

# DUMPSBOSS.COM

**Microsoft Azure Architect Technologies (beta)**

**Microsoft AZ-303**

**Version Demo**

**Total Demo Questions: 20**

**Total Premium Questions: 449**

**Buy Premium PDF**

**<https://dumpsboss.com>**

**[support@dumpsboss.com](mailto:support@dumpsboss.com)**

**dumpsboss.com**

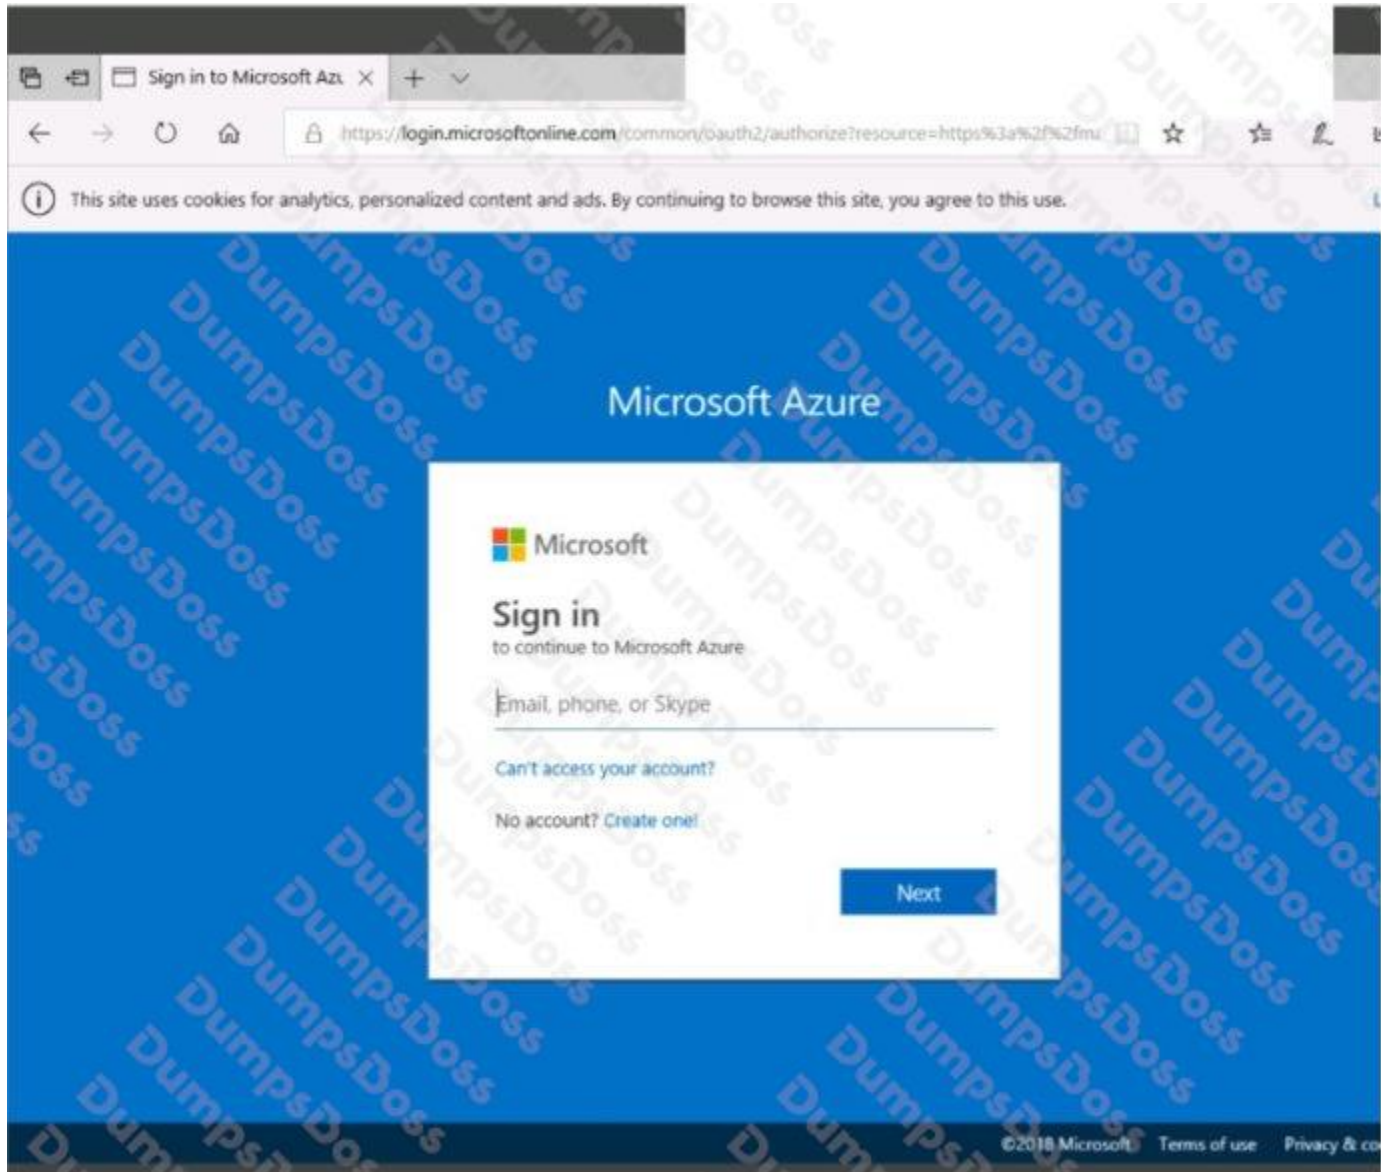
## Topic Break Down

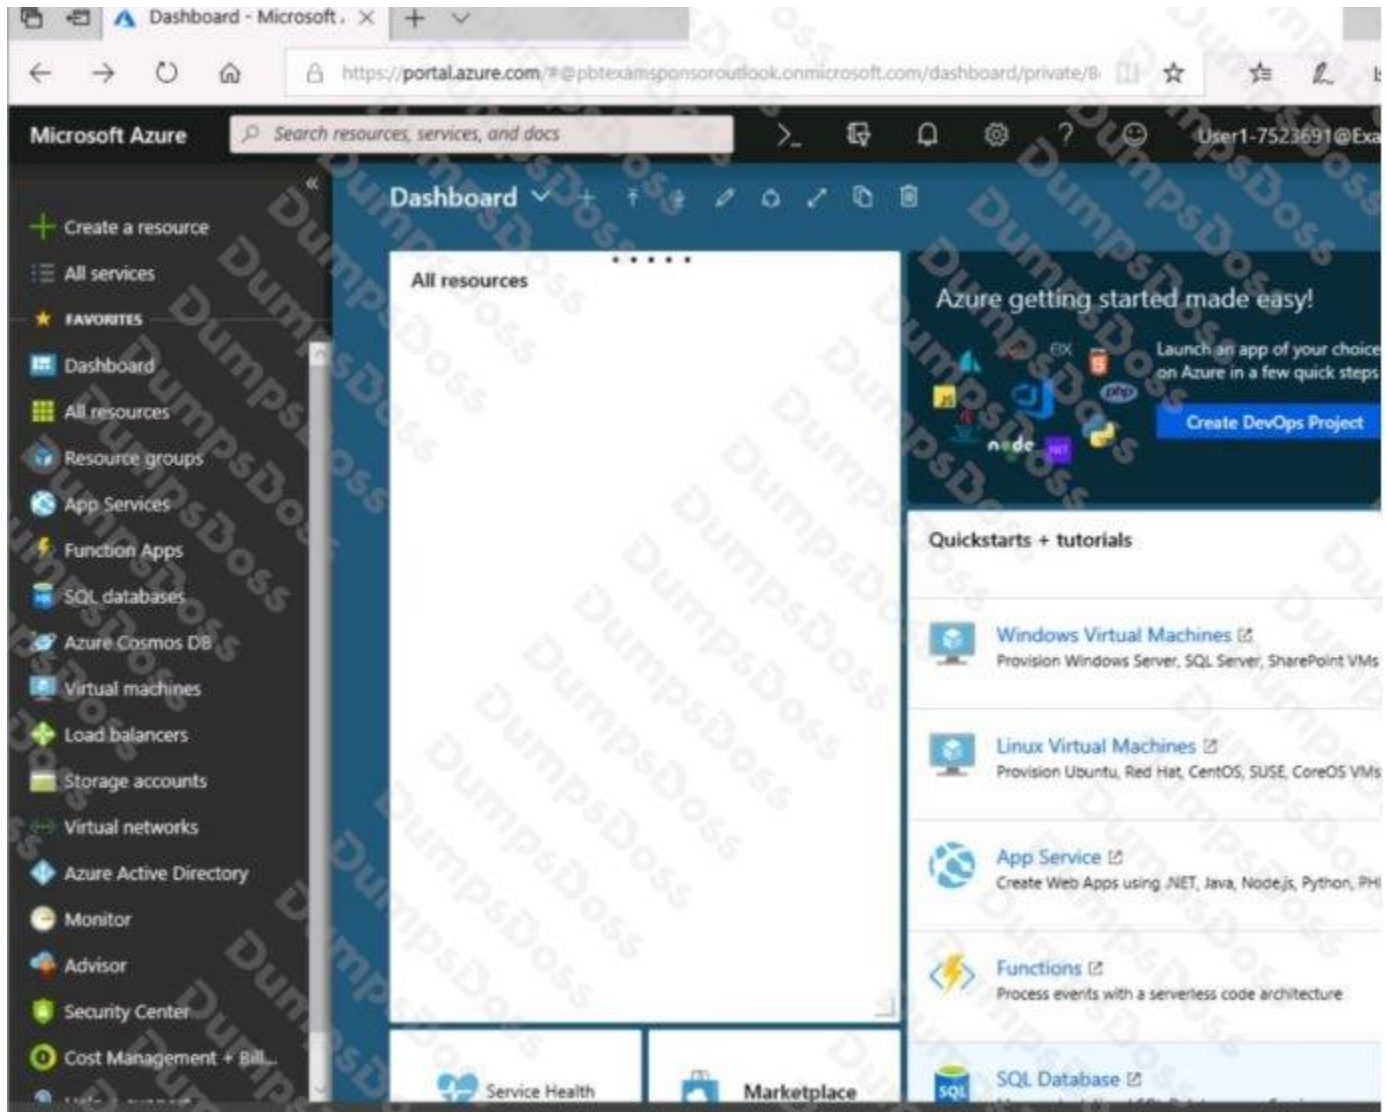
Topic	No. of Questions
Topic 1, Case Study 1	3
Topic 2, Case Study 2	4
Topic 3, Case Study 3	4
Topic 4, Case Study 4	5
Topic 5, Case Study 5	2
Topic 6, Case Study 6	2
Topic 7, Case Study 7	4
Topic 8, Case Study 8	7
Topic 9, Mixed Questions	418
Total	449



**QUESTION NO: 1 - (SIMULATION)****SIMULATION**


Click to expand each objective. To connect to the Azure portal, type <https://portal.azure.com> in the browser address bar.





[Home](#) > [Storage accounts](#) > Create storage account

## Create storage account

 Validation passed[Basics](#) [Advanced](#) [Tags](#) [Review + create](#)

### BASICS

Subscription	Microsoft AZ-300 5
Resource group	corpdatalod7523690
Location	East US
Storage account name	corpdata7523690n1
Deployment model	Resource manager
Account kind	StorageV2 (general purpose v2)
Replication	Read-access geo-redundant storage (RA-GRS)
Performance	Standard
Access tier (default)	Hot

### ADVANCED

Secure transfer required	Enabled
Hierarchical namespace	Disabled

[Create](#)[Previous](#)[Next](#)[Download a template for automation](#)

[Home](#) > [Storage accounts](#) > Create storage account

## Create storage account

**\*\*\* Submitting deployment...**

Submitting the deployment template for resource 'corpdatalod7523690'.

[Basics](#) [Advanced](#) [Tags](#) [Review + create](#)

### BASICS

Subscription	Microsoft AZ-300 5
Resource group	corpdatalod7523690
Location	East US
Storage account name	corpdata7523690n1
Deployment model	Resource manager
Account kind	StorageV2 (general purpose v2)
Replication	Read-access geo-redundant storage (RA-GRS)
Performance	Standard
Access tier (default)	Hot

### ADVANCED

Secure transfer required	Enabled
Hierarchical namespace	Disabled



[Home](#) > Microsoft.StorageAccount-20181011170335 - Overview

## Microsoft.StorageAccount-20181011170335 - Overview

Deployment

[Delete](#) [Cancel](#) [Redeploy](#) [Refresh](#)[Overview](#)[Outputs](#)[Inputs](#)[Template](#)

### Your deployment is underway

Check the status of your deployment, manage resources, or troubleshoot deployment issues. Pin this page to your dashboard to easily find it next time.



Deployment

name: Microsoft.StorageAccount-20181011170335

Subscription: [Microsoft AZ-300 5](#)Resource group: [corpdatalod7523690](#)

#### DEPLOYMENT DETAILS [\(Download\)](#)

Start time: 10/11/2018 5:04:06 PM

Duration: 17 seconds


Correlation ID: bd0806a4-d1bd-42db-be6b-55e0ec38f49b

RESOURCE	TYPE	STATUS	OPERATI...
----------	------	--------	------------

No results.			
-------------	--	--	--

Home > Virtual machines > Create a virtual machine

## Create a virtual machine

 Validation failed. Required information is missing or not valid.

Basics • Disks Networking Management Guest config Tags Review + create

### PRODUCT DETAILS

Ubuntu Server 18.04 LTS

by Canonical

[Terms of use](#) | [Privacy policy](#)


Standard D2s v3

by Microsoft

[Terms of use](#) | [Privacy policy](#)

**Pricing not available for this offering**

View [Pricing details](#) for more information.

Subscription credits apply 

**0.0960 USD/hr**

[Pricing for other VM sizes](#)

### TERMS

By clicking "Create", I (a) agree to the legal terms and privacy statement(s) associated with the Marketplace offering(s) listed above; (b) authorize Microsoft to bill my current payment method for the fees associated with the offering(s), with the same billing frequency as my Azure subscription; and (c) agree that Microsoft may share my contact, usage and transactional information with the provider(s) of the offering(s) for support, billing and other transactional activities. Microsoft does not provide rights for third-party offerings. See the [Azure Marketplace Terms](#) for additional details.

When you are finished performing all the tasks, click the 'Next' button.

Note that you cannot return to the lab once you click the 'Next' button. Scoring occurs in the background while you complete the rest of the exam.

### Overview

The following section of the exam is a lab. In this section, you will perform a set of tasks in a live environment. While most functionality will be available to you as it would be in a live environment, some functionality (e.g., copy and paste, ability to navigate to external websites) will not be possible by design.

Scoring is based on the outcome of performing the tasks stated in the lab. In other words, it doesn't matter how you accomplish the task, if you successfully perform it, you will earn credit for that task.

Labs are not timed separately, and this exam may have more than one lab that you must complete. You can use as much time as you would like to complete each lab. But, you should manage your time appropriately to ensure that you are able to complete the lab(s) and all other sections of the exam in the time provided.

Please note that once you submit your work by clicking the Next button within a lab, you will NOT be able to return to the lab.

To start the lab



You may start the lab by clicking the Next button.

You plan to store media files in the corpdata7523690n1 storage account.

You need to configure the storage account to store the media files. The solution must ensure that only users who have access keys can download the media files and that the files are accessible only over HTTPS.

What should you do from the Azure portal?

**ANSWER: See solution below.**

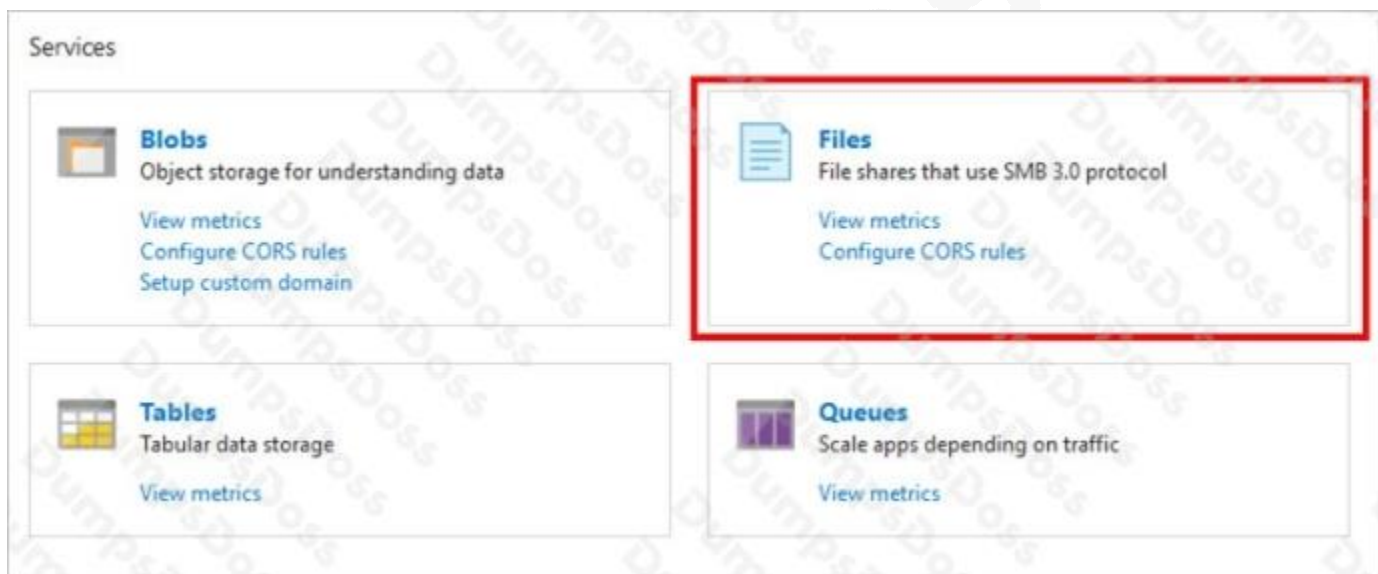
**Explanation:**

We should create an Azure file share.

Step 1: In the Azure portal, select All services. In the list of resources, type Storage Accounts. As you begin typing, the list filters based on your input. Select Storage Accounts. On the Storage Accounts window that appears.

Step 2: Locate the corpdata7523690n1 storage account.

Step 3: On the storage account page, in the Services section, select Files.



Step 4: On the menu at the top of the File service page, click + File share. The New file share page drops down.

Step 5: In Name type myshare. Click OK to create the Azure file share.

References: <https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-use-files-portal>

**QUESTION NO: 2**

An app uses a virtual network with two subnets. One subnet is used for the application server. The other subnet is used for a database server. A network virtual appliance (NVA) is used as a firewall.

Traffic destined for one specific address prefix is routed to the NVA and then to an on-premises database server that stores sensitive data. A Border Gateway Protocol (BGP) route is used for the traffic to the on-premises database server.

You need to recommend a method for creating the user-defined route.

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

NOTE: Each correct selection is worth one point.

- A. For the virtual network configuration, use a VPN.
- B. For the next hop type, use a virtual network peering.
- C. For the virtual network configuration, use Azure ExpressRoute.
- D. For the next hop type, use a virtual network gateway.

**ANSWER: A D****Explanation:**

You can create custom, or user-defined, routes in Azure to override Azure's default system routes, or to add additional routes to a subnet's route table. You can specify the following next hop types when creating a user-defined route:

- Virtual appliance: A virtual appliance is a virtual machine that typically runs a network application, such as a firewall.
- Virtual network gateway: Specify when you want traffic destined for specific address prefixes routed to a virtual network gateway. The virtual network gateway must be created with type VPN. You cannot specify a virtual network gateway created as type ExpressRoute in a user-defined route because with ExpressRoute, you must use BGP for custom routes.
- None: Specify when you want to drop traffic to an address prefix, rather than forwarding the traffic to a destination.
- Virtual network: Specify when you want to override the default routing within a virtual network.
- Internet: Specify when you want to explicitly route traffic destined to an address prefix to the Internet, or if you want traffic destined for Azure services with public IP addresses kept within the Azure backbone network.

**Incorrect Answers:**

B: You cannot specify VNet peering or VirtualNetworkServiceEndpoint as the next hop type in user-defined routes. Routes with the VNet peering or VirtualNetworkServiceEndpoint next hop types are only created by Azure, when you configure a virtual network peering, or a service endpoint.

C: You cannot specify a virtual network gateway created as type ExpressRoute in a user-defined route because with ExpressRoute, you must use BGP for custom routes.

**Reference:**

<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-udr-overview>

**QUESTION NO: 3**

You have an Azure subscription that contains two storage accounts named storagecontoso1 and storagecontoso2. Each storage account contains a queue service, a table service, and a blob service.

You develop two apps named App1 and App2. You need to configure the apps to store different types of data to all the storage services on both the storage accounts.

How many endpoints should you configure for each app?

- A. 2
- B. 3
- C. 6
- D. 12

**ANSWER: A**

**Explanation:**

Each app needs a service endpoint in each Storage Account.

References: <https://docs.microsoft.com/en-us/azure/storage/common/storage-network-security>

#### QUESTION NO: 4 - (HOTSPOT)

HOTSPOT

You need to prepare the environment to implement the planned changes for Server2.

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

NOTE: Each correct selection is worth one point.

**Hot Area:**



**Answer Area**

From the Azure portal:

Create an Azure Migrate project
Create a Recovery Services vault
Upload a management certificate
Create an Azure Import/Export job

On Server2:

Enable Hyper-V Replica
Install the Azure File Sync agent
Create a collector virtual machine
Configure Hyper-V storage migration
Install the Azure Site Recovery Provider

**ANSWER:**

**Answer Area**

From the Azure portal:

Create an Azure Migrate project
Create a Recovery Services vault
Upload a management certificate
Create an Azure Import/Export job

On Server2:

Enable Hyper-V Replica
Install the Azure File Sync agent
Create a collector virtual machine
Configure Hyper-V storage migration
Install the Azure Site Recovery Provider

**Explanation:**

Box 1: Create a Recovery Services vault

Create a Recovery Services vault on the Azure Portal.

Box 2: Install the Azure Site Recovery Provider

Azure Site Recovery can be used to manage migration of on-premises machines to Azure.

Scenario: Migrate the virtual machines hosted on Server1 and Server2 to Azure. Server2 has the Hyper-V host role.

References: <https://docs.microsoft.com/en-us/azure/site-recovery/migrate-tutorial-on-premises-azure>

**QUESTION NO: 5 - (HOTSPOT)****HOTSPOT**

You have the Azure SQL Database servers shown in the following table.

Name	Elastic pool
sqlserver1	Pool1
sqlserver2	Pool1, Pool2

You have the Azure SQL databases shown in the following table.

Name	Azure SQL Database server	Elastic pool
DB1	sqlserver1	None
DB2	sqlserver1	Pool1
DB3	sqlserver2	Pool1
DB4	sqlserver2	Pool2

You create a failover group named failover1 that has the following settings:

- Primary server: sqlserver1
- Secondary server: sqlserver2
- Read/Write failover policy Automatic
- Read/Write grace period (hours): 1 hour

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

### Statements

Yes

No

You can add DB1 to failover1.

☐☐

You can add DB3 to failover1.

☐☐

Sqlserver1 and sqlserver2 are in the same Azure region.

☐☐

ANSWER:



## Answer Area

### Statements

**Yes****No**

You can add DB1 to failover1.

☒☐

You can add DB3 to failover1.

☐☒

Sqlserver1 and sqlserver2 are in the same Azure region.

☐☒

### Explanation:

Box 1: Yes

DB1 is on the primary server

Box 2: No

DB3 is on the secondary server.

You can put all or several databases within an elastic pool into the same failover group.

Box 3: No

A failover group is a named group of databases managed by a single server or within a managed instance that can fail over as a unit to another region in case all or some primary databases become unavailable due to an outage in the primary region.

The secondary cannot be in the same region as the primary.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

## QUESTION NO: 6

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

Name	Region
RG1	East US
RG2	West US

The subscription contains the storage accounts shown in the following table.

Name	Resource group	Location	Account kind
Storage1	RG1	West US	BlobStorage
Storage2	RG2	West US	Storage (general purpose v1)
Storage3	RG1	East US	Storage V2 (general purpose v2)

You create a Recovery Services vault named Vault1 in RG1 in the West US location.

You need to identify which storage accounts can be used to archive the diagnostics logs of Vault1.

Which storage accounts should you identify?

- A. Storage1 only
- B. Storage2 only
- C. Storage3 only
- D. Storage1 or Storage2 only
- E. Storage1 or Storage3 only

**ANSWER: D E**

**Explanation:**

The same region or the same resource group.

#### QUESTION NO: 7

Note: This question is part of 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.

You are planning to create a virtual network that has a scale set that contains six virtual machines (VMs).

A monitoring solution on a different network will need access to the VMs inside the scale set.

You need to define public access to the VMs.

Solution: Use Remote Desktop Protocol (RDP) to connect to the VM in the scale set.

Does the solution meet the goal?

- A. Yes
- B. No

**ANSWER: A**

**Explanation:**

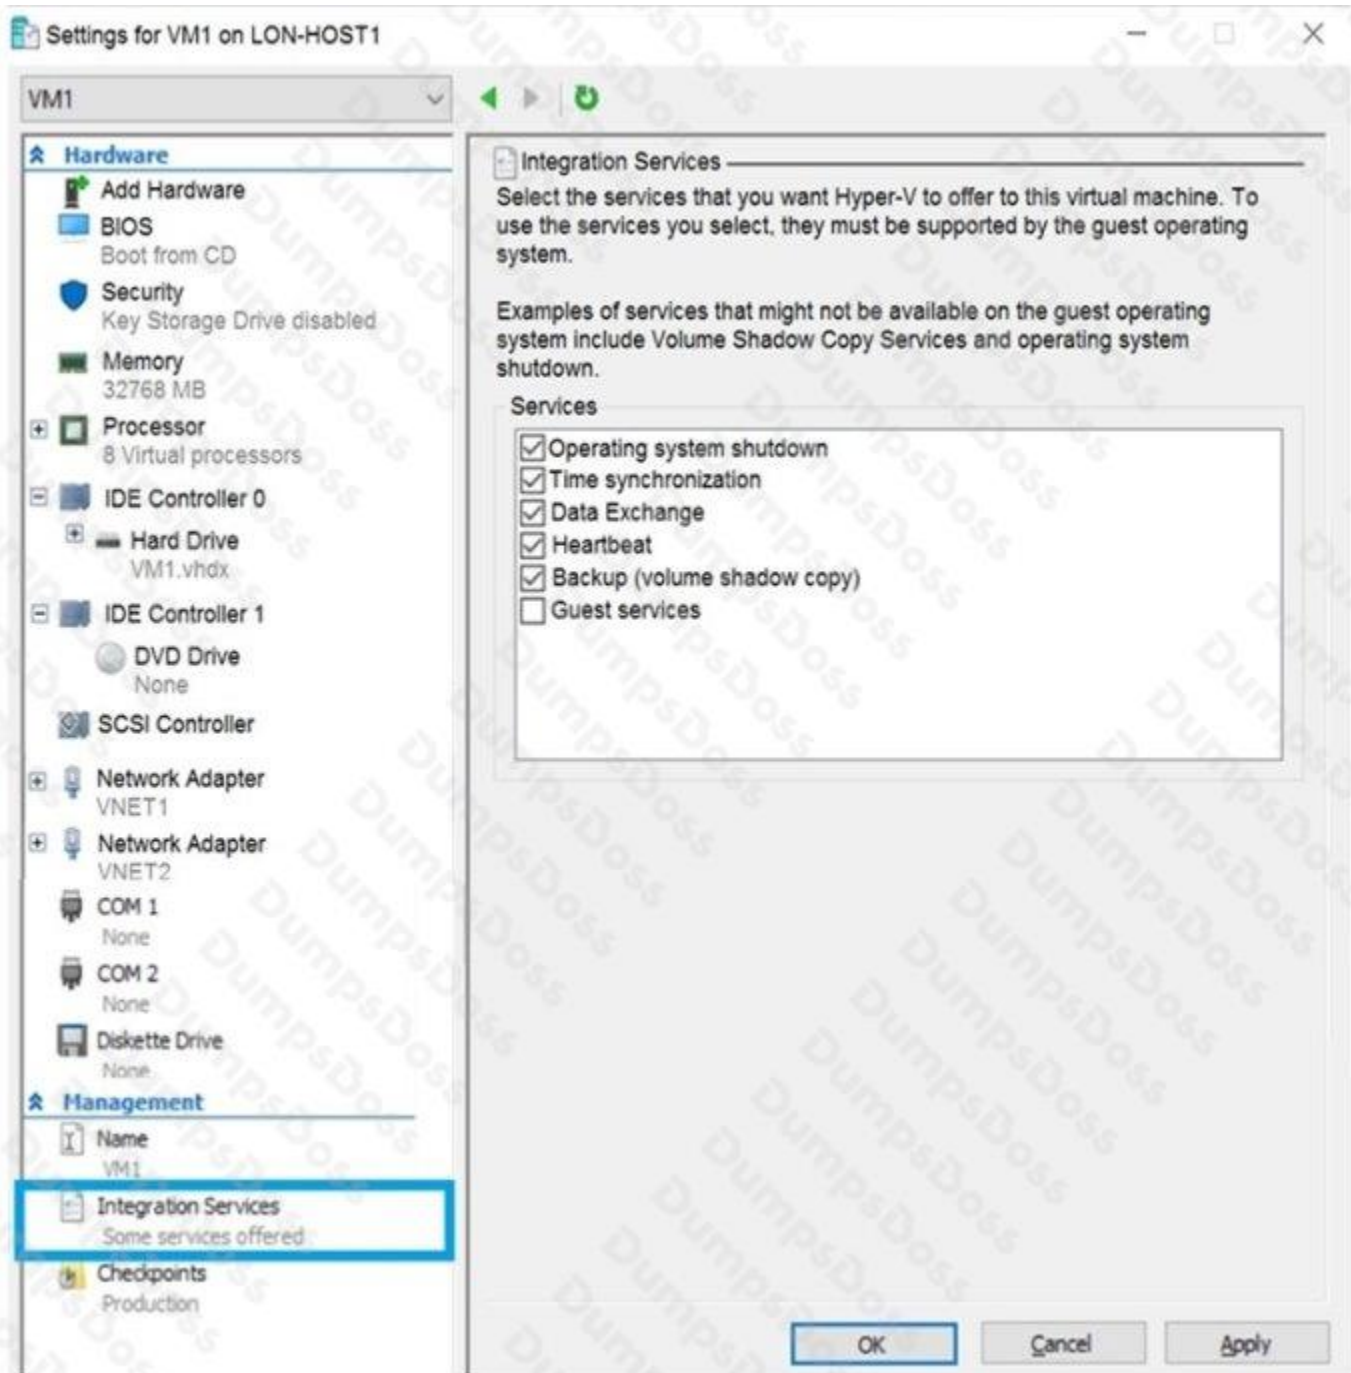
Instead, deploy a standalone VM that has a public IP address to the virtual network.

**QUESTION NO: 8**

You have an on-premises virtual machine named VM1 configured as shown in the following exhibit.

DUMPSBOSS.COM





VM1 is started.

You need to create a new virtual machine image in Azure from VM1.

Which three actions should you perform before you create the new image? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Run Add-AzVhd and specify a file share as the destination.

- B. Convert the disk type to VHD.
- C. Remove the Backup (volume shadow copy) integration service.
- D. Generalize VM1.
- E. Reduce the amount of memory to 16 GB.
- F. Run Add-AzVhd and specify a blob service container as the destination.

**ANSWER: B D F**

**Explanation:**

D: Sysprep removes all your personal account and security information, and then prepares the machine to be used as an image.

B, F: The Add-AzureVhd cmdlet uploads on-premises virtual hard disks, in .vhd file format, to a blob storage account as fixed virtual hard disks.

Reference: <https://docs.microsoft.com/en-us/powershell/module/azurerm.compute/add-azurermvhd?view=azurermps-6.13.0>  
<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource>

**QUESTION NO: 9**

You have an Azure subscription that contains three virtual networks named VNet1, VNet2, and VNet3. VNet2 contains a virtual appliance named VM2 that operates as a router.

You are configuring the virtual networks in a hub and spoke topology that uses VNet2 as the hub network.

You plan to configure peering between VNet1 and VNet2 and between VNet2 and VNet3.

You need to provide connectivity between VNet1 and VNet3 through VNet2.

Which two configurations should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. On the peering connections, allow forwarded traffic
- B. Create a route filter
- C. On the peering connections, allow gateway transit
- D. Create route tables and assign the table to subnets
- E. On the peering connections, use remote gateways

**ANSWER: C E**

**Explanation:**

Allow gateway transit: Check this box if you have a virtual network gateway attached to this virtual network and want to allow traffic from the peered virtual network to flow through the gateway.

The peered virtual network must have the Use remote gateways checkbox checked when setting up the peering from the other virtual network to this virtual network.

Note: VNet2 is the hub network.

References: <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-manage-peering#requirements-and-constraints>

## QUESTION NO: 10

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.

You have an Azure Cosmos DB database that contains a container named Container1. The partition key for Container1 is set to /day. Container1 contains the items shown in the following table.



Name	Content
Item1	{ "id": "1", "day": "Mon", "value": "10" }
Item2	{ "id": "2", "day": "Mon", "value": "15" }
Item3	{ "id": "3", "day": "Tue", "value": "10" }
Item4	{ "id": "4", "day": "Wed", "value": "15" }

You need to programmatically query Azure Cosmos DB and retrieve Item1 and Item2 only.

Solution: You run the following query.

```
SELECT day FROM c
```

WHERE c.value = "10" OR c.value = "15" You set the EnableCrossPartitionQuery property to True.

Does this meet the goal?

A. Yes

B. No

**ANSWER: B**

**Explanation:**

Returns Item1, Item2, Item3, and Item4.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/sql-query-where>

**QUESTION NO: 11**

You manage a solution in Azure that consists of a single application which runs on a virtual machine (VM). Traffic to the application has increased dramatically.

The application must not experience any downtime and scaling must be dynamically defined.

You need to define an auto-scale strategy to ensure that the VM can handle the workload.

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

NOTE: Each correct selection is worth one point.

- A. Deploy application automatic vertical scaling.
- B. Create a VM availability set.
- C. Create a VM scale set.
- D. Deploy application automatic horizontal scaling.
- E. Deploy a custom auto-scale implementation.

**ANSWER: C D E****QUESTION NO: 12**

You have an Azure Service Bus and two clients named Client1 and Client2.

You create a Service Bus queue named Queue1 as shown in the exhibit. (Click the Exhibit tab.)

### Create queue ×

Service Bus

---

Name <sup>\*</sup> ⓘ

Queue1 ✓

Max queue size

1 GB ▼

Message time to live ⓘ

Days: 14 Hours: 0 Minutes: 0 Seconds: 0

Lock duration ⓘ

Days: 0 Hours: 0 Minutes: 0 Seconds: 30

☒ Enable duplicate detection ⓘ

Duplicate detection window ⓘ

Days: 0 Hours: 0 Minutes: 10 Seconds: 0

☐ Enable dead lettering on message expiration ⓘ

☒ Enable sessions ⓘ

Client1 send messages to Queue1 as shown in the following table.

Time hh:mm:ss	Message
12:01:01	M3
12:01:02	M2
12:01:03	M1
12:01:04	M3

Client2 reads the messages from Queue1 at 12:01:05.

How will the messages be presented to Client2?



- A. Client2 will read three messages in the following order: M1, M2, and then M3.
- B. Client2 will read three messages in the following order: M3, M1, and then M2.
- C. Client2 will read four messages in the following order: M3, M1, M2 and then M3.
- D. Client2 will read four messages in the following order: M3, M2, M1 and then M3.

**ANSWER: D**

**Explanation:**

It should be M3, M2, M1 as duplicate detection is enabled, and the duplication detection window is set to 10 minutes. The second M3 message in the queue would be discarded.

Note 1: Duplicate detection enables the sender resend the same message, and the queue or topic discards any duplicate copies.

Note 2: Queues offer First In, First Out (FIFO) message delivery to one or more competing consumers. That is, receivers typically receive and process messages in the order in which they were added to the queue, and only one message consumer receives and processes each message.

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions>

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/duplicate-detection>

**QUESTION NO: 13 - (DRAG DROP)**

**DRAG DROP**

You have an Azure subscription that is used by four departments in your company. The subscription contains 10 resource groups. Each department uses resources in several resource groups.

You need to send a report to the finance department. The report must detail the costs for each department.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Select and Place:**

## Actions

## Answer Area

Assign a tag to each resource group.

Assign a tag to each resource.

Download the usage report.

From the Cost analysis blade, filter the view by tag.

Open the **Resource costs** blade of each resource group.



## ANSWER:

## Actions

## Answer Area

Assign a tag to each resource group.

Assign a tag to each resource.

Download the usage report.

From the Cost analysis blade, filter the view by tag.

Open the **Resource costs** blade of each resource group.



Assign a tag to each resource.

From the Cost analysis blade, filter the view by tag.

Download the usage report.



## Explanation:

Box 1: Assign a tag to each resource.

You apply tags to your Azure resources giving metadata to logically organize them into a taxonomy. After you apply tags, you can retrieve all the resources in your subscription with that tag name and value. Each resource or resource group can have a maximum of 15 tag name/value pairs. Tags applied to the resource group are not inherited by the resources in that resource group.

Box 2: From the Cost analysis blade, filter the view by tag

After you get your services running, regularly check how much they're costing you. You can see the current spend and burn rate in Azure portal.

1. Visit the Subscriptions blade in Azure portal and select a subscription.

1. You should see the cost breakdown and burn rate in the popup blade.

2. Click Cost analysis in the list to the left to see the cost breakdown by resource. Wait 24 hours after you add a service for the data to populate.

3. You can filter by different properties like tags, resource group, and timespan. Click Apply to confirm the filters and Download if you want to export the view to a Comma-Separated Values (.csv) file.

Box 3: Download the usage report

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-using-tags>  
<https://docs.microsoft.com/en-us/azure/billing/billing-getting-started>

#### QUESTION NO: 14

You have an Azure subscription that contains a policy-based virtual network gateway named GW1 and a virtual network named VNet1.

You need to ensure that you can configure a point-to-site connection from an on-premises computer to VNet1.

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

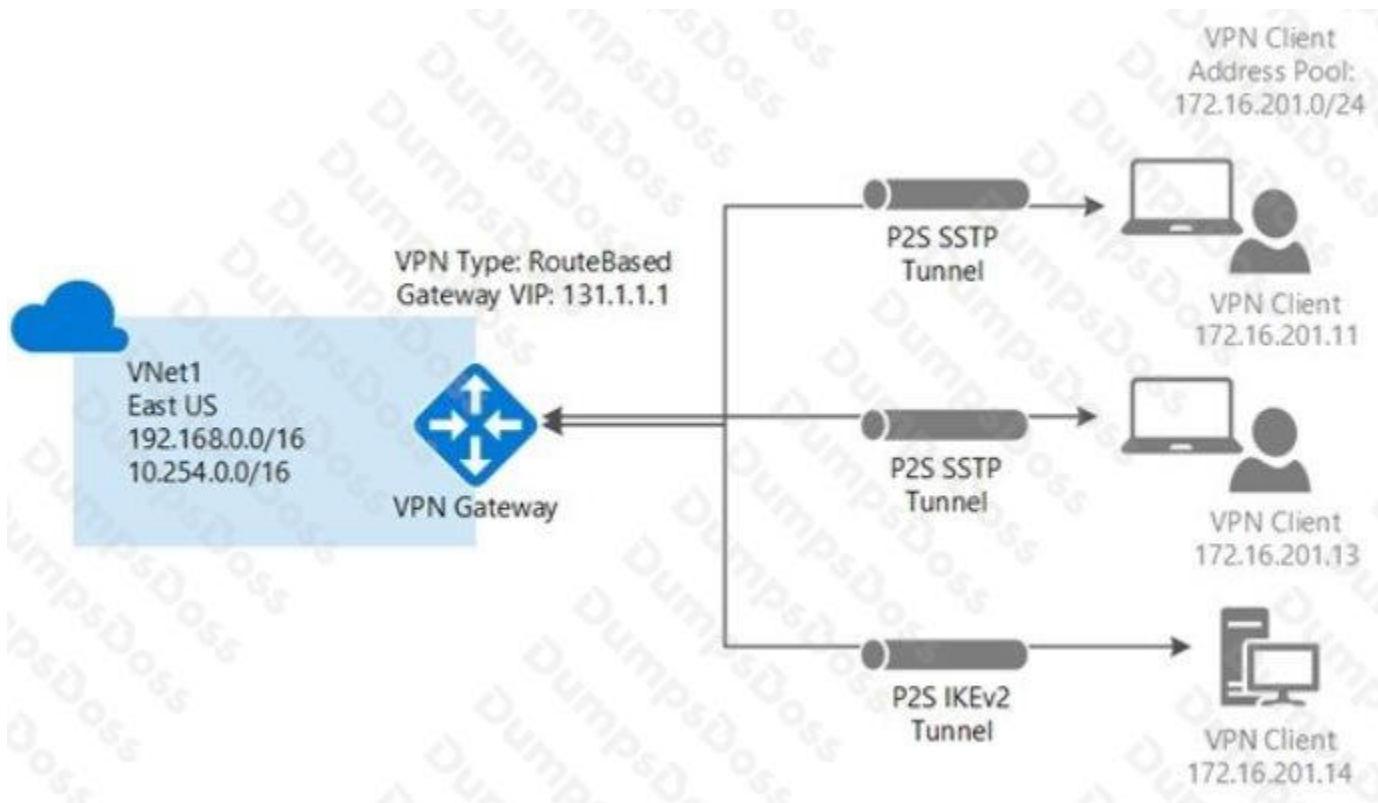
NOTE: Each correct selection is worth one point.

- A. Add a service endpoint to VNet1.
- B. Add a public IP address space to VNet1.
- C. Create a route-based virtual network gateway.
- D. Reset GW1.
- E. Delete GW1.
- F. Add a connection to GW1.

**ANSWER: C E**

**Explanation:**

Need a RouteBased VPN gateway.



References: <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/vpn-gateway/vpn-gateway-howto-point-to-site-resource-manager-portal.md>

### QUESTION NO: 15

A company is migrating an existing on-premises third-party website to Azure. The website is stateless.

The company does not have access to the source code for the website. They have the original installer.

The number of visitors at the website varies throughout the year. The on-premises infrastructure was resized to accommodate peaks but the extra capacity was not used.

You need to implement a virtual machine scale set instance.

What should you do

- A. Use an autoscale setting with unlimited maximum number of instances.
- B. Use an autoscale setting to scale instances vertically.
- C. Use only default diagnostics metrics to trigger autoscaling.
- D. Use Azure Monitor to create autoscale settings using custom metrics.

**ANSWER: D**

**Explanation:**



With Azure Monitor you can auto scale by custom metric for Virtual Machine Scale Sets.

Note: By default, Resource Manager-based Virtual Machines and Virtual Machine Scale Sets emit basic (host-level) metrics. In addition, when you configure diagnostics data collection for an Azure VM and VMSS, the Azure diagnostic extension also emits guest-OS performance counters (commonly known as "guest-OS metrics"). You use all these metrics in autoscale rules.

Note 2: In-guest VM metrics with the Azure diagnostics extension

The Azure diagnostics extension is an agent that runs inside a VM instance. The agent monitors and saves performance metrics to Azure storage. These performance metrics contain more detailed information about the status of the VM, such as AverageReadTime for disks or PercentIdleTime for CPU. You can create autoscale rules based on a more detailed awareness of the VM performance, not just the percentage of CPU usage or memory consumption.

Reference: <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/autoscale-custom-metric>  
<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/autoscale-common-metrics>

### QUESTION NO: 16

You have SQL Server on an Azure virtual machine named SQL1.

You need to automate the backup of the databases on SQL1 by using Automated Backup v2 for the virtual machines. The backups must meet the following requirements:

- Meet a recovery point objective (RPO) of 15 minutes.
- Retain the backups for 30 days. ▪ Encrypt the backups at rest.

What should you provision as part of the backup solution?

- A. Elastic Database jobs
- B. Azure Key Vault
- C. an Azure Storage account
- D. a Recovery Services vault

### ANSWER: C

#### Explanation:

An Azure storage account is used for storing Automated Backup files in blob storage. A container is created at this location to store all backup files. The backup file naming convention includes the date, time, and database GUID.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/automated-backup>

### QUESTION NO: 17

You create the following Azure role definition.

```
{
  "Name": "Role1",
  "Id": "80808080-8080-8080-8080-808080808080",
  "IsCustom": false,
  "Description": "",
  "Actions": [
    "Microsoft.Storage/*/read",
    "Microsoft.Network/*/read",
    "Microsoft.Compute/virtualMachines/start/action",
    "Microsoft.Compute/virtualMachines/restart/action",
    "Microsoft.Authorization/*/read"],
  "NotActions": [ ],
  "DataActions": [ ],
  "NotDataActions": [ ],
  "AssignableScopes": [ ]
}
```

You need to create Role1 by using the role definition.

Which two values should you modify before you create Role1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. AssignableScopes
- B. Description
- C. DataActions
- D. IsCustom
- E. Id

**ANSWER: A D**

**Explanation:**

Part of example: "IsCustom": true,

"AssignableScopes": [

"/subscriptions/{subscriptionId1}",

"/subscriptions/{subscriptionId2}", "/subscriptions/{subscriptionId3}"

The following shows what a custom role looks like as displayed in JSON format. This custom role can be used for monitoring and restarting virtual machines.

```
{
```

```
"Name": "Virtual Machine Operator",
"Id": "888888888-8888-8888-8888-888888888888",
"IsCustom": true,
"Description": "Can monitor and restart virtual machines.",
"Actions": [
  "Microsoft.Storage/*/read",
  "Microsoft.Network/*/read",
  "Microsoft.Compute/*/read",
  "Microsoft.Compute/virtualMachines/start/action",
  "Microsoft.Compute/virtualMachines/restart/action",
  "Microsoft.Authorization/*/read",
  "Microsoft.ResourceHealth/availabilityStatuses/read",
  "Microsoft.Resources/subscriptions/resourceGroups/read",
  "Microsoft.Insights/alertRules/*",
  "Microsoft.Insights/diagnosticSettings/*",
  "Microsoft.Support/*"
],
"NotActions": [],
>DataActions": [],
"NotDataActions": [],
"AssignableScopes": [
  "/subscriptions/{subscriptionId1}",
  "/subscriptions/{subscriptionId2}",
  "/subscriptions/{subscriptionId3}"
]}
```

Reference: <https://docs.microsoft.com/en-us/azure/role-based-access-control/custom-roles>

### QUESTION NO: 18 - (DRAG DROP)

DRAG DROP

You need to identify the appropriate sizes for the Azure virtual machines.

Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

## Actions

From the Azure portal, create an Azure Migrate assessment.

From the Azure portal, create an Azure Migrate project.

From VM1, connect to the collector virtual machine and run the Azure Site Recovery deployment planner.

From the Azure portal, download an OVA file.

From VM1, connect to the collector virtual machine and run the Azure Migrate Collector.

From Microsoft Download Center, download the Azure Site Recovery deployment planner

From VM1, run the Deploy OVF Template wizard.

## Answer Area



**ANSWER:****Actions****Answer Area**

From the Azure portal, create an Azure Migrate project.

From the Azure portal, download an OVA file.

From VM1, connect to the collector virtual machine and run the Azure Site Recovery deployment planner.

From VM1, run the Deploy OVF Template wizard.

From VM1, connect to the collector virtual machine and run the Azure Migrate Collector.

From the Azure portal, create an Azure Migrate assessment.

From Microsoft Download Center, download the Azure Site Recovery deployment planner

**Explanation:**



References: <https://docs.microsoft.com/en-us/azure/migrate/tutorial-assessment-vmware> Legacy AZ-300: Implement Workloads and Security

**QUESTION NO: 19**

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.

You have an app named App1 that uses data from two on-premises Microsoft SQL Server databases named DB1 and DB2.

You plan to move DB1 and DB2 to Azure.

You need to implement Azure services to host DB1 and DB2. The solution must support server-side transactions across DB1 and DB2.

Solution: You deploy DB1 and DB2 as Azure SQL databases on the same Azure SQL Database server.

Does this meet the goal?

A. Yes

B. No

**ANSWER: B****Explanation:**

Instead deploy DB1 and DB2 to SQL Server on an Azure virtual machine.

Note: Understanding distributed transactions.

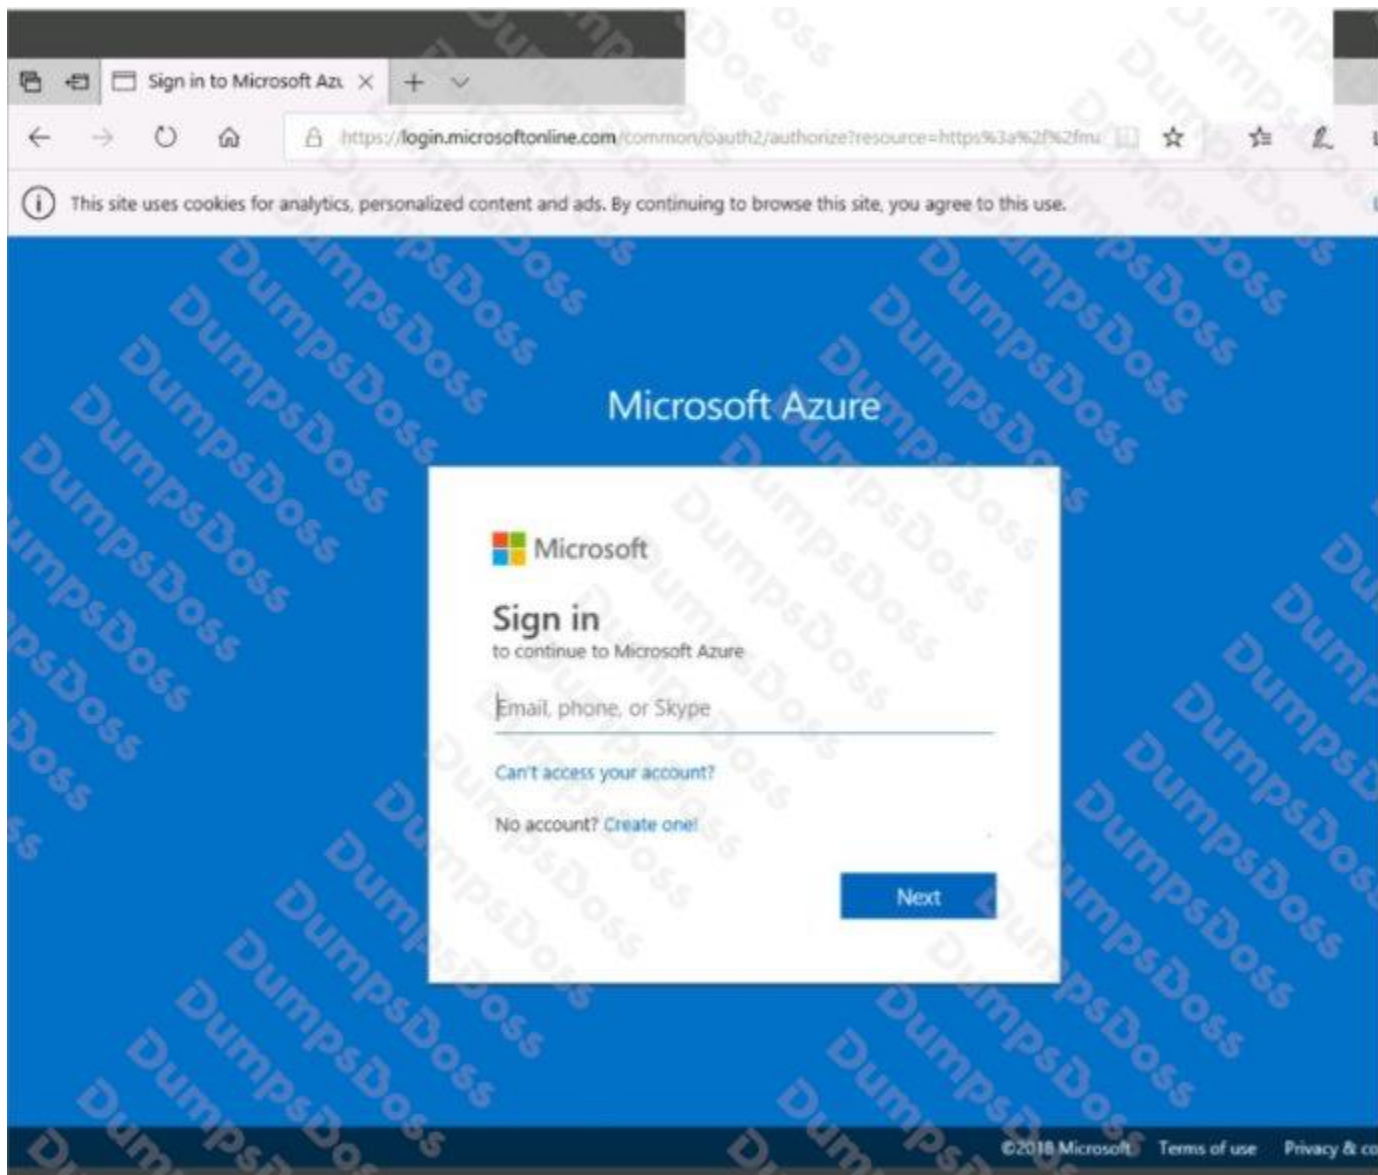
When both the database management system and client are under the same ownership (e.g. when SQL Server is deployed to a virtual machine), transactions are available and the lock duration can be controlled.


Reference:

<https://docs.particular.net/nservicebus/azure/understanding-transactionality-in-azure>

**QUESTION NO: 20 - (SIMULATION)****SIMULATION**


Click to expand each objective. To connect to the Azure portal, type <https://portal.azure.com> in the browser address bar.





[Home](#) > [Storage accounts](#) > Create storage account

## Create storage account

 Validation passed[Basics](#) [Advanced](#) [Tags](#) [Review + create](#)

### BASICS

Subscription	Microsoft AZ-300 5
Resource group	corpdatalod7523690
Location	East US
Storage account name	corpdata7523690n1
Deployment model	Resource manager
Account kind	StorageV2 (general purpose v2)
Replication	Read-access geo-redundant storage (RA-GRS)
Performance	Standard
Access tier (default)	Hot

### ADVANCED

Secure transfer required	Enabled
Hierarchical namespace	Disabled

[Create](#)[Previous](#)[Next](#)[Download a template for automation](#)

[Home](#) > [Storage accounts](#) > Create storage account

## Create storage account

\*\*\* Submitting deployment...

Submitting the deployment template for resource 'corpdatalod7523690'.

[Basics](#) [Advanced](#) [Tags](#) [Review + create](#)

### BASICS

Subscription	Microsoft AZ-300 5
Resource group	corpdatalod7523690
Location	East US
Storage account name	corpdata7523690n1
Deployment model	Resource manager
Account kind	StorageV2 (general purpose v2)
Replication	Read-access geo-redundant storage (RA-GRS)
Performance	Standard
Access tier (default)	Hot

### ADVANCED

Secure transfer required	Enabled
Hierarchical namespace	Disabled



[Home](#) > Microsoft.StorageAccount-20181011170335 - Overview

## Microsoft.StorageAccount-20181011170335 - Overview

Deployment

[Delete](#) [Cancel](#) [Redeploy](#) [Refresh](#) Overview Outputs Inputs Template

### Your deployment is underway

Check the status of your deployment, manage resources, or troubleshoot deployment issues. Pin this page to your dashboard to easily find it next time.



Deployment

name: Microsoft.StorageAccount-20181011170335

Subscription: [Microsoft AZ-300 5](#)Resource group: [corpdatalod7523690](#)

#### DEPLOYMENT DETAILS [\(Download\)](#)

Start time: 10/11/2018 5:04:06 PM

Duration: 17 seconds


Correlation ID: bd0806a4-d1bd-42db-be6b-55e0ec38f49b

RESOURCE	TYPE	STATUS	OPERATI...
----------	------	--------	------------

No results.			
-------------	--	--	--

Home > Virtual machines > Create a virtual machine

## Create a virtual machine

 Validation failed. Required information is missing or not valid.

Basics • Disks Networking Management Guest config Tags Review + create

### PRODUCT DETAILS

Ubuntu Server 18.04 LTS

by Canonical

[Terms of use](#) | [Privacy policy](#)

Standard D2s v3

by Microsoft

[Terms of use](#) | [Privacy policy](#)

**Pricing not available for this offering**

View [Pricing details](#) for more information.

Subscription credits apply ⓘ

**0.0960 USD/hr**

[Pricing for other VM sizes](#)

### TERMS

By clicking "Create", I (a) agree to the legal terms and privacy statement(s) associated with the Marketplace offering(s) listed above; (b) authorize Microsoft to bill my current payment method for the fees associated with the offering(s), with the same billing frequency as my Azure subscription; and (c) agree that Microsoft may share my contact, usage and transactional information with the provider(s) of the offering(s) for support, billing and other transactional activities. Microsoft does not provide rights for third-party offerings. See the [Azure Marketplace Terms](#) for additional details.

When you are finished performing all the tasks, click the 'Next' button.

Note that you cannot return to the lab once you click the 'Next' button. Scoring occurs in the background while you complete the rest of the exam.

### Overview

The following section of the exam is a lab. In this section, you will perform a set of tasks in a live environment. While most functionality will be available to you as it would be in a live environment, some functionality (e.g., copy and paste, ability to navigate to external websites) will not be possible by design.

Scoring is based on the outcome of performing the tasks stated in the lab. In other words, it doesn't matter how you accomplish the task, if you successfully perform it, you will earn credit for that task.

Labs are not timed separately, and this exam may have more than one lab that you must complete. You can use as much time as you would like to complete each lab. But, you should manage your time appropriately to ensure that you are able to complete the lab(s) and all other sections of the exam in the time provided.

Please note that once you submit your work by clicking the Next button within a lab, you will NOT be able to return to the lab.

To start the lab

You may start the lab by clicking the Next button.

You plan to deploy an application gateway named appgw1015 to load balance internal IP traffic to the Azure virtual machines connected to subnet0.

You need to configure a virtual network named VNET1015 to support the planned application gateway.

What should you do from the Azure portal?

**ANSWER: See explanation below.**

**Explanation:**

Step 1:

Click Networking, Virtual Network, and select VNET1015.

Step 2:

Click Subnets, and Click +Add on the VNET1015 - Subnets pane that appears.

Step 3:

On the Subnets page, click +Gateway subnet at the top to open the Add subnet page.



Step 4:

Locate subnet0 and add it.

References: <https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-howto-site-to-site-resource-manager-portal>