Question #1 *Topic 1*

HOTSPOT -

You have an Azure Migrate project that has the following assessment properties:

Target location: East US -

- Storage redundancy: Locally redundant.
- Comfort factor: 2.0
- Performance history: 1 month
- Percentile utilization: 95th
- Pricing tier: Standard
- Offer: Pay as you go

You discover the following two virtual machines:

- A virtual machine named VM1 that runs Windows Server 2016 and has 10 CPU cores at 20 percent utilization
- A virtual machine named VM2 that runs Windows Server 2012 and has four CPU cores at 50 percent utilization

How many CPU cores will Azure Migrate recommend for each virtual machine? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

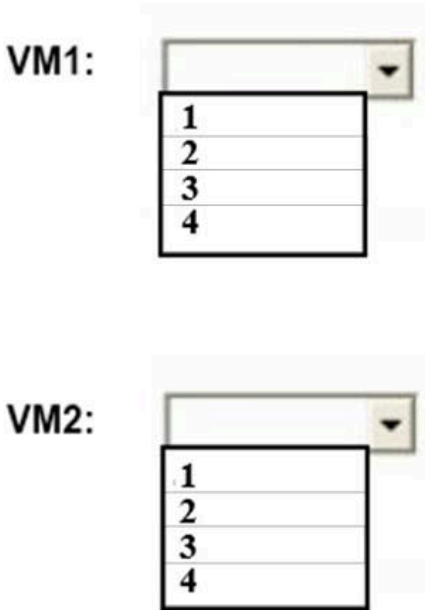

Suggested Answer:

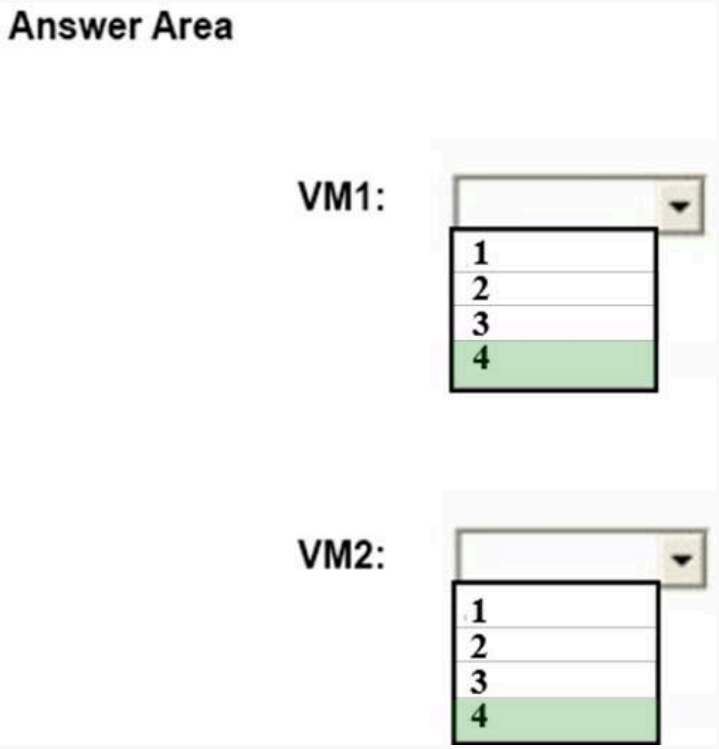

Box 2: 4 -

4 *0.50 * 0.95* 2 = 3.8

Note: The number of cores in the machines must be equal to or less than the maximum number of cores (128 cores) supported for an Azure VM.

If performance history is available, Azure Migrate considers the utilized cores for comparison. If a comfort factor is specified in the assessment settings, the number of utilized cores is multiplied by the comfort factor.

If there's no performance history, Azure Migrate uses the allocated cores, without applying the comfort factor.

References:

https://docs.microsoft.com/en-us/azure/migrate/concepts-assessment-calculation

Abbas786 Highly Voted 3 years, 7 months ago

Equation is: 'core usage x comfort factor'. The comfort factor is 2.0.

So VM 1 is 10 cores at 20% utilization which equals 2 cores. Multiply that the comfort factor and you get 4 cores.

VM 2 is 4 cores at 50% utilization which equals 2 cores. Multiply that the comfort factor and you get 4 cores.

upvoted 5 times

yidnekachew 1 year, 4 months ago

deeply elaborated, know I understood

upvoted 1 times

yidnekachew 1 year, 4 months ago

X1=20/100*10*2.0=>4

X2=50/100*4*2.0=>4

upvoted 1 times

tashakori Most Recent 1 year, 7 months ago

Given answer is correct

upvoted 1 times

AvoKikinha 2 years, 1 month ago

According to the information provided, the Azure Migrate project has a comfort factor of 2.0 1. The comfort factor is used to calculate the number of CPU cores recommended for each virtual machine. The equation is: core usage x comfort factor 2.

For VM1, which has 10 CPU cores at 20% utilization, the effective utilized cores are 2 2. Multiplying this by the comfort factor of 2.0 gives us a recommended number of 4 CPU cores 2.

For VM2, which has 4 CPU cores at 50% utilization, the effective utilized cores are 2 2. Multiplying this by the comfort factor of 2.0 gives us a recommended number of 4 CPU cores 2.

Therefore, Azure Migrate will recommend 4 CPU cores for each virtual machine 2.

upvoted 2 times

Question #2 *Topic 1*

You have an on-premises network that contains a Hyper-V host named Host1. Host1 runs Windows Server 2016 and hosts 10 virtual machines that run Windows
Server 2016.
You plan to replicate the virtual machines to Azure by using Azure Site recovery.
You create a Recovery Services vault named ASR1 and Hyper-V site named Site1.
You need to add Host1 to ASR1.
What should you do?

A. Download the installation file for the Azure Site Recovery Provider. Download the vault registration key. Install the Azure Site Recovery Provider on Host1 and register the server.

B. Download the installation file for the Azure Site Recovery Provider. Download the storage account key. Install the Azure Site Recovery Provider on each virtual machine and register the virtual machines.

C. Download the installation file for the Azure Site Recovery Provider. Download the vault registration key. Install the Azure Site Recovery Provider on each virtual machine and register the virtual machines.

D. Download the installation file for the Azure Site Recovery Provider. Download the storage account key. Install the Azure Site Recovery Provider on Host1 and register the server.

Suggested Answer: *A*
Download the Vault registration key. You need this when you install the Provider. The key is valid for five days after you generate it.
Install the Provider on each VMM server. You don't need to explicitly install anything on Hyper-V hosts.
Incorrect Answers:
B, D: Use the Vault Registration Key, not the storage account key.
References:
https://docs.microsoft.com/en-us/azure/site-recovery/migrate-tutorial-on-premises-azure

Community vote distribution
A (100%)

jwala1991 3 months, 1 week ago

Selected Answer: A

The steps to add Host1 to ASR1 involve
1.Download the Azure Site Recovery provider installation file.
2.Download the vault registration key.
3.Installing the Azure Site Recovery provider on Host1.
4.Registering the server
upvoted 1 times

Question #3 *Topic 1*

You have an Azure subscription that contains 100 virtual machines.
You regularly create and delete virtual machines.
You need to identify unattached disks that can be deleted.
What should you do?

A. From Azure Cost Management, open the Optimizer tab and create a report.

B. From Azure Cost Management, create a Cost Management report.

C. From Microsoft Azure Explorer, view the Account Management properties.

D. From the Azure portal, configure the Advisor recommendations.

Suggested Answer: *C*
You can find unused disks in the Azure Storage Explorer console. Once you drill down to the Blob containers under a storage account, you can see the lease state of the residing VHD (the lease state determines if the VHD is being used by any resource) and the VM to which it is leased out. If you find that the lease state and the VM fields are blank, it means that the VHD in question is unused.
Note: The ManagedBy property stores the Id of the VM to which Managed Disk is attached to. If the ManagedBy property is $null then it means that the Managed

Disk is not attached to a VM -
References:
https://cloud.netapp.com/blog/reduce-azure-storage-costs

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 1*

DRAG DROP -

You create an Azure Migrate project named Testis in a resource named test-migration.

You need to discover which on-premises virtual machines to assess for migration.

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:

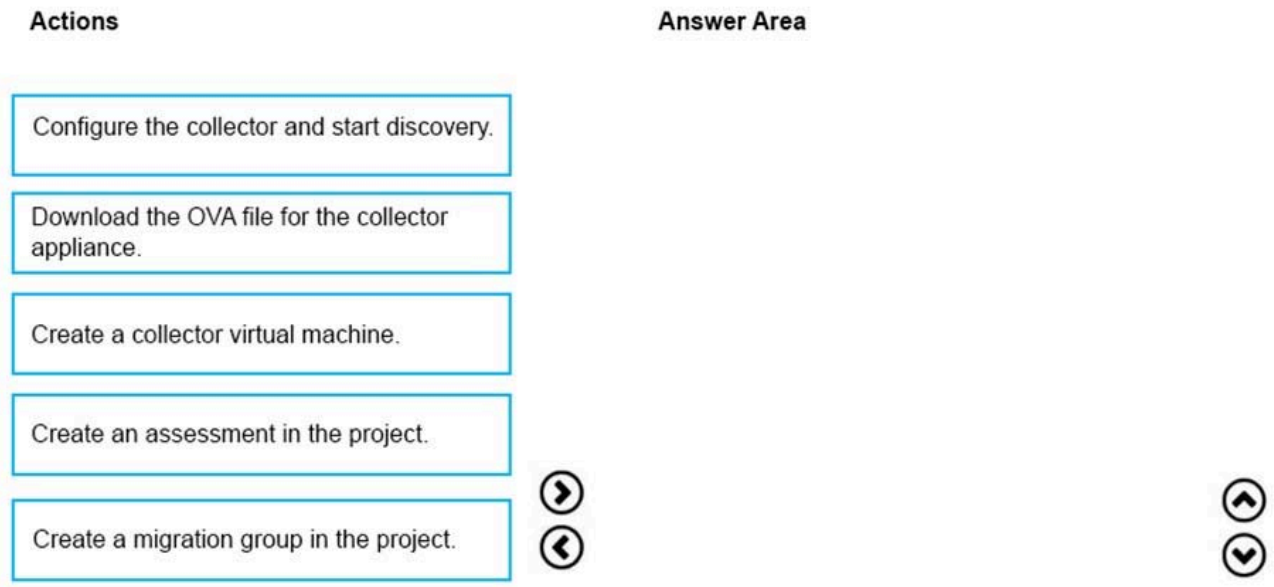

Suggested Answer:

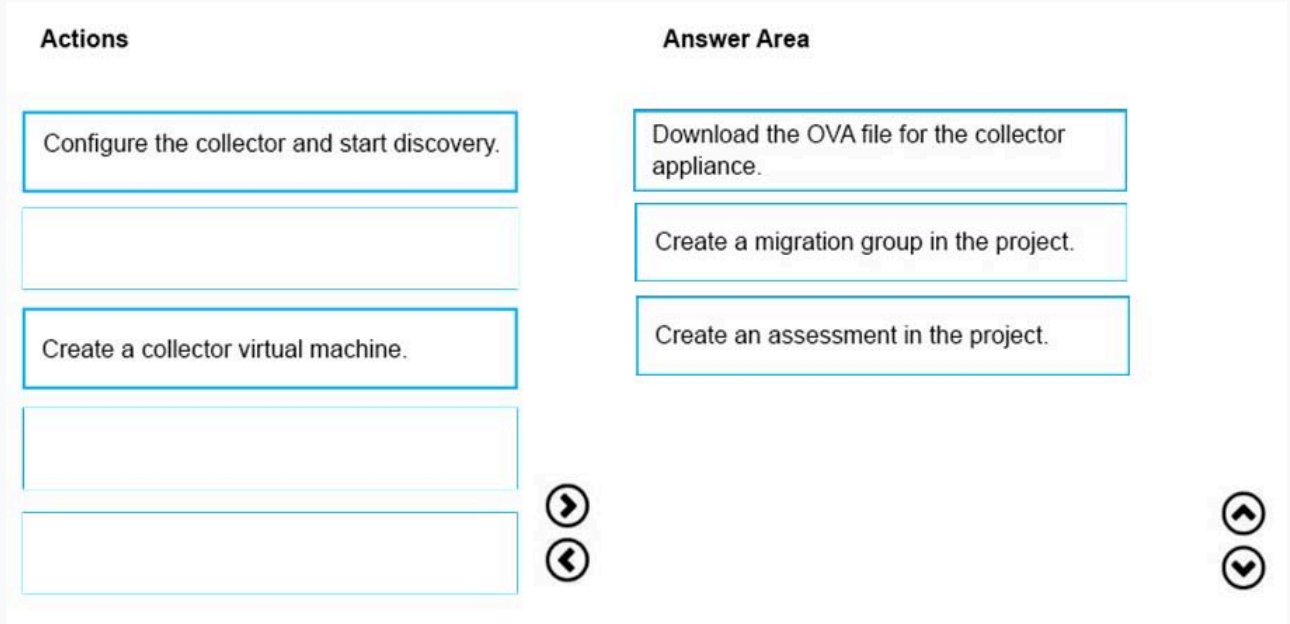

Step 1: Download the OVA file for the collection appliance

Azure Migrate uses an on-premises VM called the collector appliance, to discover information about your on-premises machines. To create the appliance, you download a setup file in Open Virtualization Appliance (.ova) format, and import it as a VM on your on-premises vCenter Server.

Step 2: Create a migration group in the project

For the purposes of assessment, you gather the discovered VMs into groups. For example, you might group VMs that run the same application. For more precise grouping, you can use dependency visualization to view dependencies of a specific machine, or for all machines in a group and refine the group.

Step 3: Create an assessment in the project

After a group is defined, you create an assessment for it.

References:

https://docs.microsoft.com/en-us/azure/migrate/migrate-overview

Currently there are no comments in this discussion, be the first to comment!

Question #5 *Topic 1*

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 is 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:

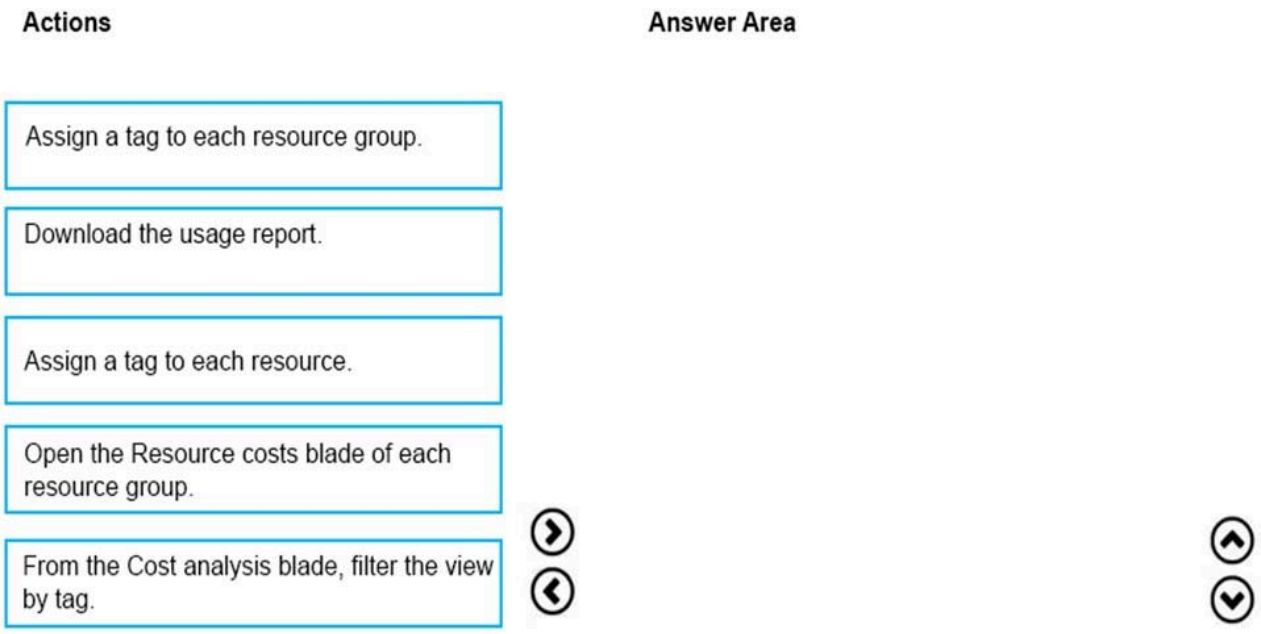

Suggested Answer:

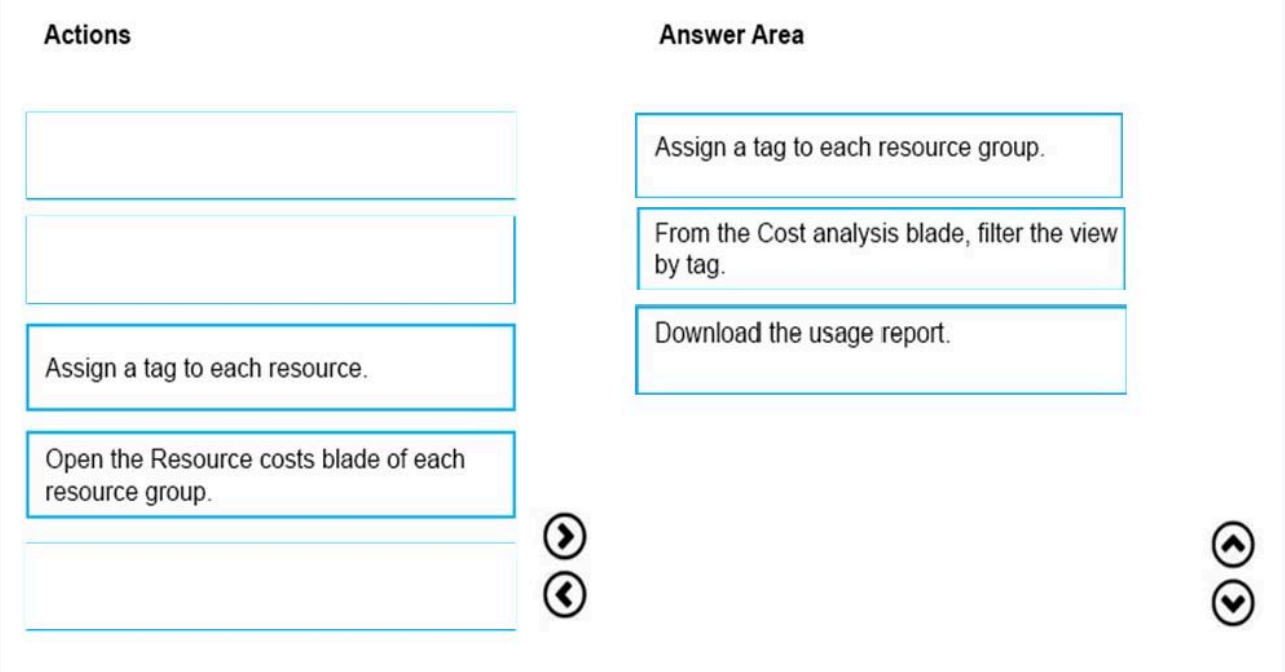

Box 1: Assign a tag to each resource.

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.
2. You should see the cost breakdown and burn rate in the popup blade.
3. 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.
4. 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

allray15 4 years, 7 months ago

Im just here to do first comment on March 2021

upvoted 1 times

HOTSPOT -

You have an Azure subscription named Subscription1.

You have a virtualization environment that contains the virtualization servers in the following table.

Name	Hypervisor	Run virtual machine
Server1	Hyper-V	VM1, VM2, VM3
Server2	VMWare	VMA, VMB, VMC

The virtual machines are configured as shown in the following table.

Name	Generation	Memory	Operating system (OS) disk	Data disk	OS
VM1	1	4 GB	200 GB	800 GB	Windows Server 2012 R2
VM2	1	12 GB	3 TB	200 GB	Red Hat Enterprise Linux 7.2
VM3	2	32 GB	100 GB	1 TB	Windows Server 2016
VMA	Not applicable	8 GB	100 GB	2 TB	Windows Server 2012 R2
VMB	Not applicable	16 GB	150 GB	1 TB	Red Hat Enterprise Linux 7.2
VMC	Not applicable	24 GB	500 GB	6 TB	Windows Server 2016

All the virtual machines use basic disks. VM1 is protected by using BitLocker Drive Encryption (BitLocker).

You plan to use Azure Site Recovery to migrate the virtual machines to Azure.

Which virtual machines can you migrate? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Virtual machines that can be migrated from Server1:

- VM1
- VM2
- VM3
- VM1 and VM2 only
- VM1 and VM3 only
- VM1 and VM2 and VM3

Virtual machines that can be migrated from Server2:

- VMA only
- VMB only
- VMC only
- VMA and VMB only
- VMA and VMC only
- VMA and VMB, and VMC

Suggested Answer:

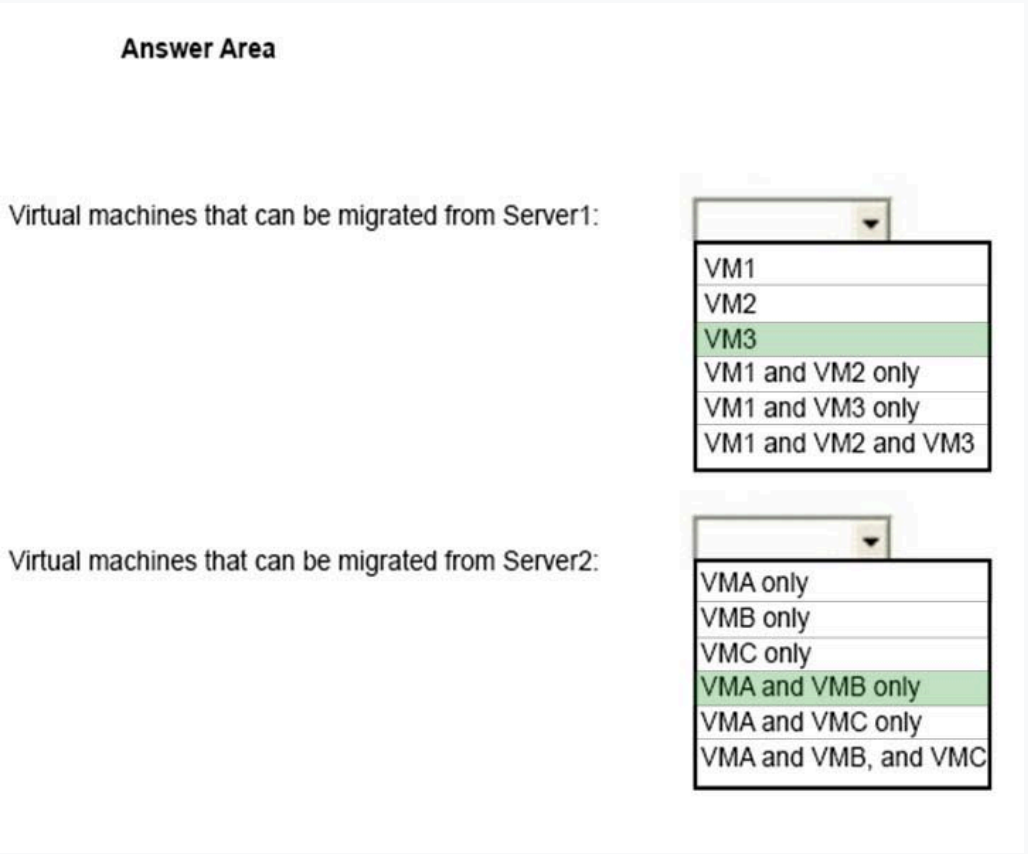

Box 1: VM3 -
Not VM1 as Bitlocker is not supported. BitLocker must be disabled before you enable replication for a VM.
Not VM2 as maximum Operating system disk size for a generation VM is 2,048 GB.

Box 2: VMA and VMB only -
Not VMC as the max data disk size is 4,095 GB
References:
https://docs.microsoft.com/en-us/azure/site-recovery/hyper-v-azure-support-matrix https://docs.microsoft.com/en-us/azure/site-recovery/vmware-physical-azure-support-matrix#azure-vm-requirements

Anirbanfiem 7 months, 2 weeks ago

Explanation for the Answer:
From Server1 (Hyper-V):

VM1: Uses BitLocker. ASR requires BitLocker to be disabled before migration. ❌

VM2:

Gen1 VM.

RHEL 7.2 is supported.

OS disk (3 TB) is ≤ 4 TB. ✅

VM3:

Gen2 VM (supported by Azure).

Windows Server 2016 is supported. ✅

From Server2 (VMware):

VM4, VMB, VMC:

All use basic disks (supported).

OS disk sizes ≤ 4 TB.

Windows Server 2012 R2, RHEL 7.2, and Windows Server 2016 are supported. ✅

How to Approach Similar Questions:
Check ASR's Supported Configurations:

Review OS versions, disk types/sizes, and hypervisor requirements.

Identify Blockers:

BitLocker, unsupported OS, oversized disks.

Map VM Properties to Requirements:

Use tables to cross-reference VM specs with ASR's documentation.

This structured approach ensures you systematically eliminate non-compliant VMs and identify valid candidates for migration.
upvoted 1 times

Question #7 *Topic 1*

Your company has an Azure subscription named Subscription1.
The company also has two on-premises servers named Server1 and Server2 that run Windows Server 2016. Server1 is configured as a DNS server that has a primary DNS zone named adatum.com. Adatum.com contains 1,000 DNS records.
You manage Server1 and Subscription1 from Server2. Server2 has the following tools installed:
⇨ The DNS Manager console
⇨ Azure PowerShell
⇨ Azure CLI 2.0
You need to move the adatum.com zone to Subscription1. The solution must minimize administrative effort.
What should you use?

A. the Azure portal

B. Azure PowerShell

C. the DNS Manager console

D. Azure CLI

Suggested Answer: *D*
Azure DNS supports importing and exporting zone files by using the Azure command-line interface (CLI). Zone file import is not currently supported via Azure
PowerShell or the Azure portal.
References:
https://docs.microsoft.com/en-us/azure/dns/dns-import-export

Currently there are no comments in this discussion, be the first to comment!

DRAG DROP -

You have an Azure Active Directory (Azure AD) tenant that has the initial domain name.

You have a domain name of contoso.com registered at a third-party registrar.

You need to ensure that you can create Azure AD users that have names containing a suffix of @contoso.com.

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

Select and Place:

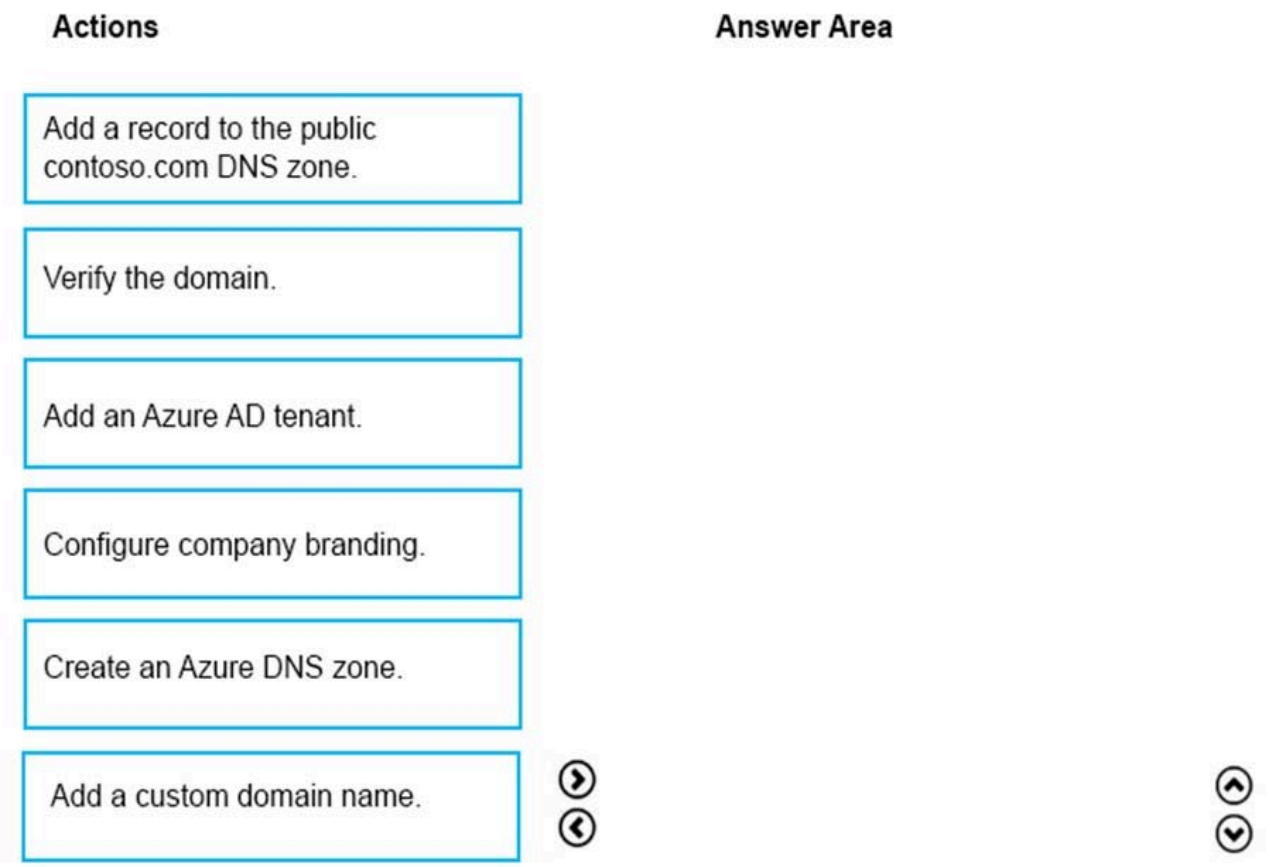

Suggested Answer:

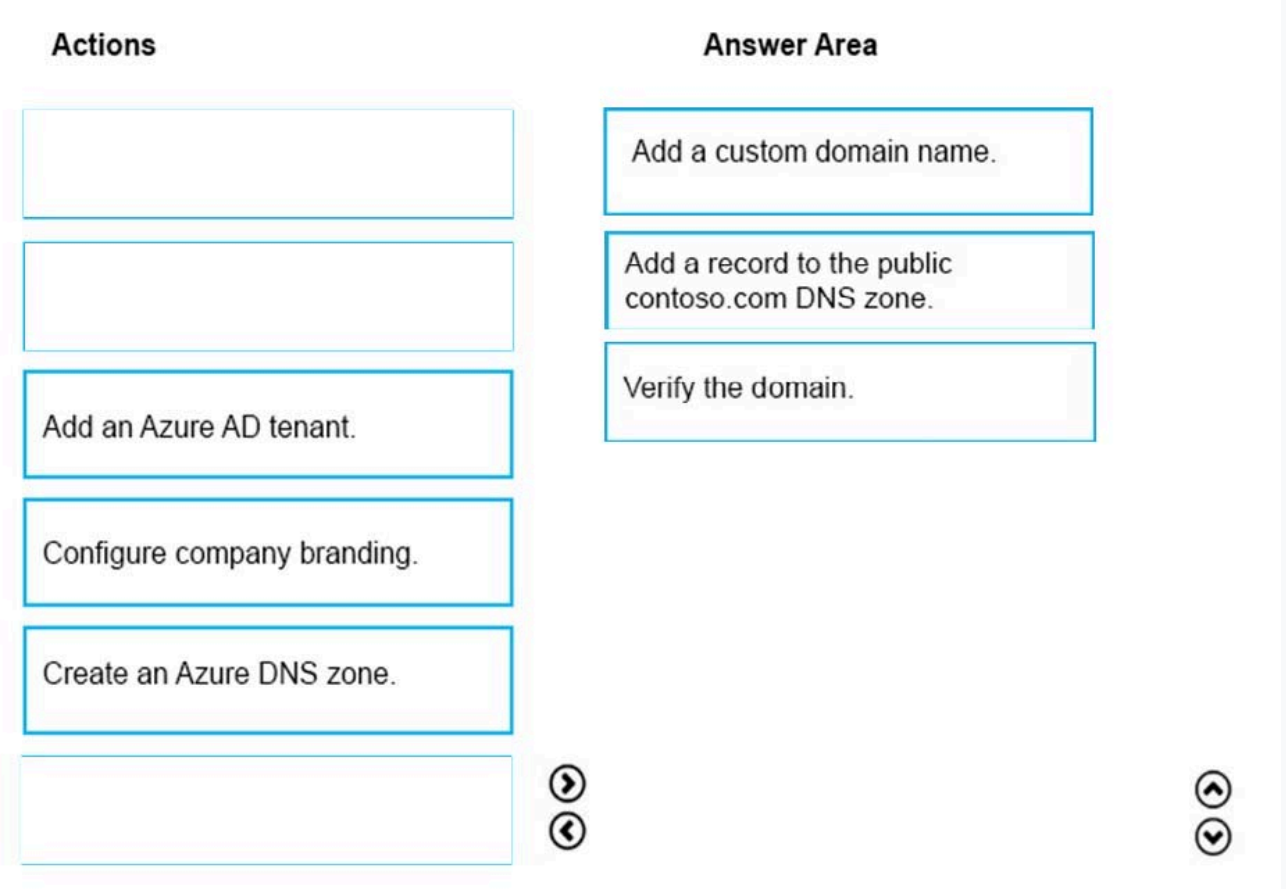

The process is simple:
1. Add the custom domain name to your directory
2. Add a DNS entry for the domain name at the domain name registrar
3. Verify the custom domain name in Azure AD

References:
https://docs.microsoft.com/en-us/azure/dns/dns-web-sites-custom-domain

praveen97 5 years, 4 months ago

Answers are
1. Add the custom domain name to your directory
2. Add a DNS entry for the domain name at the domain name registrar
3. Verify the custom domain name in Azure AD

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

upvoted 2 times

Question #9 *Topic 1*

You plan to migrate an on-premises Hyper-V environment to Azure by using Azure Site Recovery. The Hyper-V environment is managed by using Microsoft
System Center Virtual Machine Manager (VMM).
The Hyper-V environment contains the virtual machines in the following table.

Name	Operating System (OS)	OS disk size	BitLocker Drive Encryption (BitLocker) enabled on OS disks	Generation
DC1	Windows Server 2016	500 GB	No	2
FS1	Ubuntu 16.04 LTS	200 GB	No	2
CA1	Windows Server 2012 R2	1 TB	Yes	1
SQL1	Windows Server 2016	200 GB	No	2

Which virtual machine can be migrated by using Azure Site Recovery?

A. DC1

B. SQL1

C. CA1

D. FS1

Suggested Answer: *B*
Up to 300 GB OS disk size is supported for generation 2 VMs and BitLocker is not enabled.
Incorrect Answers:
A: Only up to 300 GB OS disk size is supported for generation 2 VMs.
C: BitLocker must be disabled before you enable replication for a VM.
D: Linux Generation 2 VMs aren't supported.
References:
https://docs.microsoft.com/en-us/azure/site-recovery/hyper-v-azure-support-matrix#azure-vm-requirements

Currently there are no comments in this discussion, be the first to comment!

Question #10 *Topic 1*

DRAG DROP -

You have an Azure subscription that contains the following resources:

- a virtual network named VNet1
- a replication policy named ReplPolicy1
- a Recovery Services vault named Vault1
- an Azure Storage account named Storage1

You have an Amazon Web Services (AWS) EC2 virtual machine named VM1 that runs Windows Server 2016.

You need to migrate VM1 to VNet1 by using Azure Site Recovery.

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:

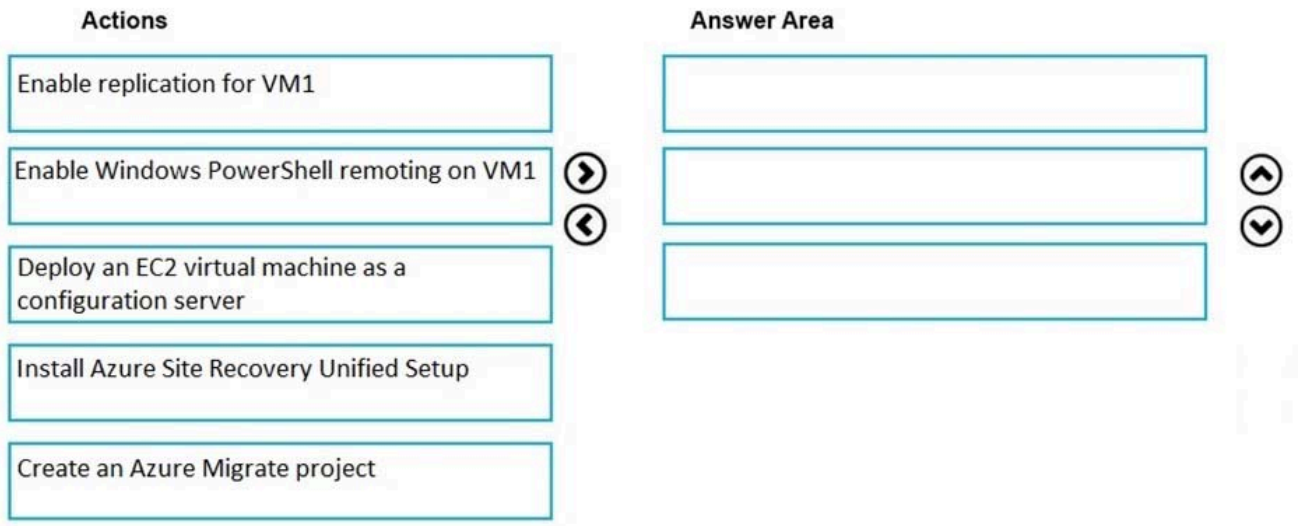

Suggested Answer:

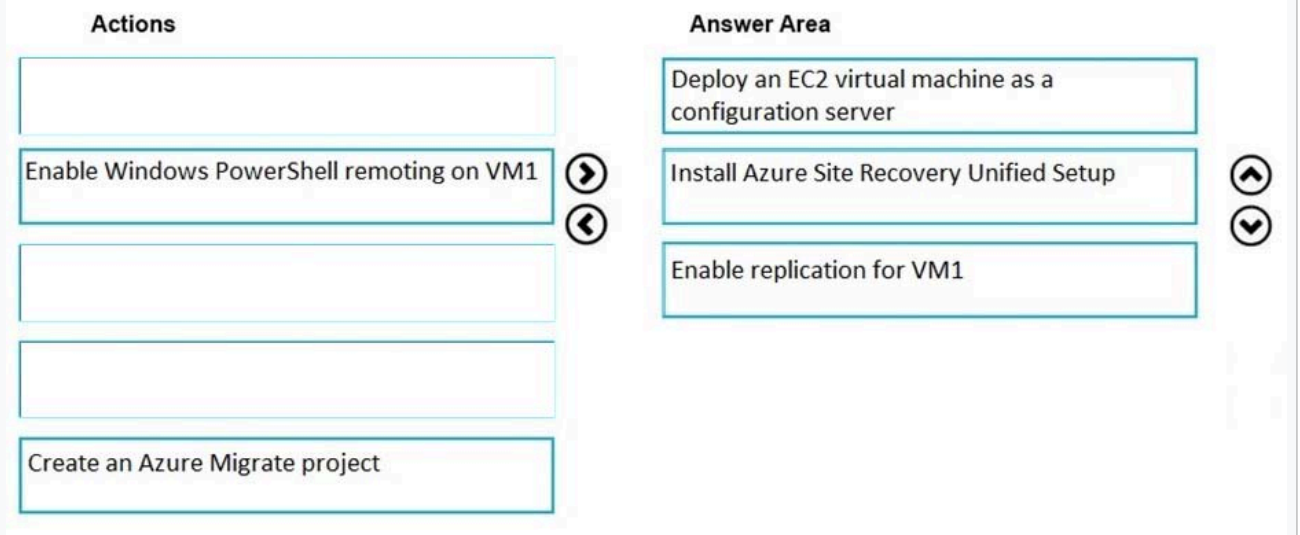

Step 1: Deploy an EC2 virtual machine as a configuration server

Use an EC2 instance that's running Windows Server 2012 R2 to create a configuration server and register it with your recovery vault.

Step 2: Install Azure Site Recovery Unified Setup

Download Microsoft Azure Site Recovery Unified Setup and install it on the VM.

Step 3: Enable replication for each VM that you want to migrate.

References:

https://docs.microsoft.com/en-us/azure/site-recovery/migrate-tutorial-aws-azure

Currently there are no comments in this discussion, be the first to comment!

Question #11 *Topic 1*

DRAG DROP -

You create a new replication policy in Azure for the physical servers.

You successfully complete the following actions:

- Create and configure a Recovery Services vault.
- Ensure Internet connectivity.
- Ensure that the required URLs are reachable.
- Ensure that the host server requirements are met.
- Ensure that the servers marked for replication comply with the requirements of the Azure virtual machines.

You need to replicate the on-premises servers to Azure.

Which four 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

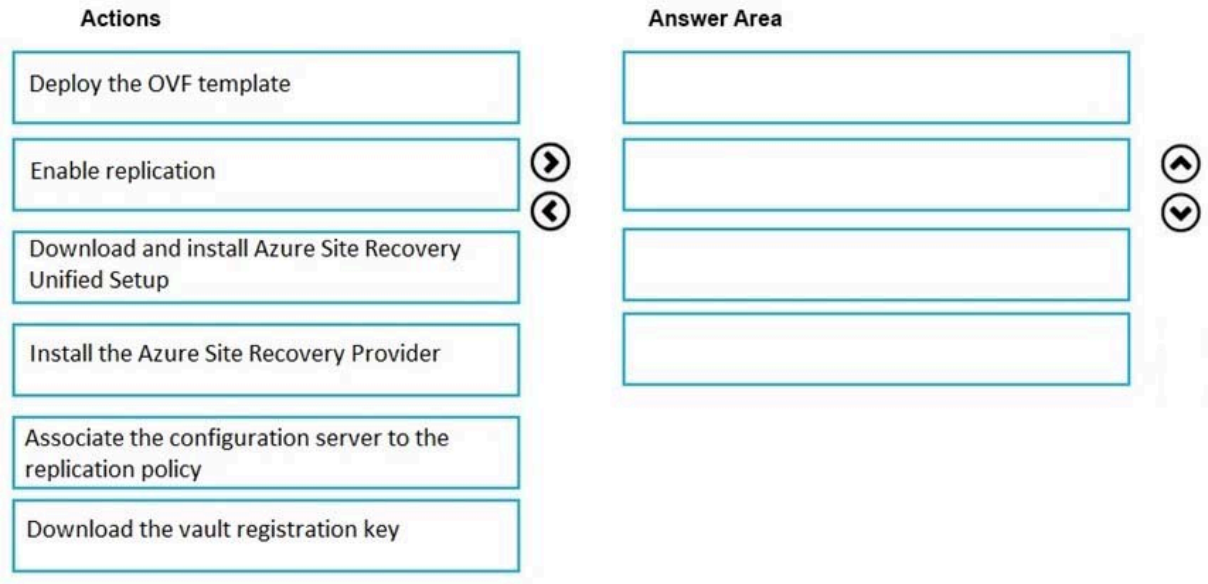

Suggested Answer:

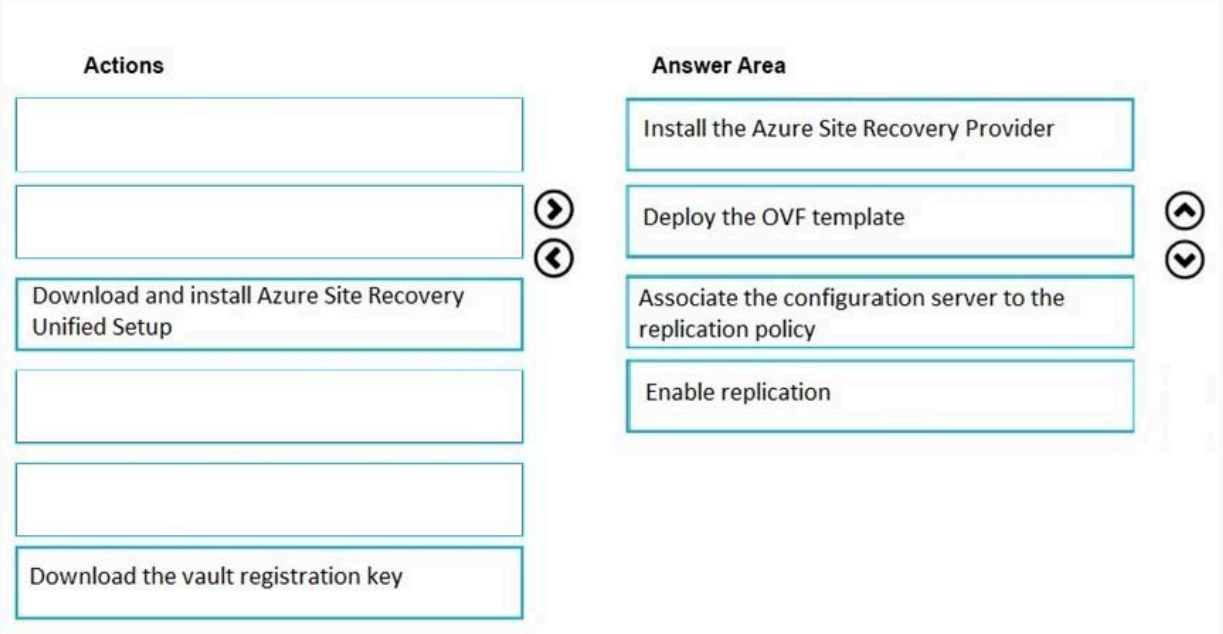

Step 1: Install the Azure Site Recovery Provider

Step 2: Deploy the OVF template -

Set up the source environment.

Download the OVF template for the configuration server, and import the template in VMware.

Note: Open Virtualization Format (OVF) template is an industry standard software distribution model for virtual machine templates. Starting January 2018, configuration server for the VMware to Azure scenario will be available to all our customers as an OVF template.
Step 3: Associate the configuration server to the replication policy
Associate the replication policy with your on-premises configuration server.

Step 4: Enable replication -
References:
https://docs.microsoft.com/en-us/azure/site-recovery/vmware-azure-set-up-replication

Currently there are no comments in this discussion, be the first to comment!

Question #12 *Topic 1*

You plan to move services from your on-premises network to Azure.

You identify several virtual machines that you believe can be hosted in Azure. The virtual machines are shown in the following table.

Name	Role	Operating system (OS)	Environment
Sea-DC01	Domain controller	Windows Server 2016	Hyper-V on Windows Server 2016
NYC-FS01	File server	Windows Server 2012 R2	VMware vCenter Server 5.1
BOS-DB01	Microsoft SQL server	Windows Server 2016	VMware vCenter Server 6
Sea-CA01	Certification authority (CA)	Windows Server 2012 R2	Hyper-V on Windows Server 2016
Hou-NW01	DHCP/DNS	Windows Server 2008 R2	VMware vCenter Server 5.5

Which two virtual machines can you access by using Azure migrate? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Sea-CA01

B. Hou-NW01

C. NYC-FS01

D. Sea-DC01

E. BOS-DB01

Suggested Answer: *CE*

The VMware VMs must be managed by vCenter Server (version 5.5, 6.0, 6.5 or 6.7).

References:

https://docs.microsoft.com/en-us/azure/migrate/migrate-overview

Amir1909 1 year, 8 months ago

B and E is correct

upvoted 2 times

haim 5 years, 8 months ago

can't be B and E vcenter 5.1 is not supported.

upvoted 3 times

jmyers43 5 years, 9 months ago

Right explanation, wrong answer. Should be B,E

upvoted 4 times

Question #1 *Topic 2*

HOTSPOT -

You have an Azure subscription named Subscription1.

In Subscription1, you create an alert rule named Alert1.

The Alert1 action group is configured as shown in the following exhibit.

```
PS Azure:\> Get-AzureRmActionGroup

ResourceGroupName: default-activitylogalerts
GroupShortName    : AG1
Enabled           : True
EmailReceivers    : {Action1_-EmailAction-}
SmsReceivers      : {Acrtion_-SMSAction-}
WebhookReceivers  : {}
Id                : /subscriptions/a4fde29b-d56a-4f6c-8298-6c53cd0b720c/
resourceGroups/default-activitylogalerts/providers/microsoft.insights/actionGrooups/ActionGroup1
Name              : ActionGroup1
Type              : Microsoft.Insights/ActionGroups
Location          : Global
Tags              : {}
```

Alert1 alert criteria is triggered every minute.

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.

Hot Area:

Answer Area

The number of email messages that Alert1 will send in an hour is **[answer choice]**.

- 0
- 4
- 6
- 12
- 60

The number of SMS messages that Alert1 will send in an hour is **[answer choice]**.

- 0
- 4
- 6
- 12
- 60

Suggested Answer:

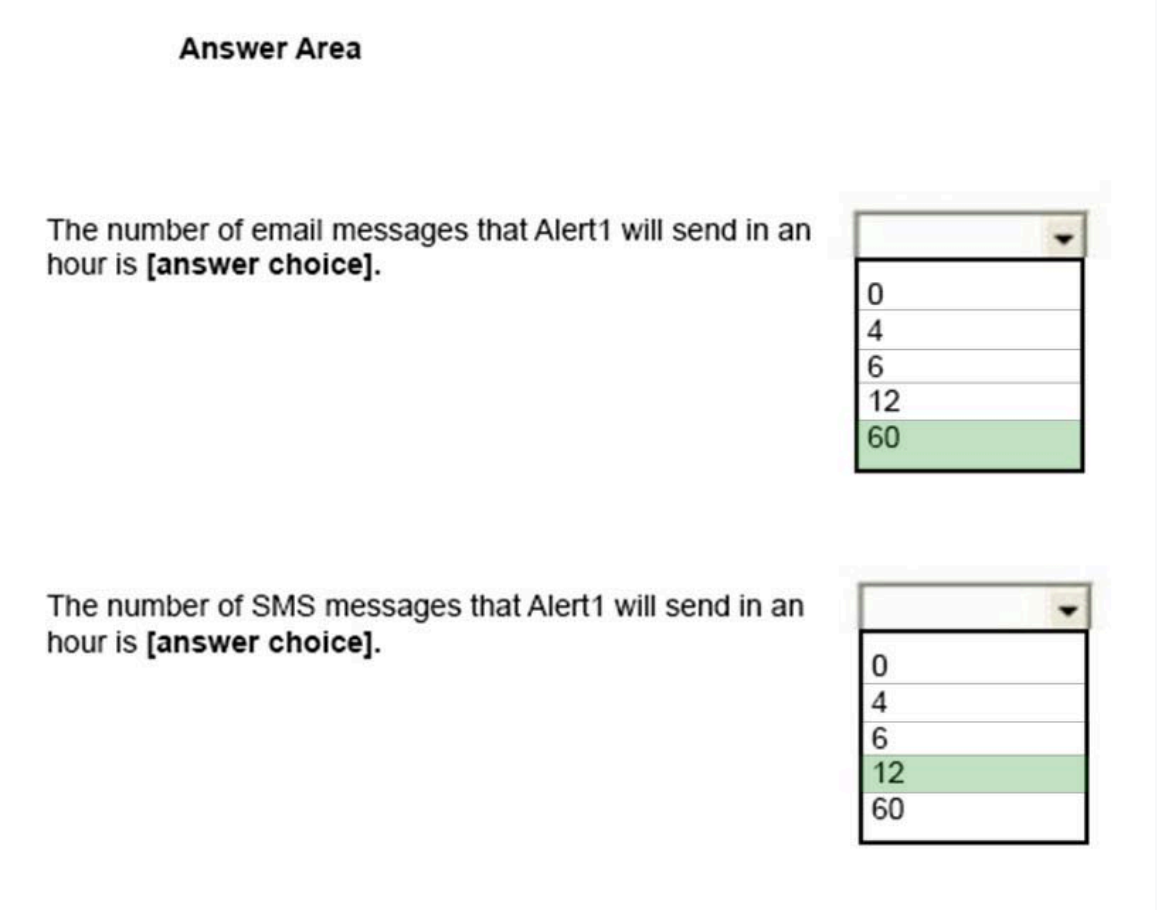

Box 1: 60 -

One alert per minute will trigger one email per minute.

Box 2: 12 -

No more than 1 SMS every 5 minutes can be send, which equals 12 per hour.

Note: Rate limiting is a suspension of notifications that occurs when too many are sent to a particular phone number, email address or device.

Rate limiting ensures that alerts are manageable and actionable.

The rate limit thresholds are:

- SMS: No more than 1 SMS every 5 minutes.
- Voice: No more than 1 Voice call every 5 minutes.
- Email: No more than 100 emails in an hour.
- Other actions are not rate limited.

References:

https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/monitoring-and-diagnostics/monitoring-overview-alerts.md

Currently there are no comments in this discussion, be the first to comment!

Question #2 *Topic 2*

DRAG DROP -

You have an on-premises file server named Server1 that runs Windows Server 2016.

You have an Azure subscription that contains an Azure file share.

You deploy an Azure File Sync Storage Sync Service, and you create a sync group.

You need to synchronize files from Server1 to Azure.

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:

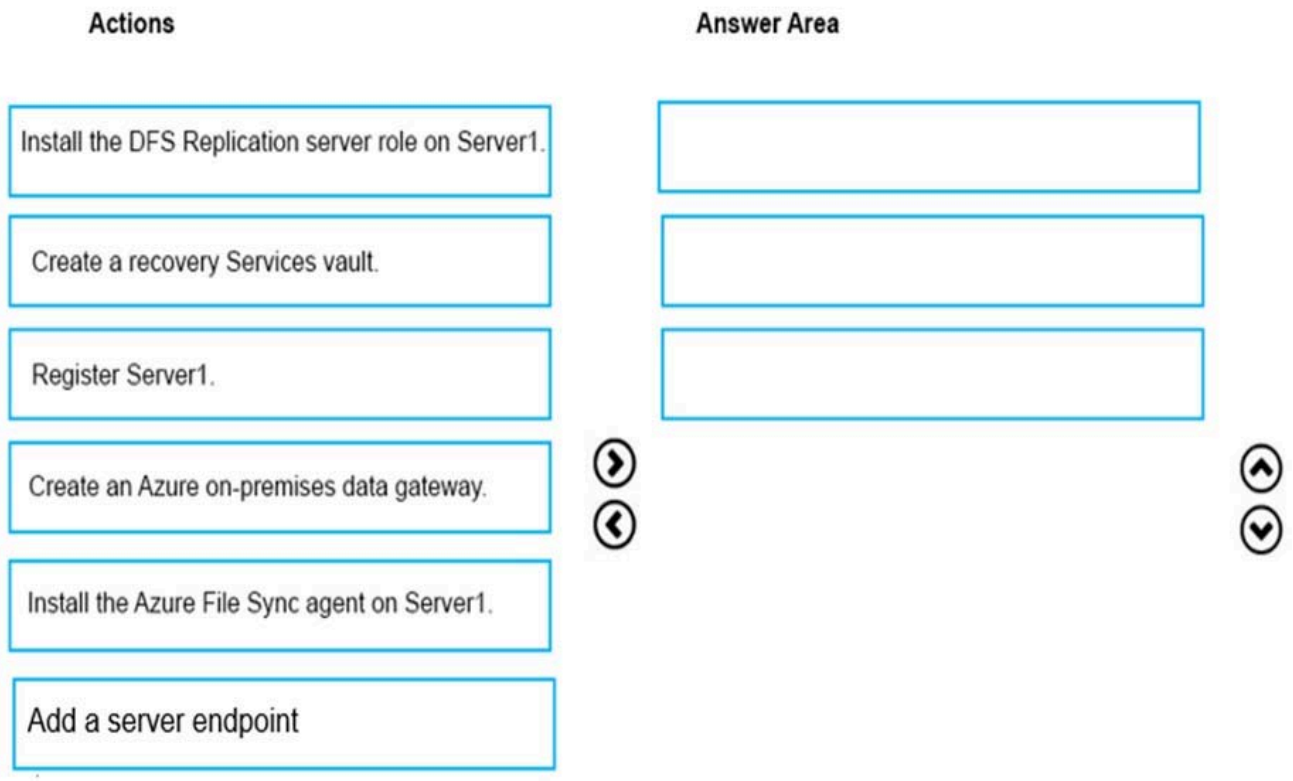

Suggested Answer:

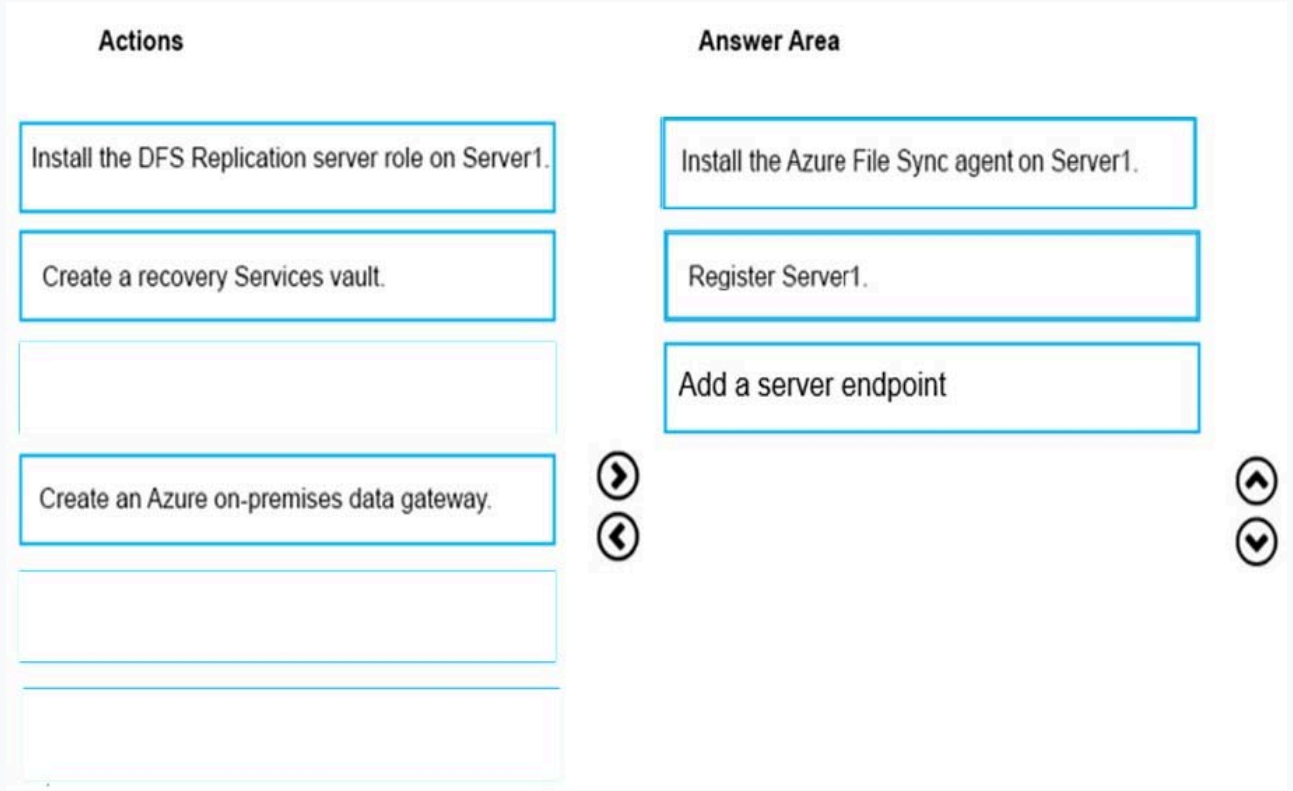

Step 1: Install the Azure File Sync agent on Server1

The Azure File Sync agent is a downloadable package that enables Windows Server to be synced with an Azure file share

Step 2: Register Server1.
Register Windows Server with Storage Sync Service
Registering your Windows Server with a Storage Sync Service establishes a trust relationship between your server (or cluster) and the Storage Sync Service.

Step 3: Add a server endpoint -
Create a sync group and a cloud endpoint.
A sync group defines the sync topology for a set of files. Endpoints within a sync group are kept in sync with each other. A sync group must contain one cloud endpoint, which represents an Azure file share and one or more server endpoints. A server endpoint represents a path on registered server.
References:
https://docs.microsoft.com/en-us/azure/storage/files/storage-sync-files-deployment-guide

Ghanekar 5 years, 1 month ago
correct answer
upvoted 1 times

Question #3 *Topic 2*

You are building a custom Azure function app to connect to Azure Event Grid.
You need to ensure that resources are allocated dynamically to the function app. Billing must be based on the executions of the app.
What should you configure when you create the function app?

A. the Windows operating system and the App Service plan hosting plan

B. the Docker container and an App Service plan that uses the B1 pricing tier

C. the Docker container and an App Service plan that uses the S1 pricing tier

D. the Windows operating system and the Consumption plan hosting plan

Suggested Answer: *D*
Azure Functions runs in two different modes: Consumption plan and Azure App Service plan. The Consumption plan automatically allocates compute power when your code is running. Your app is scaled out when needed to handle load, and scaled down when code is not running.
Incorrect Answers:
A: When you run in an App Service plan, you must manage the scaling of your function app.
References:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-first-azure-function

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 2*

DRAG DROP -

You have an Azure subscription named Subscription1.

You create an Azure Storage account named contosostorage, and then you create a file share named data.

Which UNC path should you include in a script that references files from the data file share? To answer, drag the appropriate values to the correct targets. Each value 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.

Select and Place:

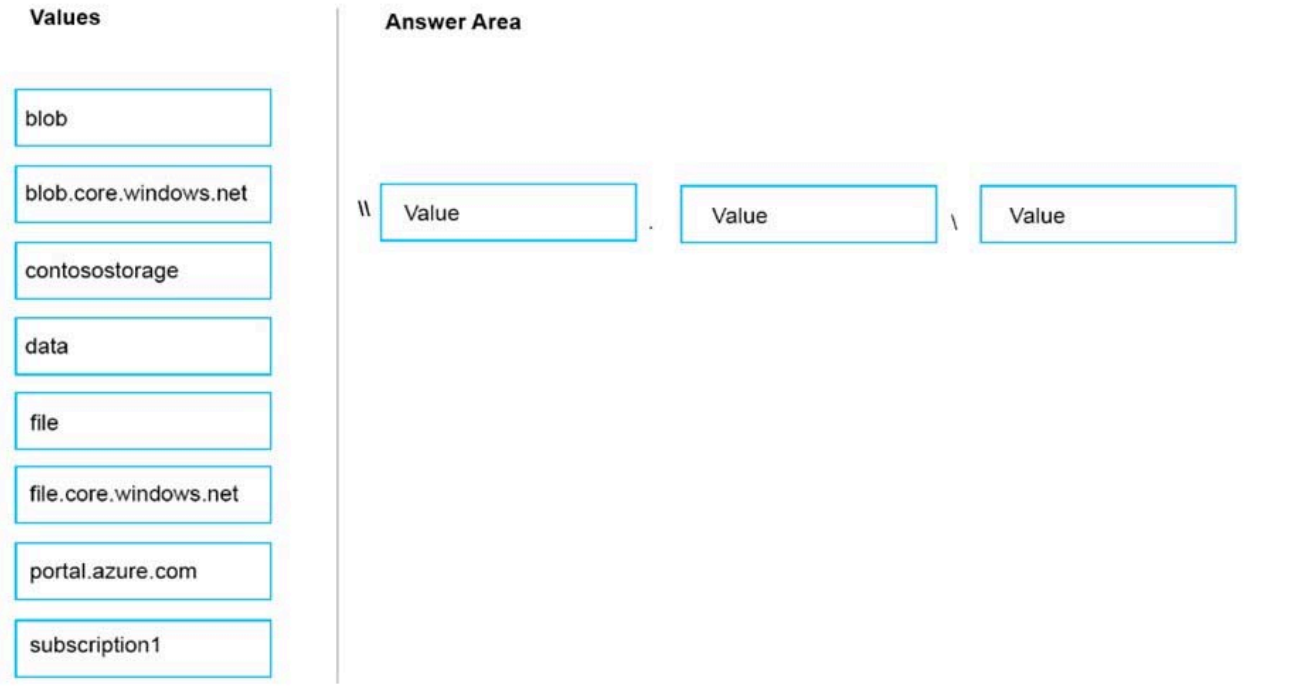

Suggested Answer:

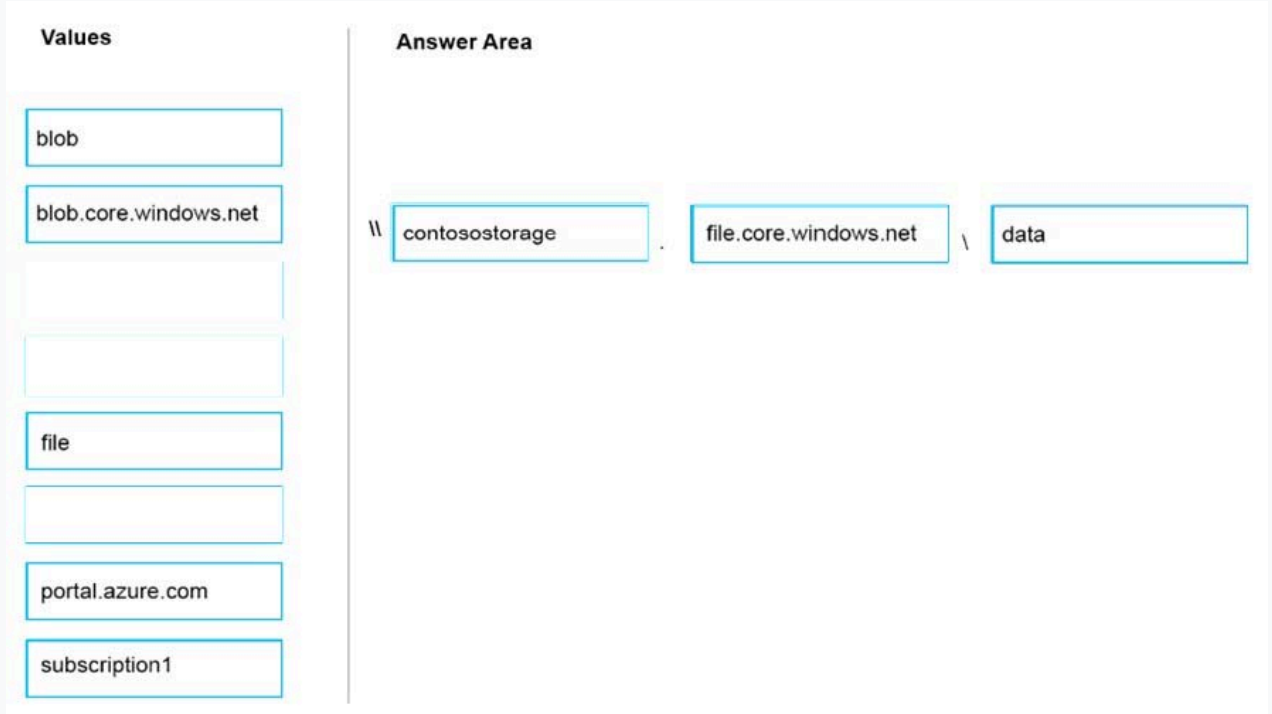

Box 1: contosostorage -

The name of account -

Box 2: file.core.windows.net -

Box 3: data -

The name of the file share is data.

Example:

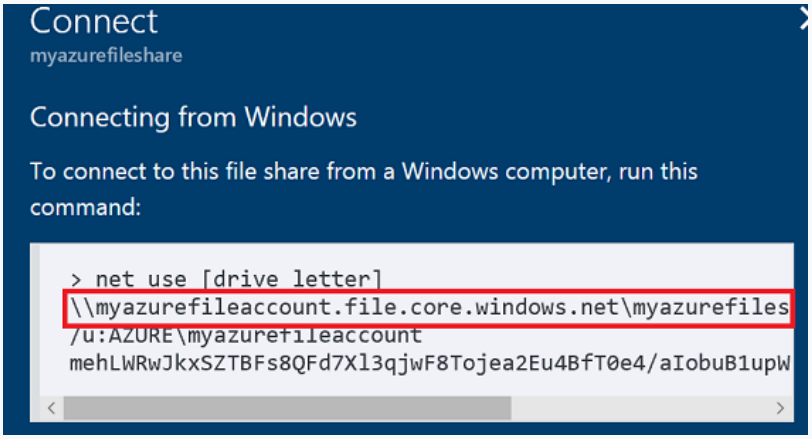

References:

https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-use-files-windows

Currently there are no comments in this discussion, be the first to comment!

Question #5 *Topic 2*

You have an Azure Service Bus.
You need to implement a Service Bus queue that guarantees first-in-first-out (FIFO) delivery messages.
What should you do?

A. Set the Max Size setting of the queue to 5 GB.

B. Enable duplicate detection.

C. Set the Lock Duration setting to 10 seconds.

D. Enable sessions.

E. Enable partitioning.

Suggested Answer: *D*
Through the use of messaging sessions you can guarantee ordering of messages, that is first-in-first-out (FIFO) delivery of messages.
References:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

Currently there are no comments in this discussion, be the first to comment!

Question #6 *Topic 2*

You are troubleshooting a performance issue for an Azure Application Gateway.
You need to compare the total requests to the failed requests during the past six hours.
What should you use?

A. Metrics in Application Gateway

B. Diagnostics logs in Application Gateway

C. NSG flow logs in Azure Network Watcher

D. Connection monitor in Azure Network Watcher

Suggested Answer: *A*
Application Gateway currently has seven metrics to view performance counters.
Metrics are a feature for certain Azure resources where you can view performance counters in the portal. For Application Gateway, the following metrics are available:
⇨ Total Requests
⇨ Failed Requests
⇨ Current Connections
⇨ Healthy Host Count
⇨ Response Status
⇨ Throughput
⇨ Unhealthy Host count
You can filter on a per backend pool basis to show healthy/unhealthy hosts in a specific backend pool
References:
https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-diagnostics#metrics

Currently there are no comments in this discussion, be the first to comment!

Question #7 *Topic 2*

DRAG DROP -

You have an Azure subscription that contains an Azure Service Bus named Bus1.

Your company plans to deploy two Azure web apps named App1 and App2. The web apps will create messages that have the following requirements:

- Each message created by App1 must be consumed by only a single consumer.
- Each message created by App2 will be consumed by multiple consumers.

Which resource should you create for each web app? To answer, drag the appropriate resources to the correct web apps. Each resource 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:

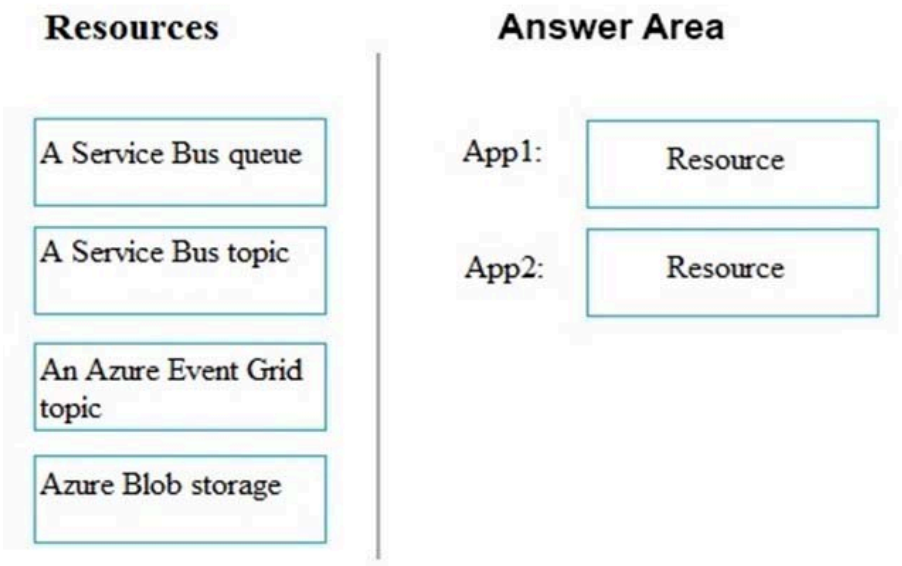

Suggested Answer:

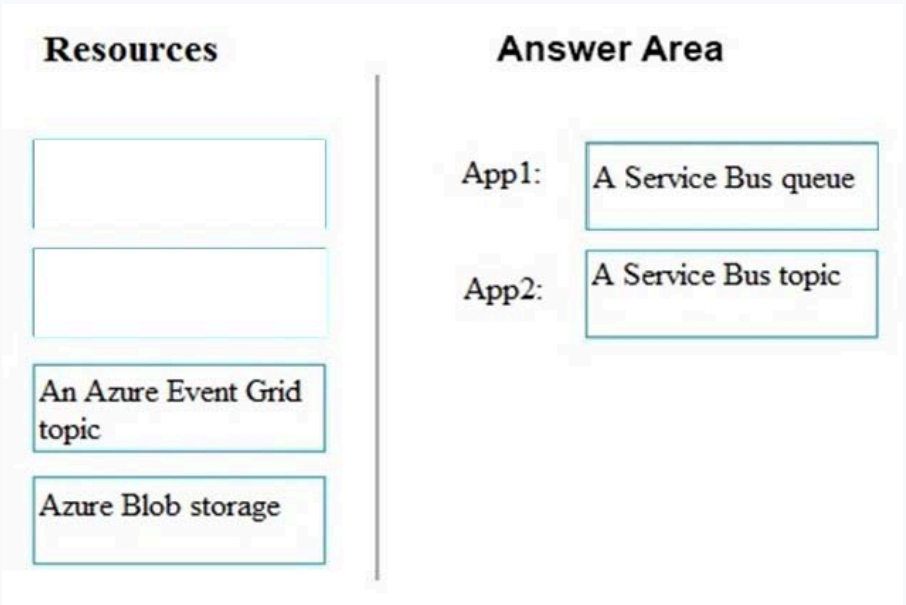

App1: A Service Bus queue -

App2: A Service Bus topic -

In contrast to queues, in which each message is processed by a single consumer, topics and subscriptions provide a one-to-many form of communication, in a publish/subscribe pattern.

References:

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

Currently there are no comments in this discussion, be the first to comment!

Question #8 *Topic 2*

HOTSPOT -

You have an Azure Service Bus.

You create a queue named Queue1. Queue1 is configured as shown in the following exhibit.

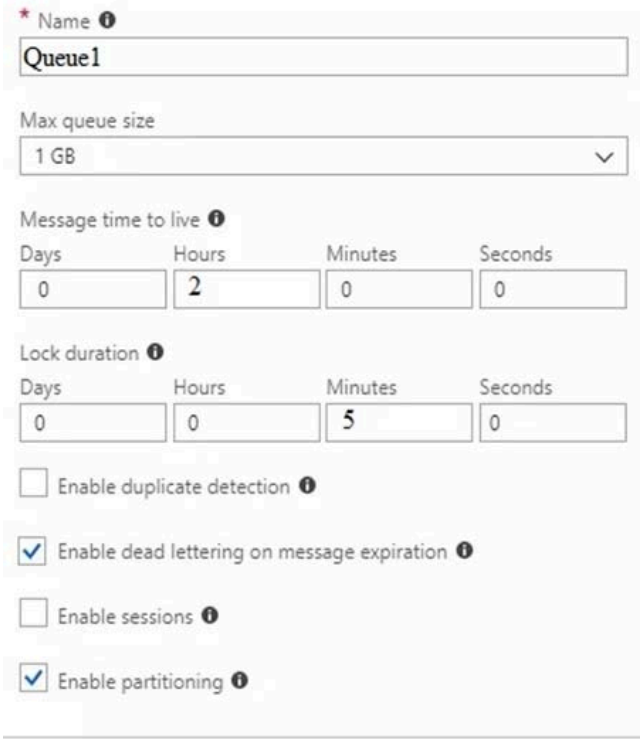

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

If a message that has a TTL of four hours is written to Queue1 and is never read, the message will be **[answer choice]**.

- deleted after two hours
- deleted after four hours
- deleted after two hours and five minutes
- retained until manually deleted

If a message that has a TTL of two hours is written to Queue1, and then read after one hour, the message will be **[answer choice]**.

- deleted immediately
- deleted in five minutes
- deleted in one hour
- retained until manually deleted

Suggested Answer:

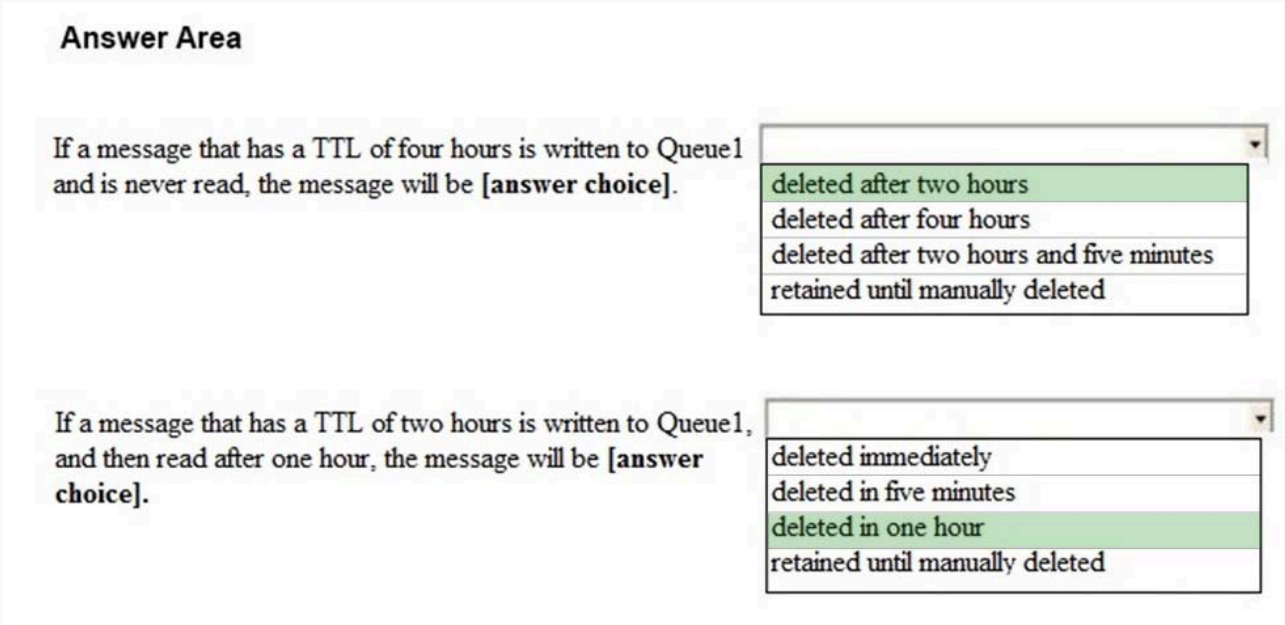

Box 1: deleted after two hours -

All messages sent into a queue or topic are subject to a default expiration that is set at the entity level with the defaultMessageTimeToLive property and which can also be set in the portal during creation and adjusted later. The default expiration is used for all messages sent to the entity where TimeToLive is not explicitly set.

The default expiration also functions as a ceiling for the TimeToLive value. Messages that have a longer TimeToLive expiration than the default value are silently adjusted to the defaultMessageTimeToLive value before being enqueued.

Box 2: deleted in one hour -

References:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/message-expiration

Chinny 5 years ago

Box 1.....retained until manually deleted

Box 2.....deleted immediately

upvoted 1 times

vrana 5 years, 6 months ago

why Box 2 answers is deleted in one hour . Should not it be deleted after it is read.

upvoted 1 times

Question #9 *Topic 2*

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

Name	Type
RG1	Resource group
Store1	Azure Storage account
Sync1	Azure File Sync

Store1 contains a file share named Data. Data contains 5,000 files.

You need to synchronize the files in Data to an on-premises server named Server1.

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

NOTE: Each correct selection is worth one point.

A. Install the Azure File Sync agent on Server1.

B. Create a container instance.

C. Download an automation script.

D. Create a sync group.

E. Register Server1.

Suggested Answer: *ADE*

Azure synchronize files on-premises agent

A: The Azure File Sync agent is a downloadable package that enables Windows Server to be synced with an Azure file share.

D: A sync group defines the sync topology for a set of files. Endpoints within a sync group are kept in sync with each other. A sync group must contain one cloud endpoint, which represents an Azure file share and one or more server endpoints. A server endpoint represents a path on registered server. A server can have server endpoints in multiple sync groups. You can create as many sync groups as you need to appropriately describe your desired sync topology.

E: Registering your Windows Server with a Storage Sync Service establishes a trust relationship between your server (or cluster) and the Storage Sync Service. A server can only be registered to one Storage Sync Service and can sync with other servers and Azure file shares associated with the same Storage Sync Service.

Incorrect Answers:

B: The Resource group is used in a container.

References:

https://docs.microsoft.com/en-us/azure/storage/files/storage-sync-files-deployment-guide

arol 5 years, 5 months ago

Answer is correct

Sequence

1. Create sync Group
2. Install Azure file sync agent
3. Register server

upvoted 1 times

Question #10 *Topic 2*

You have an Azure Logic App named App1. App1 provides a response when an HTTP POST request or an HTTP GET request is received.
During peak periods, App1 is expected to receive up to 200,000 requests in a five-minute period.
You need to ensure that App1 can handle the expected load.
What should you configure?

A. Access keys

B. Access control (IAM)

C. API connections

D. Workflow settings

Suggested Answer: *D*

Here are the limits for a single logic app definition:

Name	Limit	Notes
Actions per workflow	500	To extend this limit, you can add nested workflows as needed.
Allowed nesting depth for actions	8	To extend this limit, you can add nested workflows as needed.
Workflows per region per subscription	1,000	
Triggers per workflow	10	When working in code view, not the designer
Switch scope cases limit	25	
Variables per workflow	250	
Characters per expression	8,192	
Maximum size for `trackedProperties`	16,000 characters	
Name for `action` or `trigger`	80 characters	
Length of `description`	256 characters	
Maximum `parameters`	50	
Maximum `outputs`	10	

References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-limits-and-config

Currently there are no comments in this discussion, be the first to comment!

Question #11 *Topic 2*

DRAG DROP -

You are configuring serverless computing in Azure.

You need to receive an email message whenever a resource is created in or deleted from a resource group.

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:

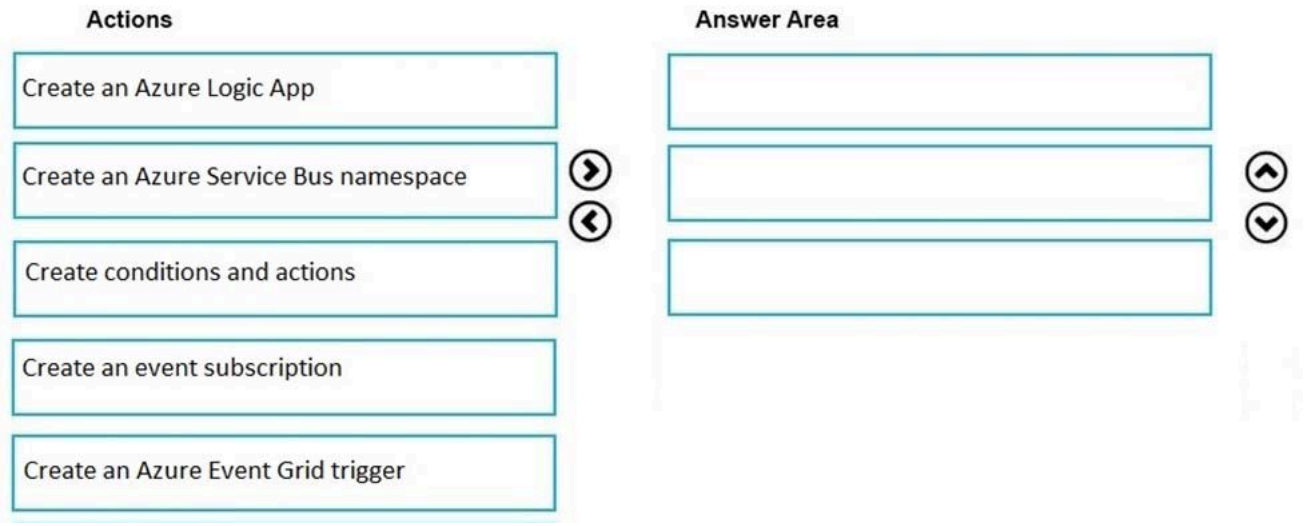

Suggested Answer:

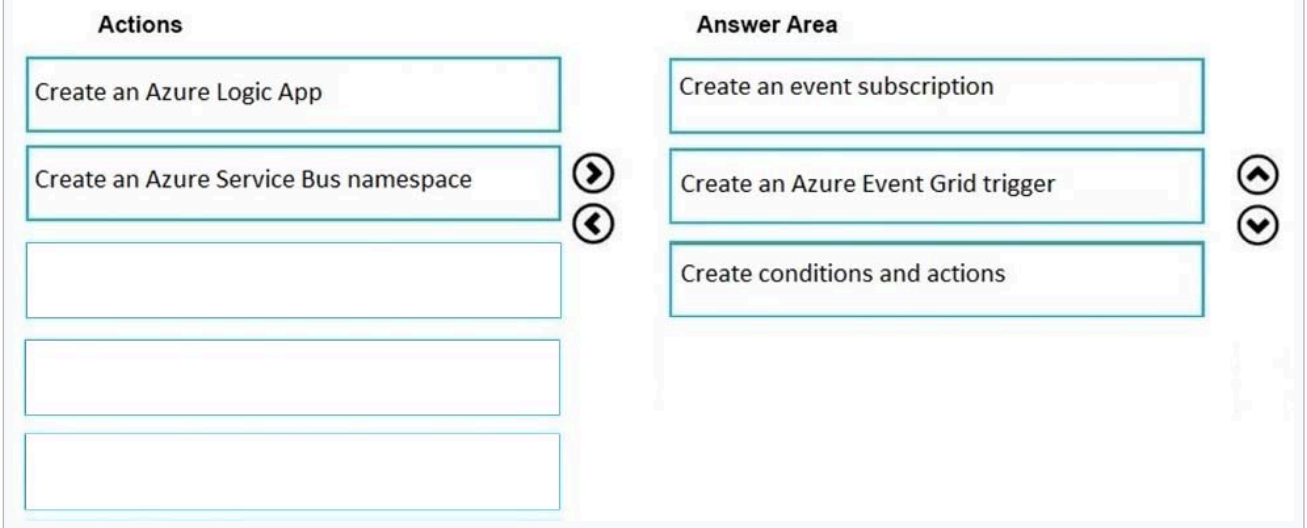

Step 1: Create an event subscription

When you subscribe to events for a resource group, your endpoint receives all events for that resource group.

Step 2: Create an Azure Event Grid trigger

Step 3: Create conditions and actions

References:

https://docs.microsoft.com/en-us/azure/event-grid/event-schema-resource-groups

Currently there are no comments in this discussion, be the first to comment!

Question #12 *Topic 2*

DRAG DROP -

You have an Azure subscription that contains an Azure file share.

You have an on-premises server named Server1 that runs Windows Server 2016.

You plan to set up Azure File Sync between Server1 and the Azure file share.

You need to prepare the subscription for the planned Azure File Sync.

Which two actions should you perform in the Azure subscription? To answer, drag the appropriate actions to the correct targets. Each action 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.

Select and Place:

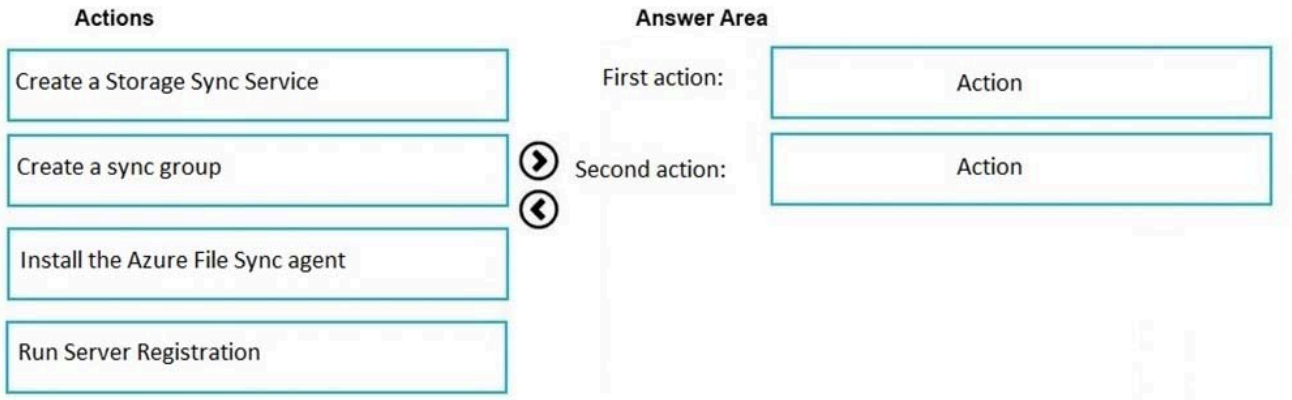

Suggested Answer:

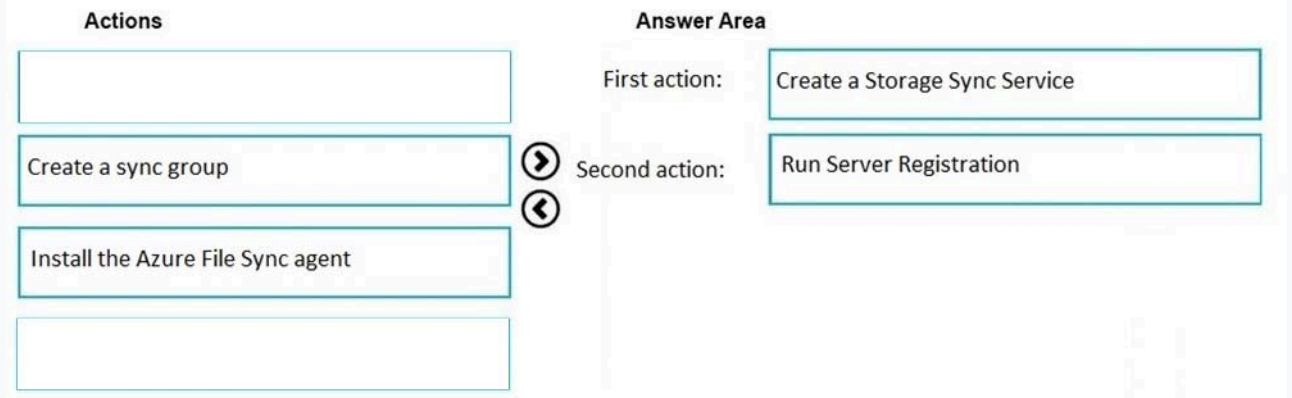

First action: Create a Storage Sync Service

The deployment of Azure File Sync starts with placing a Storage Sync Service resource into a resource group of your selected subscription.

Second action: Run Server Registration

Registering your Windows Server with a Storage Sync Service establishes a trust relationship between your server (or cluster) and the Storage Sync Service. A server can only be registered to one Storage Sync Service and can sync with other servers and Azure file shares associated with the same Storage Sync Service.

The Server Registration UI should open automatically after installation of the Azure File Sync agent.

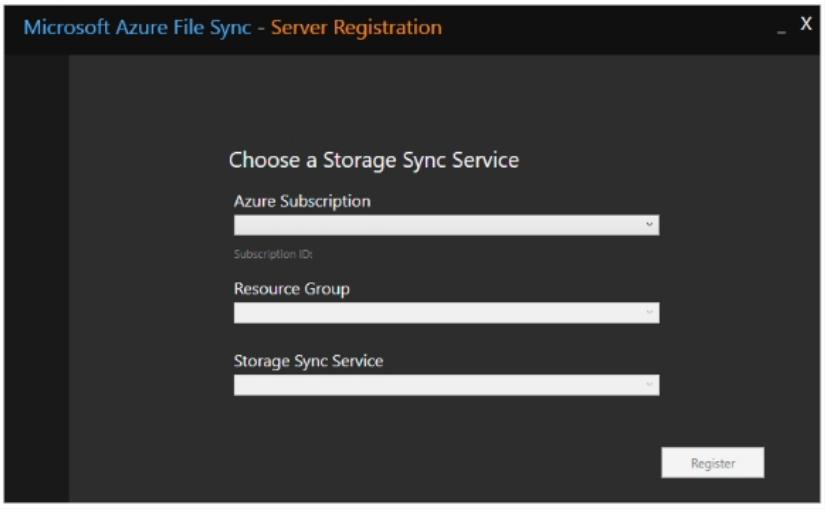

Incorrect Answers:
Not Install the Azure File Sync agent: The Azure File Sync agent is a downloadable package that enables Windows Server to be synced with an Azure file share.

jzaal19 Highly Voted 5 years, 9 months ago

First Action: Create a Storage Sync Service
Second Action: Create a sync group

These are the only actions you perform in the Azure Subscription.
Install Azure File Sync Agent and Server Registration are performed from the server endpoints, which in this case is the on-premise server named Server1.

upvoted 7 times

Question #1 *Topic 3*

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 manage a virtual network named Vnet1 that is hosted in the West US Azure region.
VNet hosts two virtual machines named VM1 and VM2 run Windows Server.
You need to inspect all the network traffic from VM1 to VM2 for a period of three hours.
Solution: From Azure Network Watcher, you create a connection monitor.
Does this meet the goal?

A. Yes

B. No

Suggested Answer: *A*
Azure Network Watcher provides tools to monitor, diagnose, view metrics, and enable or disable logs for resources in an Azure virtual network.
Use the Connection Monitor feature of Azure Network Watcher, which provides you RTT values on a per-minute granularity. You can monitor a direct TCP connection from a virtual machine to a virtual machine, FQDN, URI, or IPv4 address.
References:
https://azure.microsoft.com/en-us/updates/general-availability-azure-network-watcher-connection-monitor-in-all-public-regions/
https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-monitoring-overview

jmyers43 5 years, 9 months ago
Packet Capture
upvoted 1 times

Question #2 *Topic 3*

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 manage a virtual network named Vnet1 that is hosted in the West US Azure region.
VNet hosts two virtual machines named VM1 and VM2 run Windows Server.
You need to inspect all the network traffic from VM1 to VM2 for a period of three hours.
Solution: From Azure Monitor, you create a metric on Network In and Network Out.
Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*
You should use Azure Network Watcher.
References:
https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-monitoring-overview

Currently there are no comments in this discussion, be the first to comment!

Question #3 *Topic 3*

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 manage a virtual network named Vnet1 that is hosted in the West US Azure region.
VNet hosts two virtual machines named VM1 and VM2 run Windows Server.
You need to inspect all the network traffic from VM1 to VM2 for a period of three hours.
Solution: From Performance Monitor, you create a Data Collector Set (DCS).
Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*
You should use Azure Network Watcher.
References:
https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-monitoring-overview

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 3*

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 registers a domain name of contoso.com.
You create an Azure DNS zone named contoso.com, and then you add an A record to the zone for a host named www that has an IP address of 131.107.1.10.
You discover that Internet hosts are unable to resolve www.contoso.com to the 131.107.1.10 IP address.
You need to resolve the name resolution issue.
Solution: You modify the SOA record in the contoso.com zone.
Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*
Modify the Name Server (NS) record.
References:
https://docs.microsoft.com/en-us/azure/dns/dns-delegate-domain-azure-dns

Currently there are no comments in this discussion, be the first to comment!

Question #5 *Topic 3*

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 registers a domain name of contoso.com.
You create an Azure DNS zone named contoso.com, and then you add an A record to the zone for a host named www that has an IP address of 131.107.1.10.
You discover that Internet hosts are unable to resolve www.contoso.com to the 131.107.1.10 IP address.
You need to resolve the name resolution issue.
Solution: You modify the name servers at the domain registrar.
Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*
You should modify the Name Server record, not the Name Servers.
Note: Before you can delegate your DNS zone to Azure DNS, you need to know the name servers for your zone. The NS record set contains the names of the
Azure DNS name servers assigned to the zone.
In the registrar's DNS management page, edit the NS records and replace the NS records with the Azure DNS name servers.
References:
https://docs.microsoft.com/en-us/azure/dns/dns-delegate-domain-azure-dns

btsun 5 years, 6 months ago
The answer is A.
upvoted 2 times

Question #6 *Topic 3*

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 registers a domain name of contoso.com.
You create an Azure DNS zone named contoso.com, and then you add an A record to the zone for a host named www that has an IP address of 131.107.1.10.
You discover that Internet hosts are unable to resolve www.contoso.com to the 131.107.1.10 IP address.
You need to resolve the name resolution issue.
Solution: You create a PTR record for www in the contoso.com zone.
Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*
Modify the Name Server (NS) record.
References:
https://docs.microsoft.com/en-us/azure/dns/dns-delegate-domain-azure-dns

Currently there are no comments in this discussion, be the first to comment!

DRAG DROP -

Your network is configured as shown in the following exhibit.

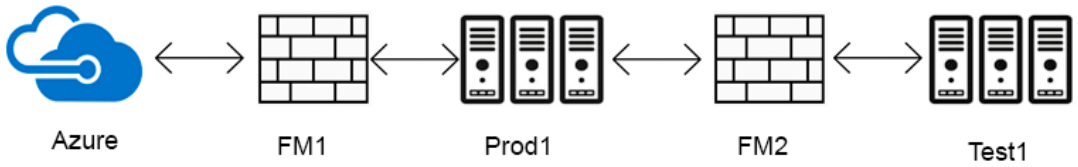

The firewalls are configured as shown in the following table.

Allowed port name	Inbound (TCP)	Outbound (TCP)
FW1	993, 3389	80, 993
FM2	443, 995, 3389	80, 995

Prod1 contains a vCenter server.

You install an Azure Migrate Collector on Test1.

You need to discover the virtual machines.

Which TCP port should be allowed on each firewall? To answer, drag the appropriate ports to the correct firewalls. Each port 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:

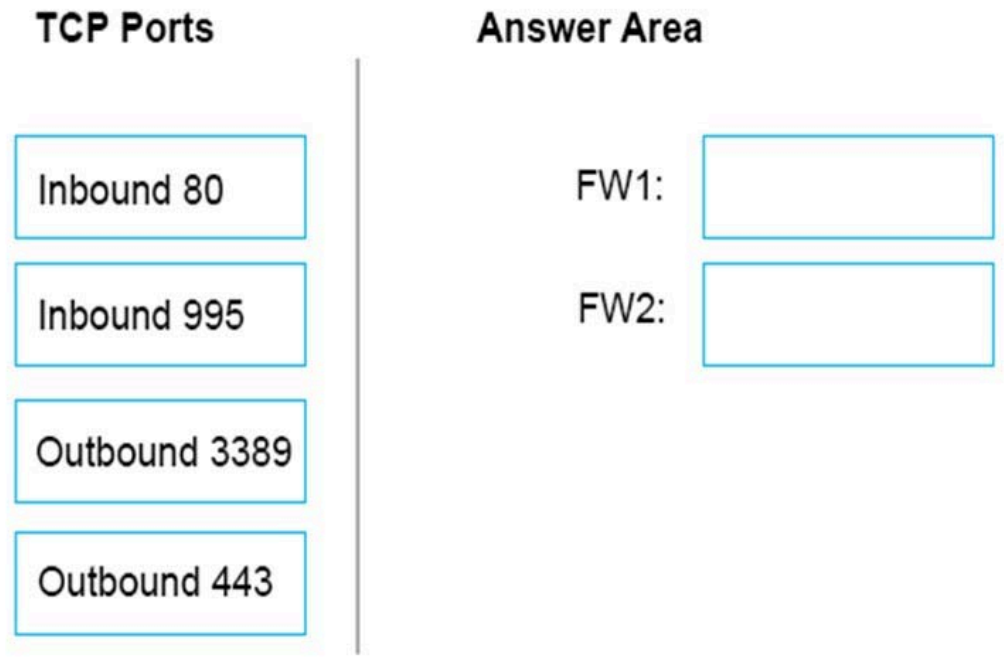

Suggested Answer:

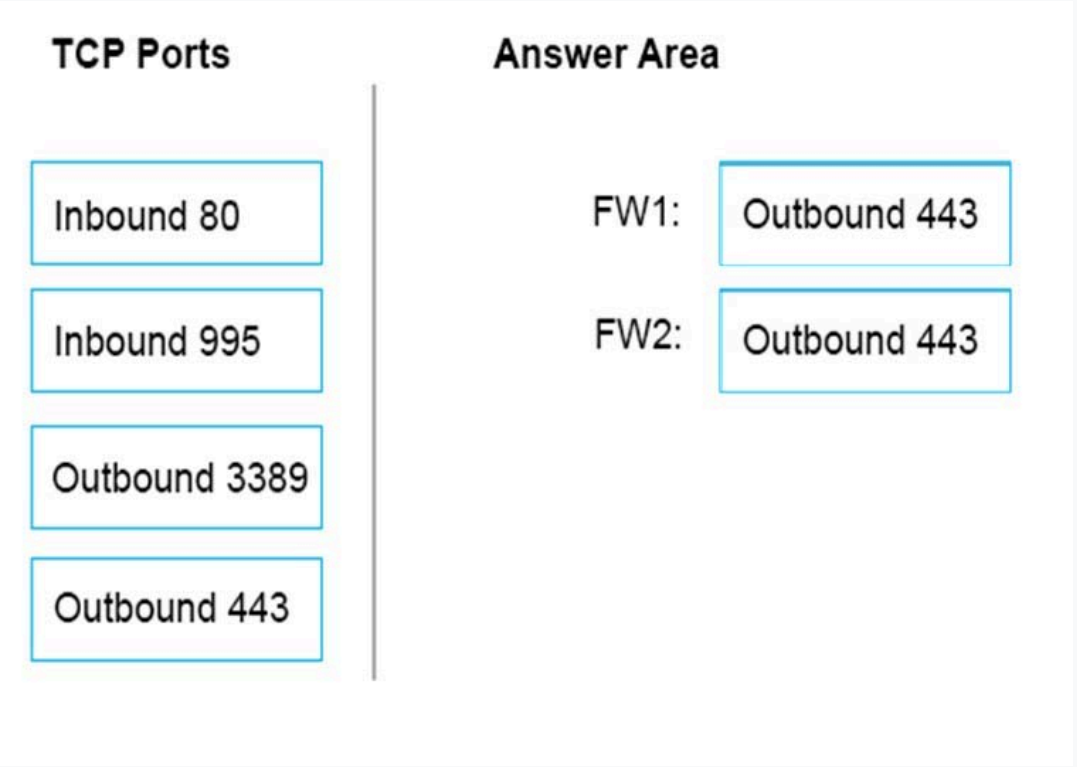

FW1: Outbound 443 -
Collector communicates with Azure Migrate service over SSL 443.

FW2: Outbound 443 -
The Collector must be able to communicate with the vCenter Server. By default, it connects to vCenter on 443.
Note: The collector communicates as summarized in the following diagram.

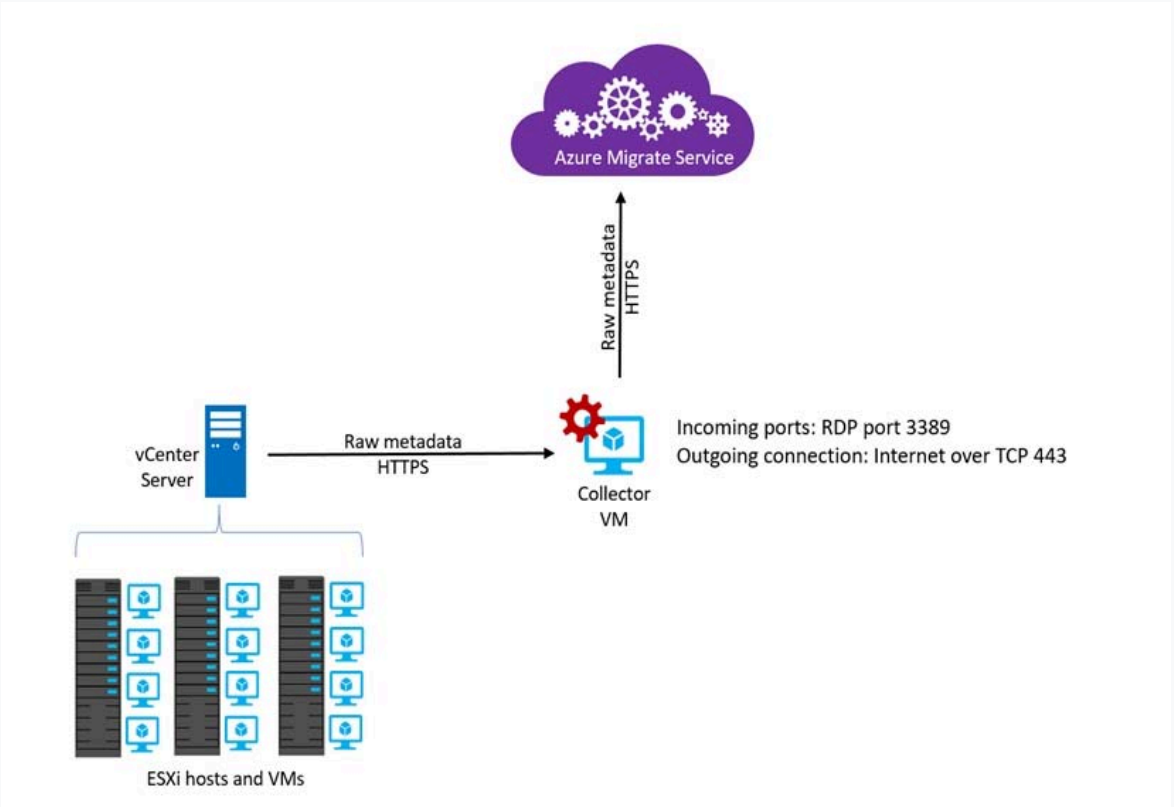

References:
https://docs.microsoft.com/en-us/azure/migrate/concepts-collector

tashakori 1 year, 7 months ago
Given answer is correct
upvoted 1 times

AzMrWhite 5 years, 2 months ago

Explanation: FW1: Outbound 443

Collector communicates with Azure Migrate service over SSL 443. FW2: Outbound 443

The Collector must be able to communicate with the vCenter Server. By default, it connects to vCenter on 443.

upvoted 1 times

deepu1207 5 years, 9 months ago

please send correct answer both were 'outbound 443' port no i think may be' outboud 3389' also one port please explain how

upvoted 1 times

samemee 3 years, 4 months ago

3389 port by default but it can be change. I think 443 is correct in this situation.

upvoted 1 times

Question #8 *Topic 3*

You create an Azure Storage account named contosostorage.
You plan to create a file share named data.
Users need to map a drive to the data file share from home computers that run Windows 10.
Which outbound port should you open between the home computers and the data file share?

A. 80

B. 443

C. 445

D. 3389

Suggested Answer: *C*
Ensure port 445 is open: The SMB protocol requires TCP port 445 to be open; connections will fail if port 445 is blocked.
References:
https://docs.microsoft.com/en-us/azure/storage/files/storage-how-to-use-files-windows

Currently there are no comments in this discussion, be the first to comment!

Question #9 *Topic 3*

DRAG DROP -

You have an on-premises network that includes a Microsoft SQL Server instance named SQL1.

You create an Azure Logic App named App1.

You need to ensure that App1 can query a database on SQL1.

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:

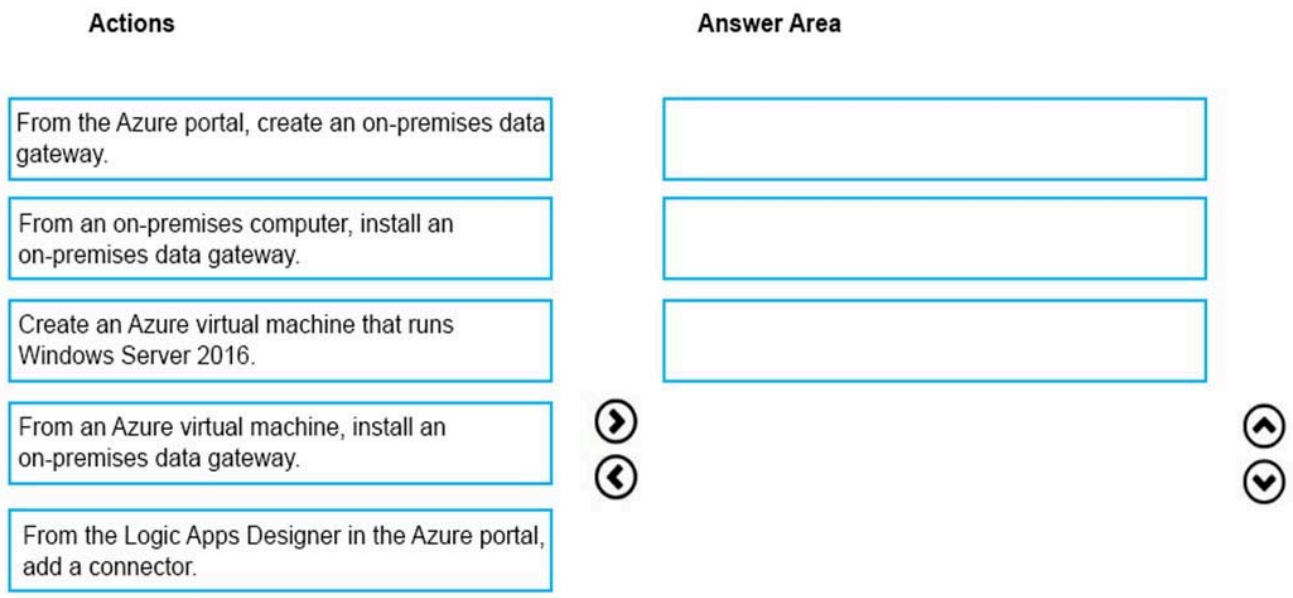

Suggested Answer:

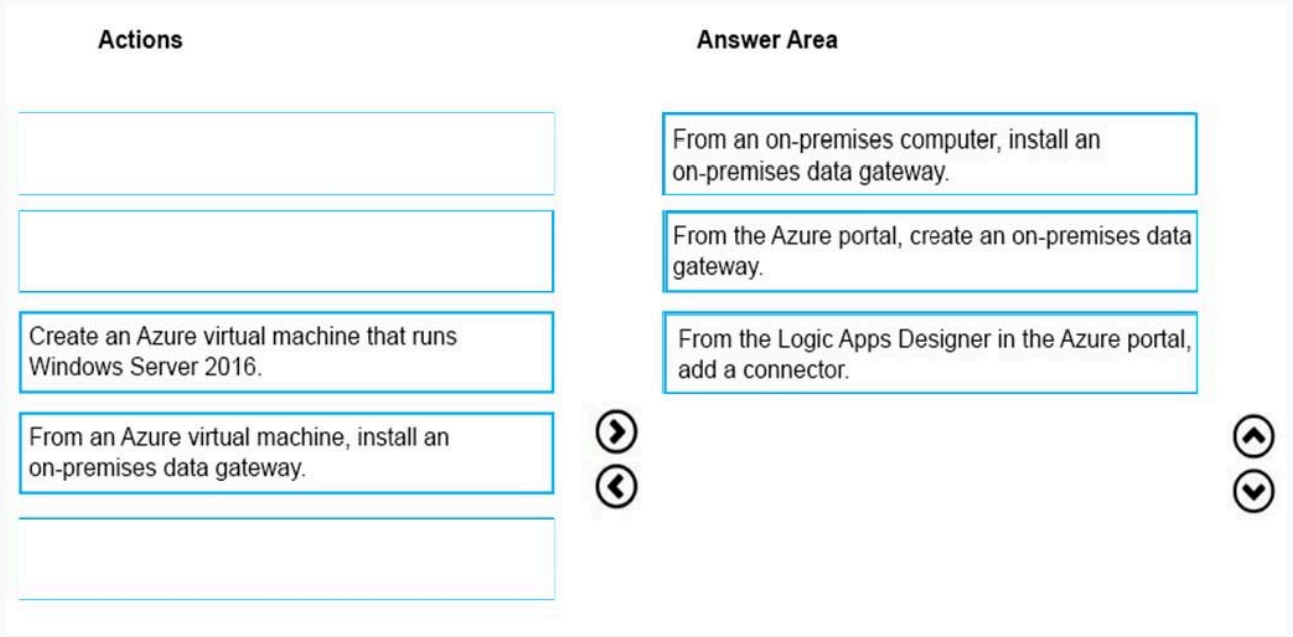

To access data sources on premises from your logic apps, you can create a data gateway resource in Azure so that your logic apps can use the on-premises connectors.

Box 1: From an on-premises computer, install an on-premises data gateway.

Before you can connect to on-premises data sources from Azure Logic Apps, download and install the on-premises data gateway on a local computer.

Box 2: From the Azure portal, create an on-premises data gateway

Create Azure resource for gateway

After you install the gateway on a local computer, you can then create an Azure resource for your gateway. This step also associates your gateway resource with your Azure subscription.

1. Sign in to the Azure portal. Make sure you use the same Azure work or school email address used to install the gateway.
2. On the main Azure menu, select Create a resource > Integration > On-premises data gateway.

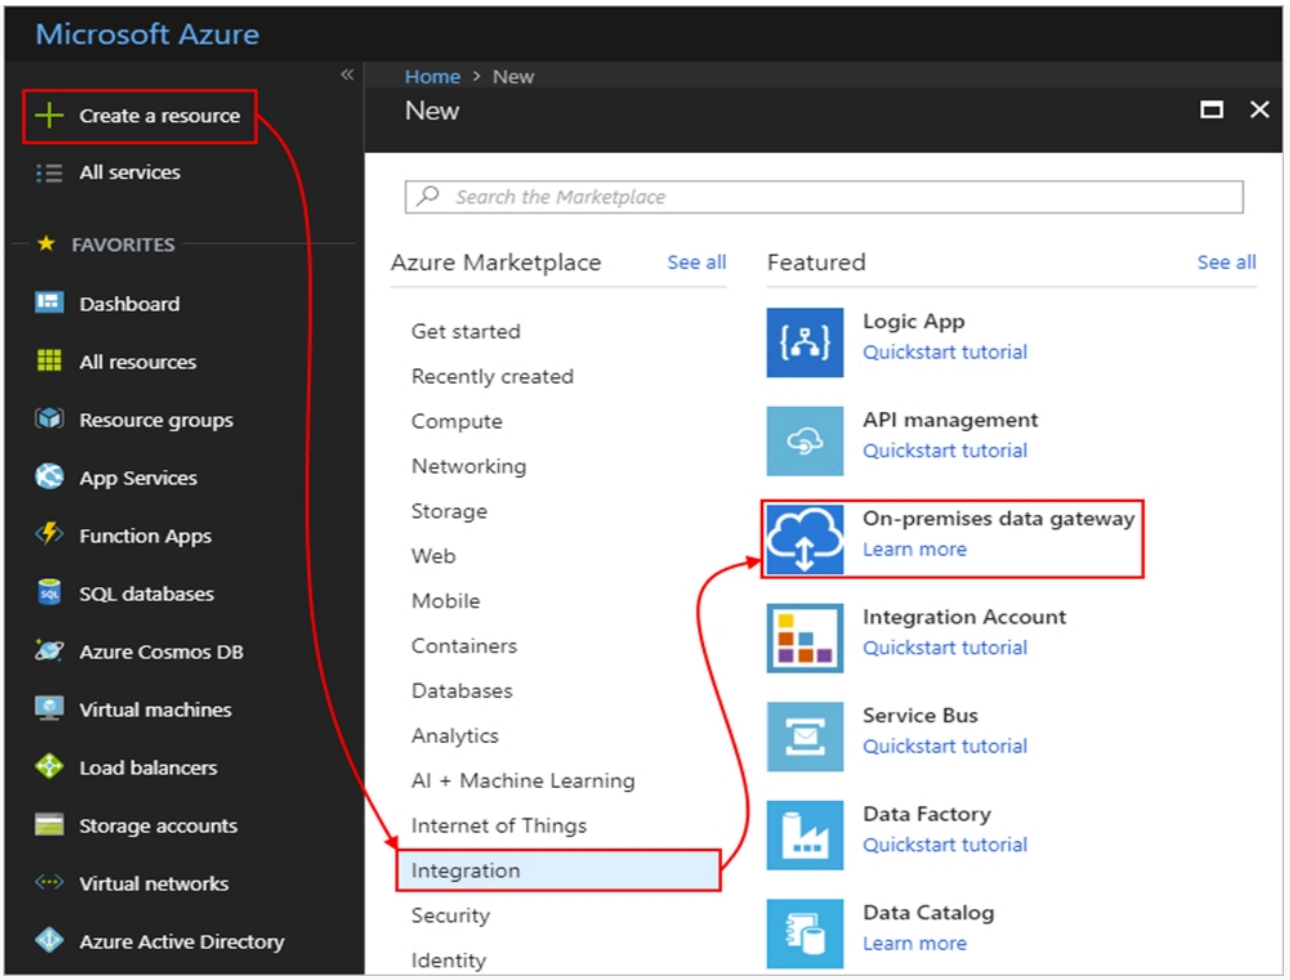

3. On the Create connection gateway page, provide this information for your gateway resource.

4. To add the gateway resource to your Azure dashboard, select Pin to dashboard. When you're done, choose Create.

Box 3: From the Logic Apps Designer in the Azure portal, add a connector

After you create your gateway resource and associate your Azure subscription with this resource, you can now create a connection between your logic app and your on-premises data source by using the gateway.

1. In the Azure portal, create or open your logic app in the Logic App Designer.

2. Add a connector that supports on-premises connections, for example, SQL Server.

3. Set up your connection.

References:

https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-gateway-connection

Currently there are no comments in this discussion, be the first to comment!

Question #10 *Topic 3*

You have a virtual network named VNet1 as shown in the exhibit.

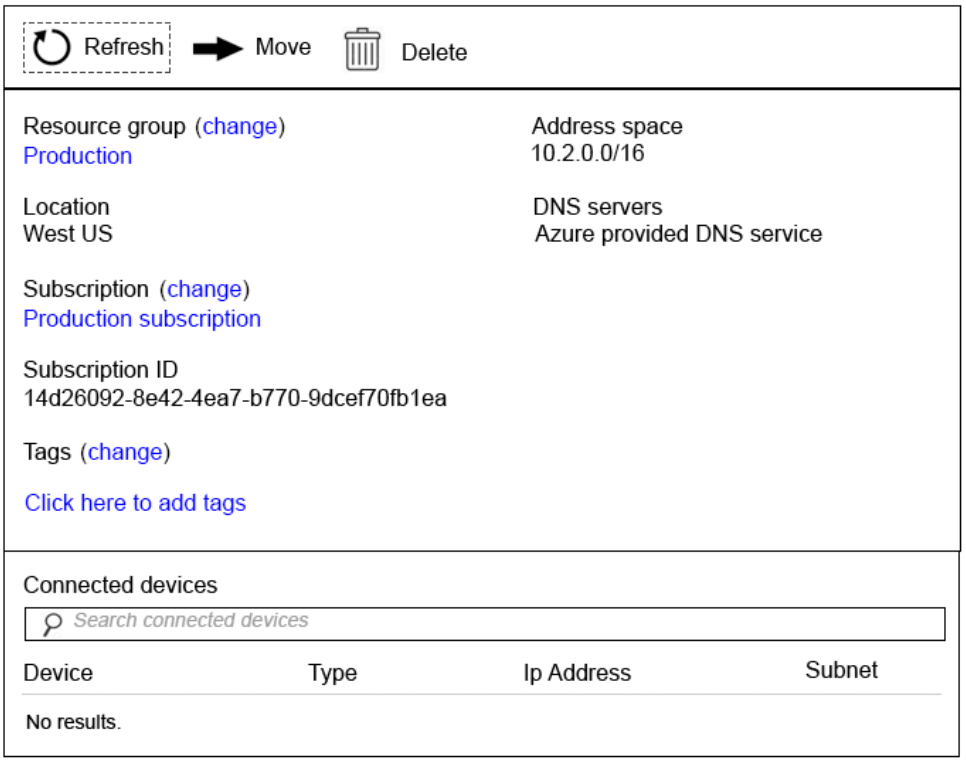

No devices are connected to VNet1.

You plan to peer VNet1 to another virtual network named Vnet2 in the same region. VNet2 has an address space of 10.2.0.0/16.

You need to create the peering.

What should you do first?

A. Modify the address space of VNet1.

B. Configure a service endpoint on VNet2

C. Add a gateway subnet to VNet1.

D. Create a subnet on VNet1 and VNet2.

Suggested Answer: *A*

The virtual networks you peer must have non-overlapping IP address spaces.

References:

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

Currently there are no comments in this discussion, be the first to comment!

Question #11 *Topic 3*

You have an Azure subscription that contains three virtual networks named VNet1, VNet2, 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. On the peering connections, allow gateway transit.

C. Create route tables and assign the table to subnets.

D. Create a route filter.

E. On the peering connections, use remote gateways.

Suggested Answer: *BE*
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.
References:
https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-manage-peering#requirements-and-constraints

Currently there are no comments in this discussion, be the first to comment!

HOTSPOT -

You have an Azure subscription named Subscription1. Subscription1 contains the virtual networks in the following table.

Name	Address space	Subnet name	Subnet address range
VNet1	10.1.0.0/16	Subnet1	10.1.1.0/24
VNet2	10.10.0.0/16	Subnet2	10.10.1.0/24
VNet3	172.16.0.0/16	Subnet3	172.16.1.0/24

Subscription1 contains the virtual machines in the following table.

Name	Network	Subnet	IP address
VM1	VNet1	Subnet1	10.1.1.4
VM2	VNet2	Subnet2	10.10.1.4
VM3	VNet3	Subnet3	172.16.1.4

The firewalls on all the virtual machines are configured to allow all ICMP traffic.

You add the peerings in the following table.

Virtual network	Peering network
VNet1	VNet3
VNet2	VNet3
VNet3	VNet1

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
VM1 can ping VM3.	○	○
VM2 can ping VM3.	○	○
VM2 can ping VM1.	○	○