

- **Vendor: Microsoft**
- **Exam Code: AZ-305**
- **Exam Name: Designing Microsoft Azure Infrastructure Solutions**
- **New Updated Questions from [Braindump2go](#)**
- **(Updated in [February/2022](#))**

[Visit Braindump2go and Download Full Version AZ-305 Exam Dumps](#)

**QUESTION 94**

Your company, named Contoso, Ltd., implements several Azure logic apps that have HTTP triggers. The logic apps provide access to an on-premises web service. Contoso establishes a partnership with another company named Fabrikam. Incl Fabrikam does not have an existing Azure Active Directory (Azure AD) tenant and uses third-party OAuth 2.0 identity management to authenticate its users. Developers at Fabrikam plan to use a subset of the logic apps to build applications that will integrate with the on-premises web service of Contoso.

You need to design a solution to provide the Fabrikam developers with access to the logic apps. The solution must meet the following requirements:

- Requests to the logic apps from the developers must be limited to lower rates than the requests from the users at Contoso.
- The developers must be able to rely on their existing OAuth 2.0 provider to gain access to the logic apps.
- The solution must NOT require changes to the logic apps.
- The solution must NOT use Azure AD guest accounts.

What should you include in the solution?

- A. Azure AD business-to-business (B2B)
- B. Azure AD Application Proxy
- C. Azure Front Door
- D. Azure API Management

**Answer: D**

**Explanation:**

API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.

You can secure API Management using the OAuth 2.0 client credentials flow.

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-key-concepts>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-features>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad#enable-oauth-20-user-authorization-in-the-developer-console>

**QUESTION 95**

You have an Azure subscription.

You need to recommend an Azure Kubernetes service (AKS) solution that will use Linux nodes. The solution must meet the following requirements:

- Minimize the time it takes to provision compute resources during scale-out operations.
- Support autoscaling of Linux containers.

[AZ-305 Exam Dumps](#) [AZ-305 Exam Questions](#) [AZ-305 PDF Dumps](#) [AZ-305 VCE Dumps](#)

<https://www.braindump2go.com/az-305.html>

- Minimize administrative effort.  
 Which scaling option should you recommend?

- A. Virtual Kubetet
- B. cluster autoscaler
- C. virtual nodes
- D. horizontal pod autoscaler

**Answer: C**

**Explanation:**

<https://docs.microsoft.com/en-us/azure/aks/virtual-nodes>

**QUESTION 96**

You have an Azure subscription.

You need to deploy an Azure Kubernetes Service (AKS) solution that will use Windows Server 2019 nodes.

The solution must meet the following requirements:

- Minimize the time it takes to provision compute resources during scale-out operations.
- Support autoscaling of Windows Server containers.

Which scaling option should you recommend?

- A. horizontal pod autoscaler
- B. Kubernetes version 1.20.2 or newer
- C. cluster autoscaler
- D. Virtual nodes
- E. with Virtual Kubelet ACI

**Answer: C**

**Explanation:**

<https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler#about-the-cluster-autoscaler>

**QUESTION 97**

You are designing an order processing system in Azure that will contain the Azure resources shown in the following table.

Name	Type	Purpose
App1	Web app	Processes customer orders
Function1	Function	Check product availability at vendor 1
Function2	Function	Check product availability at vendor 2
storage1	Storage account	Stores order processing logs

The order processing system will have the following transaction flow:

- A customer will place an order by using App1.
- When the order is received, App1 will generate a message to check for product availability at vendor 1 and vendor 2.
- An integration component will process the message, and then trigger either Function1 or Function2 depending on the type of order.
- Once a vendor confirms the product availability, a status message for App1 will be generated by Function1 or Function2.

All the steps of the transaction will be logged to storage1.

Which type of resource should you recommend for the integration component?

- A. an Azure Data Factory pipeline
- B. an Azure Service Bus queue
- C. an Azure Event Grid domain
- D. an Azure Event Hubs capture

**Answer:** A

**Explanation:**

A data factory can have one or more pipelines. A pipeline is a logical grouping of activities that together perform a task. The activities in a pipeline define actions to perform on your data.

Data Factory has three groupings of activities: data movement activities, data transformation activities, and control activities.

Azure Functions is now integrated with Azure Data Factory, allowing you to run an Azure function as a step in your data factory pipelines.

**Reference:**

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipelines-activities>

#### **QUESTION 98**

You plan to deploy 10 applications to Azure. The applications will be deployed to two Azure Kubernetes Service (AKS) clusters. Each cluster will be deployed to a separate Azure region. The application deployment must meet the following requirements:

- Ensure that the applications remain available if a single AKS cluster fails.
- Ensure that the connection traffic over the internet is encrypted by using SSL without having to configure SSL on each container.

Which service should you include in the recommendation?

- A. AKS ingress controller
- B. Azure Traffic Manager
- C. Azure Front Door
- D. Azure Load Balancer

**Answer:** C

**Explanation:**

"Azure Front Door, which focuses on global load-balancing and site acceleration, and Azure CDN Standard, which offers static content caching and acceleration. The new Azure Front Door brings together security with CDN technology for a cloud-based CDN with threat protection and additional capabilities. "

#### **QUESTION 99**

You have an Azure subscription.

You need to recommend a solution to provide developers with the ability to provision Azure virtual machines. The solution must meet the following requirements:

- Only allow the creation of the virtual machines in specific regions.
- Only allow the creation of specific sizes of virtual machines.

What should you include in the recommendation?

- A. Conditional Access policies
- B. role-based access control (RBAC)
- C. Azure Resource Manager (ARM) templates
- D. Azure Policy

**Answer:** D

**Explanation:**

<https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage>

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/manage/azure-server-management/common-policies#restrict-vm-size>

#### **QUESTION 100**

You plan to deploy an application named App1 that will run on five Azure virtual machines. Additional virtual machines will be deployed later to run App1. You need to recommend a solution to meet the following requirements for the virtual machines that will run App1:

- Ensure that the virtual machines can authenticate to Azure Active Directory (Azure AD) to gain access to an Azure key vault, Azure Logic Apps instances, and an Azure SQL database.

- Avoid assigning new roles and permissions for Azure services when you deploy additional virtual machines.
- Avoid storing secrets and certificates on the virtual machines.

Which type of identity should you include in the recommendation?

- A. a service principal that is configured to use a certificate
- B. a system-assigned managed identity
- C. a service principal that is configured to use a client secret
- D. a user-assigned managed identity

**Answer: D**

**Explanation:**

Managed identities for Azure resources is a feature of Azure Active Directory. User-assigned managed identity can be shared. The same user-assigned managed identity can be associated with more than one Azure resource.

Incorrect Answers:

B: System-assigned managed identity cannot be shared. It can only be associated with a single Azure resource.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>

### **QUESTION 101**

You have an application that is used by 6,000 users to validate their vacation requests. The application manages its own credential

Users must enter a username and password to access the application. The application does NOT support identity providers.

You plan to upgrade the application to use single sign-on (SSO) authentication by using an Azure Active Directory (Azure AD) application registration.

Which SSO method should you use?

- A. password-based
- B. OpenID Connect
- C. header-based
- D. SAML

**Answer: A**

### **QUESTION 102**

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:**

You can use Copy Activity in Azure Data Factory to copy data from and to Azure Data Lake Storage Gen2, and use Data Flow to transform data in Azure Data Lake Storage Gen2.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-data-lake-storage>

### **QUESTION 103**

You plan to deploy an app that will use an Azure Storage account.

You need to deploy the storage account.

The solution must meet the following requirements:

- Store the data of multiple users.
- Encrypt each user's data by using a separate key.
- Encrypt all the data in the storage account by using Microsoft keys or customer-managed keys.

What should you deploy?

- A. files in a general purpose v2 storage account.
- B. blobs in an Azure Data Lake Storage Gen2 account.
- C. files in a premium file share storage account.
- D. blobs in a general purpose v2 storage account

**Answer: B**

#### **QUESTION 104**

You plan to deploy an Azure SQL database that will store Personally Identifiable Information (PII).

You need to ensure that only privileged users can view the PII.

What should you include in the solution?

- A. Transparent Data Encryption (TDE)
- B. Data Discovery & Classification
- C. dynamic data masking
- D. role-based access control (RBAC)

**Answer: D**

#### **QUESTION 105**

You have the resources shown in the following table.

Name	Type
AS1	Azure Synapse Analytics instance
CDB1	Azure Cosmos DB SQL API account

CDB1 hosts a container that stores continuously updated operational data.

You are designing a solution that will use AS1 to analyze the operational data daily.

You need to recommend a solution to analyze the data without affecting the performance of the operational data store.

What should you include in the recommendation?

- A. Azure Cosmos DB change feed
- B. Azure Data Factory with Azure Cosmos DB and Azure Synapse Analytics connectors
- C. Azure Synapse Analytics with PolyBase data loading
- D. Azure Synapse Link for Azure Cosmos DB

**Answer: C**

#### **QUESTION 106**

You store web access logs data in Azure Blob storage. You plan to generate monthly reports from the access logs. You

need to recommend an automated process to upload the data to Azure SQL Database every month.

What should you include in the recommendation?

- A. Azure Data Factory
- B. Data Migration Assistant
- C. Microsoft SQL Server Migration Assistant (SSMA)
- D. AzCopy

**Answer: A**

**Explanation:**

[AZ-305 Exam Dumps](#) [AZ-305 Exam Questions](#) [AZ-305 PDF Dumps](#) [AZ-305 VCE Dumps](#)

<https://www.braindump2go.com/az-305.html>

Azure Data Factory is the platform that solves such data scenarios. It is the cloud-based ETL and data integration service that allows you to create data-driven workflows for orchestrating data movement and transforming data at scale. Using Azure Data Factory, you can create and schedule data-driven workflows (called pipelines) that can ingest data from disparate data stores. You can build complex ETL processes that transform data visually with data flows or by using compute services such as Azure HDInsight Hadoop, Azure Databricks, and Azure SQL Database.

Reference:

<https://docs.microsoft.com/en-gb/azure/data-factory/introduction>

**QUESTION 107**

Hotspot Question

You have an Azure subscription that contains a virtual network named VNET1 and 10 virtual machines. The virtual machines are connected to VNET1. You need to design a solution to manage the virtual machines from the internet. The solution must meet the following requirements:

- Incoming connections to the virtual machines must be authenticated by using Azure Multi-Factor Authentication (MFA) before network connectivity is allowed.
- Incoming connections must use TLS and connect to TCP port 443.
- The solution must support RDP and SSH.

What should you Include In the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

To provide access to virtual machines on VNET1, use:

Azure Bastion
Just-in-time (JIT) VM access
Azure Web Application Firewall (WAF) in Azure Front Door

To enforce Azure MFA, use:

An Azure Identity Governance access package
A Conditional Access policy that has the Cloud apps assignment set to Azure Windows VM Sign-In
A Conditional Access policy that has the Cloud apps assignment set to Microsoft Azure Management

Answer:

Answer Area

To provide access to virtual machines on VNET1, use:

Azure Bastion
Just-in-time (JIT) VM access
Azure Web Application Firewall (WAF) in Azure Front Door

To enforce Azure MFA, use:

An Azure Identity Governance access package
A Conditional Access policy that has the Cloud apps assignment set to Azure Windows VM Sign-In
A Conditional Access policy that has the Cloud apps assignment set to Microsoft Azure Management

**QUESTION 108**

Hotspot Question

A company plans to implement an HTTP-based API to support a web app. The web app allows customers to check the status of their orders.

The API must meet the following requirements:

- Implement Azure Functions
- Provide public read-only operations
- Do not allow write operations

You need to recommend configuration options.

What should you recommend? To answer, configure the appropriate options in the dialog box in the answer area.

NOTE: Each correct selection is worth one point.

Topic	Value
Allowed authentication methods	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>All methods</p> <p>GET only</p> <p>GET and POST only</p> <p>GET, POST, and OPTIONS only</p> </div> </div>
Authorization level	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Function</p> <p>Anonymous</p> <p>Admin</p> </div> </div>

Answer:

Topic	Value
Allowed authentication methods	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>All methods</p> <p style="background-color: #d9ead3;">GET only</p> <p>GET and POST only</p> <p>GET, POST, and OPTIONS only</p> </div> </div>
Authorization level	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Function</p> <p style="background-color: #d9ead3;">Anonymous</p> <p>Admin</p> </div> </div>

**Explanation:**

Allowed authentication methods: GET only  
 Authorization level: Anonymous

The option is Allow Anonymous requests. This option turns on authentication and authorization in App Service, but defers authorization decisions to your application code. For authenticated requests, App Service also passes along authentication information in the HTTP headers.

This option provides more flexibility in handling anonymous requests.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/overview-authentication-authorization>

**QUESTION 109**

Drag and Drop Question

You plan to import data from your on-premises environment to Azure. The data is shown in the following table.

On-premises source	Azure target
A Microsoft SQL Server 2012 database	An Azure SQL database
A table in a Microsoft SQL Server 2014 database	An Azure Cosmos DB account that uses the SQL API

What should you recommend using to migrate the data? To answer, drag the appropriate tools to the correct data sources-Each tool 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.

**Tools**

AzCopy

Azure Cosmos DB Data Migration Tool

Data Management Gateway

Data Migration Assistant

**Answer Area**

From the SQL Server 2012 database:

From the table in the SQL Server 2014 database:

**Answer:**

**Tools**

AzCopy

Data Management Gateway

**Answer Area**

From the SQL Server 2012

From the table in the SQL Server 2014

**Explanation:**

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-azure-sql> <https://docs.microsoft.com/en-us/azure/cosmos-db/import-data>

**QUESTION 110**

Drag and Drop Question

You have an on-premises network that uses an IP address space of 172.16.0.0/16.

You plan to deploy 25 virtual machines to a new Azure subscription.

You identify the following technical requirements:

- All Azure virtual machines must be placed on the same subnet named Subnet1.
- All the Azure virtual machines must be able to communicate with all on-premises servers.
- The servers must be able to communicate between the on-premises network and Azure by using a site-to-site VPN.

You need to recommend a subnet design that meets the technical requirements.

What should you include in the recommendation? To answer, drag the appropriate network addresses to the correct subnets. Each network address 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.

**Network Addresses**

172.16.0.0/16

172.16.1.0/27

192.168.0.0/24

192.168.1.0/27

**Answer Area**

Subnet1:

Gateway subnet:

**Answer:**

**Network Addresses**

172.16.0.0/16

172.16.1.0/27

**Answer Area**

Subnet1:

Gateway subnet:

**QUESTION 111**

Hotspot Question

You have an Azure web app named App1 and an Azure key vault named KV1.

[AZ-305 Exam Dumps](#) [AZ-305 Exam Questions](#) [AZ-305 PDF Dumps](#) [AZ-305 VCE Dumps](#)

<https://www.braindump2go.com/az-305.html>



App1 stores database connection strings in KV1.  
App1 performs the following types of requests to KV1:

- Get
- List
- Wrap
- Delete
- Unwrap
- Backup
- Decrypt
- Encrypt

You are evaluating the continuity of service for App1.  
You need to identify the following if the Azure region that hosts KV1 becomes unavailable:

- To where will KV1 fail over?
- During the failover, which request type will be unavailable?

What should you identify? To answer, select the appropriate options in the answer area.  
NOTE: Each correct selection is worth one point.

To where will KV1 fail over?

<input type="text"/>
A server in the same Availability Set
A server in the same fault domain
A server in the same paired region
A virtual machine in a scale set

During the failover, which request type will be unavailable?

<input type="text"/>
Backup
Decrypt
Delete
Encrypt
Get
List
Unwrap
Wrap

**Answer:**

To where will KV1 fail over?

A server in the same Availability Set
A server in the same fault domain
A server in the same paired region
A virtual machine in a scale set

During the failover, which request type will be unavailable?

Backup
Decrypt
Delete
Encrypt
Get
List
Unwrap
Wrap

**Explanation:**

Box 1: A server in the same paired region

The contents of your key vault are replicated within the region and to a secondary region at least 150 miles away, but within the same geography to maintain high durability of your keys and secrets.

Box 2: Delete

During failover, your key vault is in read-only mode. Requests that are supported in this mode are:

- List certificates
- Get certificates
- List secrets
- Get secrets
- List keys
- Get (properties of) keys
- Encrypt
- Decrypt
- Wrap
- Unwrap
- Verify
- Sign
- Backup

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/disaster-recovery-guidance>

**QUESTION 112**

Hotspot Question

You have an on-premises file server that stores 2 TB of data files.

You plan to move the data files to Azure Blob Storage in the West Europe Azure region.

You need to recommend a storage account type to store the data files and a replication solution for the storage account.

The solution must meet the following requirements:

- Be available if a single Azure datacenter fails.
- Support storage tiers.
- Minimize cost.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Storage Account type:	<table border="1"> <tr><td>Premium block blobs</td></tr> <tr><td>Standard general-purpose v1</td></tr> <tr><td>Standard general-purpose v2</td></tr> </table>	Premium block blobs	Standard general-purpose v1	Standard general-purpose v2	
Premium block blobs					
Standard general-purpose v1					
Standard general-purpose v2					
Redundancy:	<table border="1"> <tr><td>Geo-redundant storage (GRS)</td></tr> <tr><td>Zone-redundant storage (ZRS)</td></tr> <tr><td>Locally-redundant storage (LRS)</td></tr> <tr><td>Read-access geo-redundant storage (RA-GRS)</td></tr> </table>	Geo-redundant storage (GRS)	Zone-redundant storage (ZRS)	Locally-redundant storage (LRS)	Read-access geo-redundant storage (RA-GRS)
Geo-redundant storage (GRS)					
Zone-redundant storage (ZRS)					
Locally-redundant storage (LRS)					
Read-access geo-redundant storage (RA-GRS)					

**Answer:**

**Answer Area**

Storage Account type:	<table border="1"> <tr><td>Premium block blobs</td></tr> <tr><td>Standard general-purpose v1</td></tr> <tr style="background-color: #e0ffe0;"><td>Standard general-purpose v2</td></tr> </table>	Premium block blobs	Standard general-purpose v1	Standard general-purpose v2	
Premium block blobs					
Standard general-purpose v1					
Standard general-purpose v2					
Redundancy:	<table border="1"> <tr><td>Geo-redundant storage (GRS)</td></tr> <tr style="background-color: #e0ffe0;"><td>Zone-redundant storage (ZRS)</td></tr> <tr><td>Locally-redundant storage (LRS)</td></tr> <tr><td>Read-access geo-redundant storage (RA-GRS)</td></tr> </table>	Geo-redundant storage (GRS)	Zone-redundant storage (ZRS)	Locally-redundant storage (LRS)	Read-access geo-redundant storage (RA-GRS)
Geo-redundant storage (GRS)					
Zone-redundant storage (ZRS)					
Locally-redundant storage (LRS)					
Read-access geo-redundant storage (RA-GRS)					

**Explanation:**

Account Type: StorageV2

Replication solution: Zone-redundant storage (ZRS)

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy#supported-azure-storage-services>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview#types-of-storage-accounts>

Data must be available if a single Azure datacenter fails. It means the storage account must support ZRS replication.

Also, solution should support storage tiers. Only General-purpose V2 supports ZRS and storage tiers.

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**QUESTION 113**

Hotspot Question

Your company has two on-premises sites in New York and Los Angeles and Azure virtual networks in the East US Azure region and the West US Azure region. Each on-premises site has Azure ExpressRoute circuits to both regions.

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

- Outbound traffic to the Internet from workloads hosted on the virtual networks must be routed through the closest available on-premises site.
- If an on-premises site fails, traffic from the workloads on the virtual networks to the Internet must reroute automatically to the other site.

What should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Routing from the virtual networks to the on-premises locations must be configured by using:

▼
Azure default routes
Border Gateway Protocol (BGP)
User-defined routes

The automatic routing configuration following a failover must be handled by using:

▼
Border Gateway Protocol (BGP)
Hot Standby Routing Protocol (HSRP)
Virtual Router Redundancy Protocol (VRRP)

**Answer:**

Routing from the virtual networks to the on-premises locations must be configured by using:

▼
Azure default routes
Border Gateway Protocol (BGP)
User-defined routes

The automatic routing configuration following a failover must be handled by using:

▼
Border Gateway Protocol (BGP)
Hot Standby Routing Protocol (HSRP)
Virtual Router Redundancy Protocol (VRRP)

**Explanation:**

An on-premises network gateway can exchange routes with an Azure virtual network gateway using the border gateway protocol (BGP). Using BGP with an Azure virtual network gateway is dependent on the type you selected when you created the gateway. If the type you selected were: ExpressRoute:

You must use BGP to advertise on-premises routes to the Microsoft Edge router. You cannot create user-defined routes to force traffic to the ExpressRoute virtual network gateway if you deploy a virtual network gateway deployed as type: ExpressRoute. You can use user-defined routes for forcing traffic from the Express Route to, for example, a Network Virtual Appliance.

<https://docs.microsoft.com/ja-jp/azure/expressroute/designing-for-disaster-recovery-with-expressroute-privatepeering>  
<https://docs.microsoft.com/en-us/azure/expressroute/expressroute-optimize-routing#suboptimal-routing-from-customer-to-microsoft>

**QUESTION 114**

Hotspot Question

You deploy several Azure SQL Database instances.

You plan to configure the Diagnostics settings on the databases as shown in the following exhibit.

Diagnostics settings

Save Discard Delete

**\***

Diagnostics

Archive to a storage account

Storage account: csa14d260928e42x4ea7xb77

Stream to an event hub

Send to Log Analytics

Log Analytics: fabrikamproductionworkspace

LOG

Feature	Retention (days)
<input checked="" type="checkbox"/> SQLInsights	90
<input checked="" type="checkbox"/> AutomaticTuning	30
<input type="checkbox"/> QueryStoreRuntimeStatistics	0
<input type="checkbox"/> QueryStoreWaitStatistics	0

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

The amount of time that SQLInsights data will be stored in blob storage is **[answer choice]**.

	▼
30 days	
90 days	
730 days	
indefinite	

The maximum amount of time that SQLInsights data can be stored in Azure Log Analytics is **[answer choice]**.

	▼
30 days	
90 days	
730 days	
indefinite	

**Answer:**

The amount of time that SQLInsights data will be stored in blob storage is **[answer choice]**.

	▼
30 days	
90 days	
730 days	
indefinite	

The maximum amount of time that SQLInsights data can be stored in Azure Log Analytics is **[answer choice]**.

	▼
30 days	
90 days	
730 days	
indefinite	

**Explanation:**

In the exhibit, the SQLInsights data is configured to be stored in Azure Log Analytics for 90 days. However, the question is asking for the "maximum" amount of time that the data can be stored which is 730 days.