establish the trust relationship between Microsoft Entra ID and the application, including identifiers, reply URLs, claims, and signing certificates, allowing users to authenticate through the tenant rather than directly to the application.

Conditional Access policies can then target that enterprise application and apply access requirements according to sign-in context. A policy can use named locations to identify trusted corporate network ranges and include all other locations, or explicitly exclude the trusted locations. The policy's grant controls can require multifactor authentication, so users signing in from an unknown or untrusted location must complete MFA before access is granted. This produces the requested location-aware MFA enforcement without changing the LOB application's authentication logic. For implementation details, see Configure SAML-based single sign-on for an enterprise application and Conditional Access location conditions.

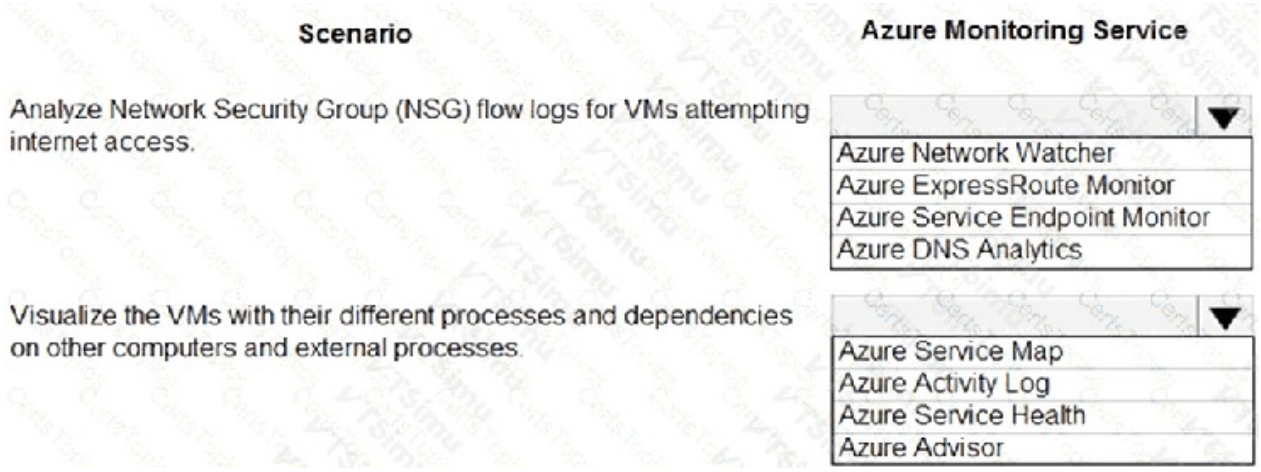
QUESTION NO: 30 - (SIMULATION)

Your company deploys several Linux and Windows virtual machines (VMs) to Azure. The VMs are deployed with the Microsoft Dependency Agent and the Microsoft Monitoring Agent installed by using Azure VM extensions. On-premises connectivity has been enabled by using Azure ExpressRoute.

You need to design a solution to monitor the VMs.

Which Azure monitoring services should you use? To answer, select the appropriate Azure monitoring services in the answer area.

NOTE: Each correct selection is worth one point.



ANSWER: See the explanation for the answer

Explanation:

Select the following Azure monitoring services:

Network Watcher provides Traffic Analytics for NSG flow logs, enabling analysis of VM traffic such as attempted Internet connections. Service Map uses the installed Microsoft Monitoring Agent and Dependency Agent to discover processes and visualize TCP dependencies and connections between Windows/Linux machines and external systems.

QUESTION NO: 31

You have several Azure virtual machines that run a line-of-business application.

You need to notify an IT Support distribution group when the average Percentage CPU metric for any virtual machine exceeds 80 percent for 15 minutes.

Which two Azure Monitor components should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. an Azure Monitor metric alert rule

B. an Azure Monitor action group

C. an Azure Activity Log alert rule

D. Azure Advisor

E. an Azure Service Health alert

ANSWER: A B

Explanation:

An Azure Monitor metric alert rule is required to evaluate the Percentage CPU platform metric for the virtual machines. The rule can be configured with an aggregation period of 15 minutes and a threshold of 80 percent. Metric alerts evaluate Azure Monitor metrics and can target one or more virtual machines.

An action group is required to define the notification action. The action group can contain the email address of the IT Support distribution group and can be associated with the metric alert rule. When the alert condition is met, Azure Monitor invokes the action group and sends the notification. Azure Activity Log records subscription-level management operations rather than sustained guest-resource metric conditions, and Azure Advisor provides recommendations rather than threshold-based alert notifications. See Azure Monitor metric alerts and Azure Monitor action groups.

QUESTION NO: 32

You are designing an application that will aggregate content for users.

You need to recommend a database solution for the application. The solution must meet the following requirements:

What should you include in the recommendation?

A. Azure Cosmos DB SQL API

B. Azure SQL Database that uses active geo-replication

C. Azure SQL Database Hyperscale

D. Azure Database for PostgreSQL

ANSWER: A

Explanation:

Azure Cosmos DB SQL API is appropriate for a globally distributed content-aggregation application because it provides a fully managed NoSQL database with flexible JSON document storage and low-latency access at global scale. Content items and user-facing aggregated data can be represented naturally as JSON documents, while the SQL API provides a familiar query language for querying those documents.

Azure Cosmos DB can replicate data to any number of Azure regions and supports multi-region writes, allowing applications to accept writes close to users in multiple locations. Its global distribution capabilities are designed to provide high availability and predictable low-latency reads and writes. The service also supports configurable consistency levels, enabling the architecture to select an appropriate balance among consistency, availability, latency, and throughput for the application's content experience.

Partitioning data by a suitable logical key enables the solution to scale storage and throughput as demand grows. For more detail, see the Azure Cosmos DB global distribution documentation and the Azure Cosmos DB for NoSQL query documentation.

QUESTION NO: 33

You have an Azure Active Directory (Azure AD) tenant that syncs with an on-premises Active Directory domain.

You have an internal web app named WebApp1 that is hosted on-premises. WebApp1 uses Integrated Windows authentication.

Some users work remotely and do NOT have VPN access to the on-premises network.

You need to provide the remote users with single sign-on (SSO) access to WebApp1.

Which two features should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure AD Application Proxy

B. Azure AD Privileged Identity Management (PIM)

C. Conditional Access policies

D. Azure Arc

E. Azure AD enterprise applications

F. Azure Application Gateway

ANSWER: A C

Explanation:

Azure AD Application Proxy is the core feature for securely publishing an on-premises web application to users outside the corporate network without requiring an inbound firewall opening or a VPN. An on-premises Application Proxy connector establishes an outbound connection to the cloud service, and users access the application through an external URL. For an application that uses Integrated Windows authentication, the proxy can be configured for Azure AD preauthentication and Kerberos Constrained Delegation, allowing the user's Azure AD sign-in to be translated into the Kerberos authentication that WebApp1 expects. This provides the required remote SSO experience.

Conditional Access policies complement Application Proxy by applying Azure AD access controls before users reach the published application. For example, the organization can require multifactor authentication, enforce device or location conditions, or block risky sign-ins. Application Proxy applications support Conditional Access because users authenticate through Azure AD before access is passed to the on-premises application. Together, Azure AD Application Proxy provides the remote publishing and authentication path, while Conditional Access policies provide centralized, identity-based protection for that remote access. See Configure single sign-on for Application Proxy and Conditional Access overview.

QUESTION NO: 34

You are designing an application that will aggregate content for users.

You need to recommend a database solution for the application. The solution must meet the following requirements:

Support SQL commands.

Support multi-master writes.

Guarantee low latency read operations.

What should you include in the recommendation?

A. Azure Cosmos DB SQL API

B. Azure SQL Database that uses active geo-replication

C. Azure SQL Database Hyperscale

D. Azure Database for PostgreSQL

ANSWER: A

Explanation:

Azure Cosmos DB SQL API is the appropriate recommendation because it combines a SQL-based query interface with globally distributed, active-active data access. The SQL API stores schema-flexible JSON items and supports SQL query syntax for querying those items, allowing the application to use familiar SQL constructs while benefiting from a horizontally scalable NoSQL database service. Cosmos DB can be configured for multi-region writes, also called multi-write or multi-master operation. With this configuration, every selected Azure region can accept writes, rather than routing all writes to one primary region. This is well suited to an aggregation application whose users or content sources are geographically distributed.

Cosmos DB also replicates data across configured regions and serves reads locally from those regions, which minimizes network distance to users. Its service-level agreement includes low-latency guarantees for reads, with read operations typically guaranteed at less than 10 ms at the 99th percentile when the documented conditions are met. Together, SQL querying, multi-region write capability, and low-latency regional reads directly satisfy the stated design requirements. See Query data in Azure Cosmos DB for NoSQL and Global distribution in Azure Cosmos DB.

QUESTION NO: 35

You have an on-premises network to which you deploy a virtual appliance.

You plan to deploy several Azure virtual machines and connect the on-premises network to Azure by using a Site-to-Site connection.

All network traffic that will be directed from the Azure virtual machines to a specific subnet must flow through the virtual appliance.

You need to recommend solutions to manage network traffic.

Which two options should you recommend? Each correct answer presents a complete solution.

A. Configure Azure Traffic Manager.

B. Implement an Azure virtual network.

C. Implement Azure ExpressRoute.

D. Configure a routing table.

ANSWER: B D

Explanation:

Implementing an Azure virtual network provides the private network boundary in which the virtual machines, the virtual appliance, and the VPN gateway can be deployed and connected. A virtual network supplies the address space and subnets required for the workload and supports connectivity to the on-premises network through a Site-to-Site VPN gateway. It is therefore the essential Azure networking construct for hosting and connecting the resources involved in this design.

Configuring a routing table enables deterministic traffic steering to the virtual appliance. An Azure route table can contain a user-defined route whose destination prefix is the specific subnet and whose next-hop type is *Virtual appliance*. The next-hop IP address is the private IP address of the appliance's relevant network interface. When the route table is associated with the subnet containing the Azure virtual machines, packets destined for that prefix are sent to the appliance rather than following the system route directly. This supports inspection, filtering, or other network functions before traffic continues toward its destination. The appliance must also be configured to forward traffic, as appropriate for the selected appliance and topology.

For route-table and virtual-appliance routing guidance, see Azure virtual network traffic routing and About VPN Gateway.

QUESTION NO: 36

You have an Azure subscription.

You plan to deploy a monitoring solution that will include the following:

• Azure Monitor Network Insights

• Application Insights

• Microsoft Sentinel

• VM insights

The monitoring solution will be managed by a single team.

What is the minimum number of Azure Monitor workspaces required?

A. 1

B. 2

C. 3

D. 4

ANSWER: A

Explanation:

1 is the minimum required because a single Log Analytics workspace can serve as the shared data store and management boundary for all the listed services when they are administered by one team. Microsoft Sentinel is enabled on a Log Analytics workspace and uses that workspace to collect, analyze, and retain its security data. VM insights can send machine monitoring data to the same workspace, allowing the team to query performance, dependency, and health telemetry centrally. Azure Monitor Network Insights also uses a Log Analytics workspace for its network monitoring data and experience. In addition, a workspace-based Application Insights resource can be connected to that same workspace, which centralizes application telemetry alongside the operational and security data.

Using one workspace supports the stated single-team ownership model by providing one location for data access, Kusto queries, retention configuration, and role-based access control. Separate workspaces might be selected for requirements such as isolation, differing retention, regional design, or delegated administration, but none of those requirements are specified here. Therefore, one shared workspace meets all stated monitoring workloads with the minimum number of workspaces. See Microsoft Sentinel onboarding and VM insights overview.

QUESTION NO: 37

You have an Azure virtual network named VNet1 that contains an Azure virtual machine.

You create a private endpoint for an Azure Storage account that contains blobs. The virtual machine must resolve the storage account blob endpoint to the private IP address of the private endpoint.

You need to recommend a private DNS zone for the solution.

Which private DNS zone should you recommend?

A. privatelink.blob.core.windows.net

B. privatelink.file.core.windows.net

C. privatelink.database.windows.net

D. blob.core.windows.net

ANSWER: A

Explanation:

The appropriate private DNS zone is **privatelink.blob.core.windows.net**. Azure Private Link uses this zone for private endpoint DNS records for the Blob service. A private DNS zone group associated with the private endpoint can create the required DNS record automatically, and the zone can be linked to VNet1.

After the private DNS zone is linked to VNet1, a DNS query from the virtual machine for the Blob service endpoint resolves to the private endpoint IP address rather than the public endpoint. This allows the virtual machine to access the storage account over the private endpoint while applications continue to use the standard storage account DNS name.

See Azure private endpoint DNS configuration.

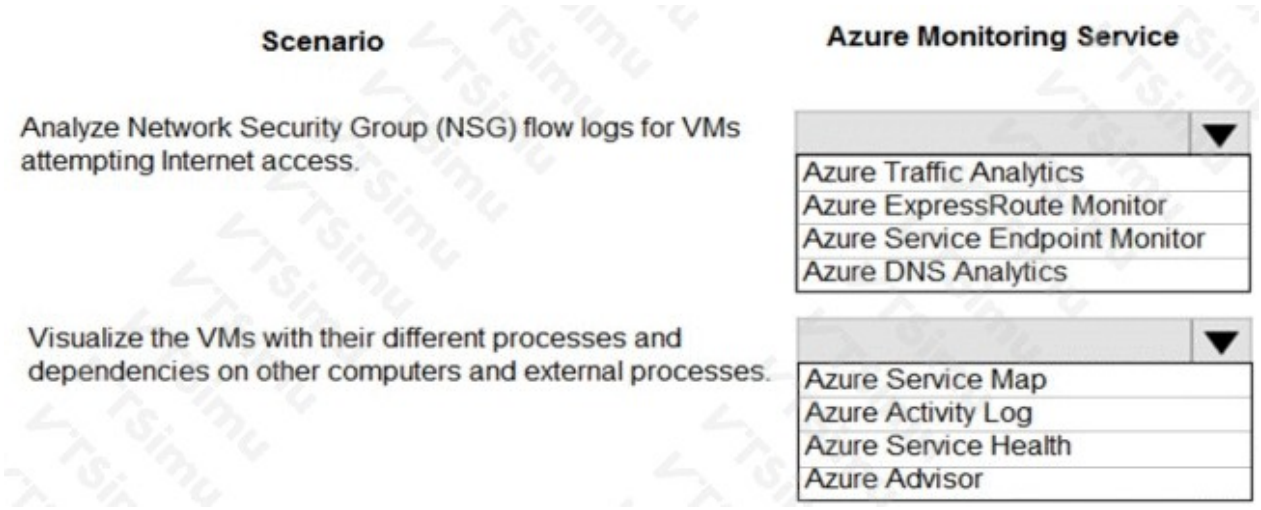
QUESTION NO: 38 - (HOTSPOT)

Your company deploys several Linux and Windows virtual machines (VMs) to Azure. The VMs are deployed with the Microsoft Dependency Agent and the Log Analytics Agent installed by using Azure VM extensions. On-premises connectivity has been enabled by using Azure ExpressRoute.

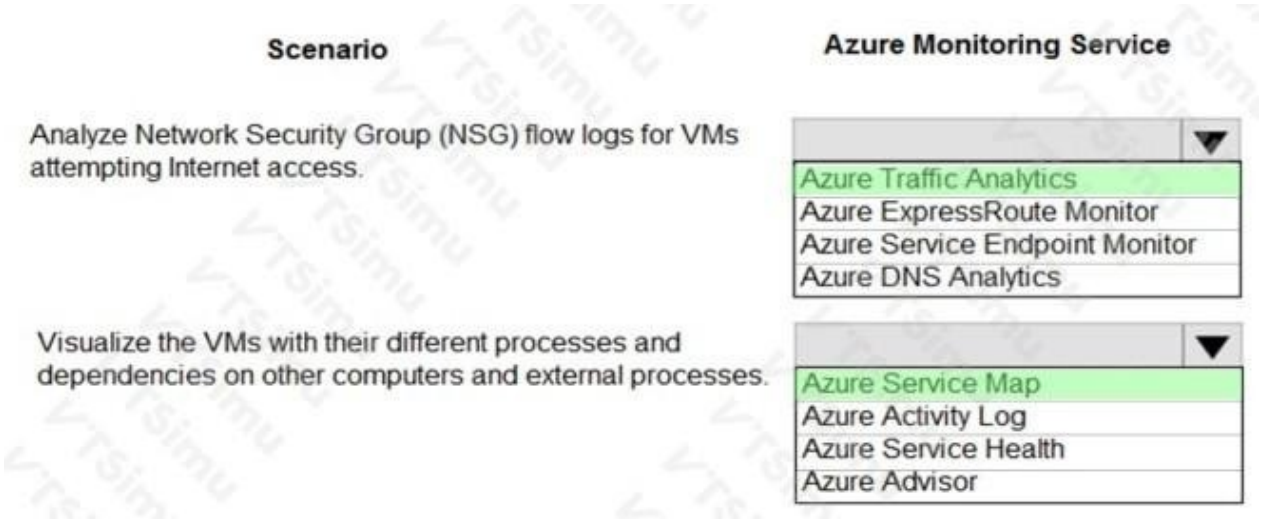
You need to design a solution to monitor the VMs.

Which Azure monitoring services should you use? To answer, select the appropriate Azure monitoring services in the answer area.

NOTE: Each correct selection is worth one point.



ANSWER:



Explanation:

Azure Traffic Analytics is the appropriate service for analyzing Network Security Group flow logs associated with virtual machines that attempt internet access. Traffic Analytics is built on Azure Network Watcher flow logs and provides a centralized view of network communication patterns. It processes the logged flow data and turns it into useful monitoring information, such as source and destination traffic, allowed and denied flows, traffic volume, and communication patterns. This makes it suitable for identifying and investigating traffic that leaves a VM toward internet destinations, while retaining the network-security context supplied by the NSG rules. The service's purpose and data flow are described in the Azure Traffic Analytics documentation.

Azure Service Map is the appropriate service for visualizing the deployed Windows and Linux virtual machines, their running processes, and their dependencies on other systems and external processes. Service Map relies on the Dependency Agent together with Log Analytics data collection to discover TCP communication relationships. It presents those relationships as an application and machine dependency map, allowing architects and operations teams to understand which processes communicate with which remote computers and services. This view is valuable for assessing the effect of infrastructure changes, troubleshooting an application path, and documenting interconnected workloads across Azure and on-premises environments connected through ExpressRoute. The existing agent deployment in the scenario supplies the required discovery capability, so Service Map directly meets the visualization requirement. Microsoft documents this dependency-mapping capability in the VM insights Map documentation.

QUESTION NO: 39

You have an on-premises application named App1 that uses an Oracle database.

You plan to use Azure Databricks to transform and load data from App1 to an Azure Synapse Analytics instance.

You need to ensure that the App1 data is available to Databricks.

Which two Azure services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Data Box Edge

B. Azure Data Lake Storage

C. Azure Data Factory

D. Azure Data Box Gateway

E. Azure Import/Export service

ANSWER: B C

Explanation:

Azure Data Factory and Azure Data Lake Storage provide the appropriate ingestion and storage path for making on-premises Oracle data available to Azure Databricks. Azure Data Factory can connect to Oracle by using its Oracle connector and a self-hosted integration runtime. The self-hosted integration runtime runs in the on-premises environment and enables the service to securely access data sources behind the organization's firewall. A copy activity can then extract the required Oracle data and write it into Azure Data Lake Storage.

Azure Data Lake Storage is a suitable persistent landing zone for this pipeline because Azure Databricks can access data stored in Azure Data Lake Storage for scalable transformation workloads. Databricks can read the staged files from the lake, transform them by using Spark, and load curated results into Azure Synapse Analytics. This separation of ingestion, durable storage, and transformation also supports repeatable scheduled loads rather than a one-time offline transfer. Microsoft documents Oracle as a supported Data Factory connector and describes using a self-hosted integration runtime for on-premises connectivity. Azure Databricks documentation also covers access to ADLS Gen2 storage. See Oracle connector in Azure Data Factory and Azure Databricks access to Azure Storage.