

70-532.exam.167q

Number: 70-532
Passing Score: 800
Time Limit: 120 min



<https://www.gratisexam.com/>

70-532

Developing Microsoft Azure Solutions

<https://www.gratisexam.com/>

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

You are a developer for Fabrikam, a company that specializes in payment processing. Fabrikam is developing a solution to process payments for various events, such as music concerts. You develop an ASP.NET MVC website that is hosted in Azure to support an upcoming music concert. The music concert is expected to generate a large volume of ticket sales in a short amount of time.

The website uploads information to an Azure storage queue. A worker role in Azure retrieves information from the queue and generates the concert tickets in a PDF file format after the financial transaction is approved.

You observe a delay between the time the website adds a message to a queue and the time it becomes available to read from the queue. After examining the queue, you determine that no queue messages have a DequeueCount value greater than zero. The website does not throw any errors.

Business Requirements

Payments

The music concert website must be able to submit event payment information for processing. The website must remain responsive while submitting payment information. Customers must be able to add notes about their orders to a free-form control on the website. These notes must be submitted with the payment when the customer submits an order.

Customers often enter notes that exceed 7 KB in size.

Technical Requirements

Payment Submission and Processing

Event payment information must be sent from the website to a Windows Communication Foundation (WCF) service worker role. The worker role must submit the information to the payment processor in JSON format.

Payment Processing

You have the following payment processing requirements:

- If the number of messages in a queue goes above or below a specified threshold, worker role instances must be created or deleted as needed. This process must be completed by using the least amount of effort. It must be easy to reconfigure role instance thresholds.
- Payments must be retrieved from the queue in the maximum batch sizes that are allowed by the queue and pulled from the queue for 5 minutes.
- The payment queue must not be re-created when processing payments.
- During single Payment processing, the number of tickets available for an event must be updated. The update operation must be retried for 30 seconds or 5 retry attempts, whichever occurs first. Each retry should pause for at least two seconds and for one second longer than the previous attempt. If the update fails, the payment should be placed in the poison queue.

Storage

You have the following storage requirements:

- Payment information must be stored by using Azure Queue storage. Connection to the Azure storage account has been established in a configured setting named `StorageConnectionString`, which is configured for the web and worker roles.
- A payment processing queue and a poison payment queue must be used when processing payments.
- Azure Queue message content must be XML-safe and UTF-8 encoded.
- An Azure storage account must be established for diagnostic information in a configured setting named `DiagnosticsStorageConnectionString`, which is configured for both the web and worker roles.

Security and Monitoring

Security

The web role must be secured by using HTTPS.

Monitoring

You must collect diagnostic data for both the web and worker roles by using the Diagnostics module. Diagnostics configuration changes must not require the code of the roles to be rebuilt. The diagnostic data is used for debugging and troubleshooting, measuring performance, monitoring resource usage, traffic analysis and capacity planning, and auditing.

Performance testing must evaluate the roles under normal and stress conditions without incurring changes for running Azure. Memory allocation, function time, and multithreading concurrency issues must be evaluated.

Deployment

You purchase a custom domain name fabrikamfunding.com to host the website, web role, and worker roles. You must deploy an HTTPS certificate with the web role, and you must update associated configuration files accordingly.

Web role and worker role instance sizes must be specified as Medium. You must deploy one web role instance named FabrikamFundingPaymentGenerator, and worker role instances named FabrikamFundingPaymentProcessor.

Application Structure

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

CustomRetryPolicy.cs

```
CR01 public class CustomRetryPolicy : IRetryPolicy
CR02 {
CR03     int _retryCount = 0;
CR04     readonly TimeSpan _baseInterval= TimeSpan.FromSeconds(1);
CR05     readonly string _poisonPaymentQueueName;
CR06     private readonly CloudQueueClient _queueClient;
CR07     private readonly EventPayment _eventPayment;
CR08     public CustomRetryPolicy(string poisonPaymentQueueName, CloudQueueClient
queueClient, EventPayment eventPayment)
CR09     {
CR10         _poisonPaymentQueueName = poisonPaymentQueueName;
CR11         _queueClient = queueClient;
CR12         _eventPayment = eventPayment;
CR13     }
CR14     public IRetryPolicy CreateInstance()
CR15     {
CR16         return new CustomRetryPolicy(_poisonPaymentQueueName, _queueClient,
_eventPayment);
CR17     }
CR18 }
```

Event.cs

```
EV01 public class Event : TableEntity
EV02 {
EV03     public int AvailableTickets { get; set; }
EV04 }
```

EventPayment.cs

```
EP01 [DataContract]
EP02 public class EventPayment
EP03 {
EP04     [DataMember]
EP05     public int EventId { get; set; }
EP06     [DataMember]
EP07     public string Email { get; set; }
EP08     [DataMember]
EP09     public string Notes { get; set; }
EP10     [DataMember]
EP11     public int TicketCount { get; set; }
EP12     [DataMember]
EP13     public DateTime OrderDate { get; set; }
EP14     [DataMember]
EP15     public Guid EventPaymentId { get; set; }
EP16 }
```

QueueManager.cs

```
QM01 public class QueueManager
QM02 {
QM03     private readonly CloudQueueClient _queueClient;
QM04     private readonly CloudTableClient _tableClient;
QM05     private const string PaymentQueueName = "paymentqueue";
QM06     private const string PoisonPaymentQueueName = "poisonpaymentqueue";
QM07     public QueueManager()
QM08     {
QM09         var storageAccount = CloudStorageAccount.Parse(
QM10             CloudConfigurationManager.GetSetting("StorageConnectionString"));
QM11         _queueClient = storageAccount.CreateCloudQueueClient();
QM12         _tableClient = storageAccount.CreateCloudTableClient();
QM13     }
QM14     public async Task SendMessageAsync(EventPayment eventPayment)
QM15     {
QM16         ...
QM17     }
QM18     public async Task ProcessMessagesAsync()
QM19     {
QM20         ...
QM21     }
QM22     public async Task ProcessPayment(EventPayment eventPayment)
QM23     {
QM24         var events = _tableClient.GetTableReference("events");
QM25         var key = eventPayment.EventId.ToString();
QM26         var operation = await events.ExecuteAsync(TableOperation.Retrieve<Event>(key, key));
QM27         var @event = operation.Result as Event;
QM28         @event.AvailableTickets = @event.AvailableTickets - eventPayment.TicketCount;
QM29         var requestOptions = new TableRequestOptions
QM30         {
QM31             RetryPolicy = new CustomRetryPolicy(
QM32                 PoisonPaymentQueueName,
QM33                 _queueClient,
QM34                 eventPayment),
QM35         };
QM36         var context = new OperationContext
QM37         {
QM38             StartTime = DateTime.Now,
QM39         };
QM40         await events.ExecuteAsync(TableOperation.Replace(@event),
requestOptions, context;
QM41     }
QM42 }
```

QUESTION 1

The **SendMessageAsync** method of the **QueueManager** class occasionally throws errors.

You need to correct the errors.

What should you do?

- A. Remove all attributes from the **EventPayment** class.
- B. Encode the **notes** field content by using UTF-32 encoding.
- C. Update the **notes** field to a byte array. Binary encode and decode the **notes** content when sending or receiving an **EventPayment** class.
- D. Update the **SendMessageAsync** method of the **QueueManager** class to store the notes field in BLOB storage. Update the **EventPayment** class to store the BLOB uniform resource identifier (URI). Extract the **notes** BLOB information by using the BLOB URI in the **ProcessMessagesAsync** method of the **QueueManager** class.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 2

You need to diagnose the source of the performance issues when preparing concert tickets.

Which two actions should you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.



- A. Examine the Storage Logging logs for any queue operations that have higher than expected latency.
- B. Examine the Storage Client Library logs to determine whether there is a decrease in the total number of requests for storage operations.

- C. Examine the diagnostic message logs for the worker role to determine whether the worker role is failing to process messages.
- D. Examine the Storage Client Library logs to determine whether there are repeated retries for storage operations.

Correct Answer: AD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/storage/storage-monitoring-diagnosing-troubleshooting#you-are-experiencing-unexpected-delays-in-message-delivery>

QUESTION 3

The SendMessageAsync method of the QueueManager class occasionally throws errors.

You need to correct the errors.

What should you do?

- A. Update the QueueManager to use the Put Message operation of the Queue Service REST API. Use HTTP compression for all calls made to the REST API.
- B. Encode the notesfield content by using UTF-32 encoding.
- C. UpdateSendMessageAsyncmethod of the QueueManagerclass to store the notesfield in BLOB storage. Update the EventPaymentclass to store the BLOB uniform resource identifier (URI). Extract the notes BLOB information by using the BLOB URI in the ProcessMessagesAsyncmethod of the QueueManagerclass.
- D. Update the notesfield to a byte array. Binary encode and decode the notescontent when sending or receiving an EventPaymentclass.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

You are a developer for LitWare, Inc., a game development company. You are developing a backend service for an online social gaming platform named GamerData. The game is built around point generators, which are associated with physical landmarks. Players claim point generators which give them a set amount of points per day.

Business Requirements

Mobile App

The game itself runs on various mobile devices and is developed by TailSpin Toys, a company that specializes in mobile game development. The mobile app will periodically make calls to the GamerData service to find the five closest point generators that are located less than the specified distance from the player's current location. If no point generators are found, the search distance increases until one is found.

The mobile app shows all the point generators owned by each player. The mobile app allows for each player to search for claimed point generators by player name. This search does not require exact spelling of names. The details for each claimed generator is shown in the app.

When a player claims a point generator, they should receive an email notification. An Azure Function named EmailPlayer has been developed to email players with details about recently claimed point generators.

Sponsors

The platform allows business to sponsor point generators within a business location.

Reports

A report named Daily Sponsor Report must be generated each day at midnight. The report must contain a section for each sponsor. Each sponsor section must contain two subsections.

The first subsection of the report contains the names of the point generators for that sponsor, ordered by the last time the point generator was claimed. The second subsection contains the current owners for each of the point generators for the sponsor. Generation of reports must not impact the GamerData service.

Technical Requirements

GamerData Service

All data for the GamerData service is stored in an Azure DocumentDB instance named GamerData. Business and players interact with the service by using a REST API.

The REST API must:

- Produce valid Swagger API specifications for non-obsolete actions.
- Be optimized for loading specific point generators.
- Follow REST best practices.
- Include appropriate terms of service.

Costs for all Azure services must be minimized.

Build and Deployment

The GamerData service will be deployed to Azure in a private VNet.

Security

Sponsors have accounts in an Azure Active Directory (Azure AD) with business-to-consumer (B2C) enabled named litwaregamerdata.onmicrosoft.com managed by Litware, Inc. for both GamerData and LitWare, Inc. services.

Only Litware, Inc. developers and automated testing tools should be able to directly access the GamerData service. All other use of the service must be through Azure API Management. A description of the security practices used during development, available on Microsoft SharePoint, must be available to users of the API under the terms of service.

Reporting

Azure Search will be used as the source for running reports. The properties of indexes in Azure Search must match the names of the properties in DocumentDB.

Performance

The Azure DocumentDB must not be used for reporting purposes. All services must perform queries in the data store when possible.

Application Structure

Startup.cs

Relevant portions of the app files are shown below. (Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.)

```
SP01 public class Startup
SP02 {
SP03     public IConfigurationRoot Configuration { get; }
SP04     public Startup(IHostingEnvironment env)
SP05     {
SP06         var builder = new ConfigurationBuilder().SetBasePath(env.ContentRoot-
SP07         Path).AddJsonFile ("appsettings.json");
SP08         Configuration = builder.Build();
SP09     }
SP10     public void ConfigureServices(IServiceCollection services)
SP11     {
SP12         services.AddMvc();
SP13         Services.AddSwaggerGen();
SP14     }
SP15     public void Configure(IApplicationBuilder app, IHostingEnvironment env,
SP16     ILoggerFactory loggerFactory)
SP17     {
SP18         app.UseMvc();
SP19         app.UseSwagger();
SP20     }
```

PointController.cs

Relevant portions of the app files are shown below. (Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.)

```

PC01 [Route("api/pointgen")]
PC02     public class PointGeneratorController : Controller
PC03     {
PC04         private static readonly string DatabaseName = "GamerData";
PC05         private static readonly string CollectionName = "PointGenerators";
PC06         private static readonly string EndpointUrl = "...";
PC07         private static readonly string AuthorizationKey = "...";
PC08
PC09         [HttpGet("{name}")]
PC10         public async Task<PointGenerator> Get(string name)
PC11         {
PC12             using (var client = new DocumentClient(new Uri(EndpointUrl),
AuthorizationKey))
PC13             {
PC14                 var response = await client.ReadDocumentAsync(UriFactory.Create-
DocumentUri(DatabaseName, CollectionName, name));
PC15                 return (PointGenerator) (dynamic) response.Resource;
PC16             }
PC17         }
PC18
PC19         [Route("nearby")]
PC20         [HttpGet]
PC21         public IEnumerable<pointGenerator> Nearby(double longitude, double
latitude, long minDistance)
PC22         {
PC23             var location = new Point(longitude, latitude);
PC24             using (var client = new DocumentClient(new Uri(EndpointUrl),
AuthorizationKey))
PC25             {
PC26
PC27             }
PC28         }
PC29
PC30         public async Task<PointGenerator> Update[FromBody] PointGenerator pg)
PC31         {
PC32             using (var client = new DocumentClient(new Uri(EndpointUrl),
AuthorizationKey))
PC33             {
PC34                 var collection = await GetCollection();
PC35                 await client.UpsertDocumentAsync(collection.SelfLink, pg);
PC36                 return pg;
PC37             }
PC38         }
PC39         private static async Task<DocumentCollection> GetCollection()
PC40         {
PC41             ...
PC42         }
PC43     }

```

PointGenerator.cs

Relevant portions of the app files are shown below. (Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.)

```
PG01 public class PointGenerator
PG02     {
PG03
PG04     public string Name { get; set; }
PG05     [JsonProperty("currentOwner")]
PG06     public string CurrentOwner { get; set; }
PG07     [JsonProperty("sponsor")]
PG08     public string Sponsor { get; set; }
PG09     [JsonProperty("dateLastClaimed")]
PG10     public DateTimeOffset DateLastClaimed { get; set; }
PG11     [JsonProperty("location")]
PG12     public Point Location { get; set; }
PG13 }
```

QUESTION 1

HOTSPOT

You need to create the Azure Search index.

How should you configure the Azure Search index? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

FIELD NAME	TYPE	RETRIEVABLE	FILTERABLE	SORTABLE
currentOwner	Edm.String	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
sponsor	Edm.String	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
dateLastClaimed	Edm.DateTimeOffset	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Correct Answer:

FIELD NAME	TYPE	RETRIEVABLE	FILTERABLE	SORTABLE
currentOwner	Edm.String	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
sponsor	Edm.String	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
dateLastClaimed	Edm.DateTimeOffset	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Section: [none]

Explanation

Explanation/Reference:

QUESTION 2

DRAG DROP

You need to write a method to return the email address for a given sponsor.

What should you do? To answer, drag the appropriate code segment to the correct location. Each code segment 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:

Code segments

```
https://graph.windows.net/litwaregamerdata
```

```
https://litwaregamerdata.onmicrosoft.com/graph
```

```
https://graph.windows.net/litwareinc
```

```
https://litwareinc.onmicrosoft.com/graph
```

```
/users?api-version=1.6&$filter=displayName eq \' {pg.Sponsor} \'
```

```
/users?api-version=1.6&$filter=displayName eq \' {pg.CurrentOwner} \'
```

```
result.value[0].mail
```

```
result.value[0].sipProxyAddress
```



Answer Area

```
private async Task<string> getEmail(PointGenerator pg, string accessToken)
```

```
{    var url = " " ;
```

```
    url += $" " ;
```


Correct Answer:

Code segments

```
https://litwaregamerdata.onmicrosoft.com/graph
```

```
https://graph.windows.net/litwareinc
```

```
https://litwareinc.onmicrosoft.com/graph
```

```
/users?api-version=1.6&$filter=displayName eq \' {pg.CurrentOwner} \'
```

```
result.value[0].sipProxyAddress
```



Answer Area

```
private async Task<string> getEmail(PointGenerator pg, string accessToken)
{
    var url = "https://graph.windows.net/litwaregamerdata";
    url += $" /users?api-version=1.6&$filter=displayName eq \' {pg.Sponsor} \'
```

Section: [none]

Explanation

Explanation/Reference:



<https://www.gratisexam.com/>

QUESTION 3

You need to add a routing constraint.

Which code segment should you add at line PC29?

- A. [HttpDelete]
- B. [HttpPost]
- C. [HttpOptions]
- D. [HttpsHead]

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 4

DRAG DROP

You need to ensure that sponsors can interact with the GamerData service by using the same credentials as they use for other LitWare, Inc. services.

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

Create a new Azure AD named litware.onmicrosoft.com.

Enable OAuth 2.0 user authorization in Azure AD.

Configure an API Management OAuth 2.0 authorization server.

Configure an Azure AD OAuth 2.0 authorization server.

Enable OAuth 2.0 user authorization in API Management.

Register the API Management developer portal as an Azure AD application in litware.onmicrosoft.com.

Register the API Management developer portal as an Azure AD application in litware2.onmicrosoft.com.

Register the API Management developer portal as an Azure AD application in litwaregamedata.onmicrosoft.com.

Answer Area



Correct Answer:

Actions

Create a new Azure AD named litware.onmicrosoft.com.

Configure an API Management OAuth 2.0 authorization server.

Enable OAuth 2.0 user authorization in API Management.

Register the API Management developer portal as an Azure AD application in litware.onmicrosoft.com.

Register the API Management developer portal as an Azure AD application in litware2.onmicrosoft.com.

Answer Area

Enable OAuth 2.0 user authorization in Azure AD.

Register the API Management developer portal as an Azure AD application in litwaregamedata.onmicrosoft.com.

Configure an Azure AD OAuth 2.0 authorization server.



Section: [none]
Explanation

Explanation/Reference:

QUESTION 5

You need to write an Azure Search Query to return data for the first subsection of the Daily Sponsor Report.

Which query string should you use?

- A. facets=currentOwner&sort=dateLastClaimed
- B. \$filter=sponsor&sort=dateLastClaimed
- C. search=currentOwner&sort=dateLastClaimed
- D. group=sponsor&sort=dateLastClaimed
- E. facets=sponsor&sort=dateLastClaimed

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 6

You need to decrease the amount of time it takes to query point generators by configuring API management caching.

In the Azure portal, which value should you use for the Vary by Query string parameters setting?

- A. name
- B. longitude;latitude;minDistance
- C. longitude;latitude;dateLastClaimed
- D. Id

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Testlet 1

Background:

You are a developer for ProseWare Inc., a software-as-a-service (SaaS) company that provides a comment system that websites use to allow for end users to post comments associated with a webpage or topic on a customer's website.

Business requirements

Moderation:

The moderation of comments is a feature of the software, and usually involves the editing of a comment.

Only users who have accounts in a group in Azure Active Directory (Azure AD) have the ability moderate. External users can also become moderators, but only by explicit invitation.

Any moderation action must include the name of the moderator.

Comment navigation:

Each comment is identified by a unique string consisting of a random string of characters.

Within the body of a comment, internal links to other comment threads can be specified using the link format: “/<parent comment id> / <child comment id>”

Comment search:

Comments can be searched using Azure Search. Searches must do the following:

- Searching for email addresses must match email addresses in comments.
- Searching must work for the client's language.
- Internal links to other comments using the link format should be searched.

Content screening:

Comment content is screened for inappropriate language, length, and topic using content analysis. Content must be screened, but can appear prior to be screened.

Mobile App:

The moderation functionality can be accessed using a Universal Windows Platform (UWP) app named ProsewareApp. The app includes functionality that notifies moderators when changes are made to a comment they modified.

Export:

Customers can perform an export of all comments to a customer supplied Microsoft OneDrive folder on demand. The export functionality is implemented as an Azure Logic App, and it must be able to be triggered by the customer from their local network.

Interaction agents:

Interaction agents are parts of the system that interact with comment threads. The main purpose is to modify a comment's body based on the contents of the comment. For example, one of the agents is WikiAgent, which adds links to Wikipedia articles when it sees text in the comment body that exactly matches a Wikipedia article title. Interaction Agents are implemented in Service Fabric.

Interaction agents must meet the following requirements:

- Only successfully process each comment once
- Any errors encountered during the processing of a comment should be retried
- Must run on systems that allow for custom applications to be installed
- Must run in a VNet or private network space
- Must be run on a system that can scale up and down based on demand
- A single user's usage of Interaction Agents must not impact other users' usage of Interaction Agents

Technical requirements

Authentication:

ProseWare Inc. allows for user authentication through Azure AD and Twitter.

Storage:

The application runs as a Web App on Azure. Comments are stored in an Azure DocumentDB database named "Proseware".

Performance:

The product includes a service level agreement (SLA) for individual method performance. All data retrieval methods must return within 100ms 99% of the time.

API:

The ProseWare Inc. API is made available to public callers using an Azure API App. Azure AD and Twitter are the Authentication Providers.

Application structure

CommentController.cs:


```

CC01 [Route("api/[controller]")]
CC02 public class CommentController : Controller
CC03 {
CC04     private IDatabase _redis;
CC05     private DataStore _dataStore;
CC06     private CloudQueue _queue
CC07
CC08     public CommentController ()
CC09     {
CC10         _queue = CloudStorageAccount.Parse(" ").Create-
CloudQueueClient().GetQueueReference ("commentQueue");
CC11         _redis = ConnectionMultiplexer.Connect("...").GetData-
base();
CC12         _dataStore = new DataStore();
CC13     }
CC14
CC15     [HttpGet("{commentId}")]
CC16     public async Task<Comment> Get(string commentId)
CC17     {
CC18         var cached = await _redis.StringGetAsync(commentId);
CC19         if (cached.HasValue)
CC20         {
CC21             return JsonConvert.DeserializeObject<Com-
ment>(cached.ToString());
CC22         }
CC23         return await _dataStore.LoadAsync(commentId);
CC24     }
CC25
CC26     [HttpGet]
CC27     public IEnumerable<Comment> GetChildComments(string com-
mentId)

```

```

CC28 {
CC29     IEnumerable<Comment> result = null;
CC30
CC31     if (result == null)
CC32     {
CC33         result = _dataStore.LoadThread(commentId);
CC34     }
CC35     return results;
CC36 }
CC37
CC38 [HttpPost]
CC39 public async Task<IActionResult> New([FromBody]Comment
comment)
CC40 {
CC41     await Save(comment);
CC42     return Ok();
CC43 }
CC44
CC45 [HttpPost]
CC46 public async Task<IActionResult> Reply(string inRe-
sponseTo, [FromBody]Comment comment)
CC47 {
CC48     comment.InResponseTo = inResponseTo;
CC49     await Save(comment);
CC50     return View();
CC51 }
CC52
CC53 private static Comment Convert(string json)
CC54 {
CC55     return JsonConvert.DeserializeObject<Comment>(json);
CC56 }
CC57
CC58 private async Task Save(Comment comment, string moderator-
Name = null)
CC59 {
CC60     comment.Moderator = moderatorName;
CC61     var json = JsonConvert.SerializeObject(comment);
CC62     _redis.StringSet(comment.Id, json);
CC62
CC63
CC64     await _queue.AddMessageAsync(new CloudQueueMessage(com-
ment.Id));
CC65     _dataStore.Save(comment);
CC66 }

```

cleaner.csx:

```
CL01 #r "Newtonsoft.Json"
CL02
CL03 using System;
CL04 using Newtonsoft.Json;
CL05 using Newtonsoft.Json.Linq;
CL06 public static void Run(string commentId, object result,
TraceWriter log)
CL07 {
CL08     dynamic comment = JObject.Parse(item);
CL09     ...
CL10     result = comment;
CL11 }
```

ICommentAgent.cs:

```
CA01 public interface ICommentAgent: IActor
CA02     {
CA03     Task<string> ModifyCommentText(string id, string body,
string title);
CA04     }
```

WikiAgent.cs:

```

WA01 [StatePersistence(StatePersistence.Persisted)]
WA02 internal class WikiAgent : Agent, ICommentAgent
WA03 {
WA04 public WikiAgent(ActorService, ActorId id) : base(service,
id) {}
WA05 public async Task<string> ModifyCommentText(string id,
string body, string title)
WA06 {
WA07     try
WA08     {
WA09         var newBody = scanForLinks(body);
WA10
WA11         return newBody;
WA12     }
WA13     catch
WA14     {
WA15     }
WA16
WA17     throw
WA18 }
WA19 }
WA20 }

```

Comment.cs:

```

CO01 public class Comment
CO02 {
CO03     public string Id {get; set;}
CO04     public string UserId {get; set;}
CO05     public string InResponseTo {get; set;}
CO06     public string Title {get; set;}
CO07     public DateTimeOffset Date {get; set;}
CO08     public string Body {get; set;}
CO09     public string Moderator {get; internal; set;}
CO10 }

```

DataStore.cs:

```

DS01 public class DataStore
DS02 {
DS03     private const string EndpointUrl = "https:
//proseware.documents.azure.com:443/";
DS04     private const string PrimaryKey = "";
DS05     private const string db = "Proseware";
DS06     private const string col = "Comments";
DS07     private DocumentClient client;
DS08
DS09     public DataStore()
DS10     {
DS11         client = new DocumentClient(new Uri(EndpointUrl), Pri-
maryKey);
DS12     }
DS13
DS14     public async Task<Comment> LoadAsync(string commentId)
DS15     {
DS16         var uri = UriFactory.CreateDocumentCollectionUri(db,
col);
DS17         return await client.ReadDocumentAsync<Comment>(UriFac-
tory.CreateDocumentUri(db, col, commentId));
DS18     }
DS19
DS20     public async void Save(Comment comment)
DS21     {
DS22         var uri = UriFactory.CreateDocumentCollectionUri(db,
col, comment.Id);
DS23         await client.UpsertDocumentAsync(uri, comment);
DS24     }
DS25     public IEnumerable<Comment> LoadThread(string commentId)
DS26     {
DS27         var uri = UriFactory.CreateDocumentCollectionUri(db,
col);
DS28         return client.CreateDocumentQuery<Comment>(uri).Where(f
=> f.Id == commentId);
DS29     }
}

```

MainPage.xaml.cs:

```
MP01 public sealed partial class MainPage : Page
MP02 {
MP03     public MainPage()
MP04     {
MP05         InitializeComponent();
MP06     }
MP07
MP08     private async void StartNotify()
MP09     {
MP10     }
MP11
MP12     private void UpdateUI()
MP13     {
MP14     }
MP15 }
```

QUESTION 1

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 need to implement the infrastructure for the Interaction Agents.

Solution: Create an Azure Container Service cluster and create a container for running Service Fabric.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 2

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 need to implement the infrastructure for the Interaction Agents.

Solution: Create a set of Azure virtual machines (VMs) using Azure Resource Manager (ARM) templates, and use Chef to install the Service Fabric runtime.

Does the solution meet the goal?



A. Yes

B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 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 need to implement the infrastructure for the Interaction Agents.

Solution: Create an Azure virtual machine (VM) scale set and use Azure Desired State Configuration (DSC) extension handler to install Service Fabric runtime.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 4

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 need to implement the infrastructure for the Interaction Agents.

Solution: Create a Service Fabric cluster with Bronze durability and reliability tiers.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 5

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 need to ensure that moderators can be added to the system.

Which authentication approach should you use?

- A. Microsoft Office 365 directory
- B. Azure AD self-service signup
- C. Azure AD Organizational Units (OU)
- D. Active Directory Federation

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Testlet 1

Background

Overview

Trey Research Inc. is a software as a service (SaaS) company that provides solutions to the legal industry including a mobile app named Finder. The company supports users globally.

Case files

Case files are documents about legal proceedings, such as trial transcripts, legal research, and summaries of trial results. Case files are used to capture trial information for future use.

Before starting a trial, a new case file is created. Researchers link other potentially relevant case files to the new case file.

During a trial, users create transcripts and notes related to the trial. Users may create documents on a variety of platforms including Linux, Mac, OS, and Windows. The files may be created in a variety of formulas including Microsoft Word documents, images, and text files.

All documents are uploaded to a file share. Once documents are uploaded, the system imports the documents as case files.

When a trial is concluded, the case file for that trial is marked as finished and made available as research for future trials. During this process, all licensed case files are removed from case file association.

Finder

Lawyers use Finder to search for and view information and documents that relate to a case. The interface includes an auto-complete text box where users can enter search terms. Searches must be performed using combination of characters in the search input.

IntelligentAssist

You plan to build an application named IntelligentAssist that lawyers will use during trials. IntelligentAssist uses several third-party servers to automatically generate a set of case files using licensed data. These third-party case files have an expiration data, after which they must not be accessed.

Requirements

Transcripts and notes

The transcripts and notes created during a trial are uploaded from machines that do not have consistent Internet connections. Documents may be uploaded multiple times with different file names. This action must not result in multiple case files being generated. The processing of uploaded documents must be started by an Azure function.

Case Files

All operations on case files must be billed when an operation is performed. Processing of an uploaded file is considered a single operation.

Security

When a Service Fabric Actor is activated, the information about the activation must be stored in Cosmos DB.

Users of the system will authenticate using their organization's Active Directory. When a user is removed from an organization's Active Directory, their access to the IntelligentAssist product must be revoked as soon as possible.

Application Structure

CaseFile.cs

Relevant portions of the app files are shown below. Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

CaseFile.cs

```
CF01 public class CaseFile
CF02 {
CF03     public string Id { get ; set; }
CF04     public List <string> RelatedCaseFiles { get; set ; }
CF05     public DateTimeOffset ? Expires { get ; set ; }
CF06     public bool InProgress { get ; set; }
CF07     public string Body { get ; set; }
CF08 }
```

MainPage.cs

Relevant portions of the app files are shown below. Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

MainPage.cs

```
MP01 sealed partial class MainPage : Page
MP02 {
MP03     private IMobileServiceSyncTable <CaseFile> caseFileTable = App.MobileService.GetSyncTable<CaseFile>();
MP04     public MainPage()
MP05     {
MP06         this.InitializeComponent();
MP07     }
MP08     private async void ButtonRefresh_Click( object sender, RoutedEventArgs e)
MP09     {
MP10
MP11     }
MP12 }
```

CaseFileController.cs

Relevant portions of the app files are shown below. Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

CaseFileController.cs

```
CC01 public class CaseFileController : TableController <CaseFileData >
CC02 {
CC03     private ISubscriber _subscriber;
CC04     private Dictionary <string, CaseFile> Cache = new Dictionary<string , CaseFile>();
CC05     public CaseFileController()
CC06     {
CC07
CC08     }
CC09 }
```

ActorProgram.cs

Relevant portions of the app files are shown below. Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

ActorProgram.cs

```
AP01 public static class Program
AP02 {
AP03     private static void Main()
AP04     {
AP05         ActorRuntime .RegisterActorAsync< CaseFileActor>((context, actorType) =>
AP06             new ActorService (context, actorType, settings: new ActorServiceSettings
AP07             {
AP08
AP09             }));
AP10         Thread.Sleep( Timeout.Infinite);
AP11     }
AP12 public async Task SetupSearch()
AP13 {
AP14     var client = new SearchServiceClient("..." , new SearchCredentials ("..."));
AP15     client.Indexes.Create(new Index
AP16     {
AP17         Tokenizers = new [] {
AP18
AP19         }
AP20     });
AP21 }
AP22 }
```

CaseFileActor.cs

Relevant portions of the app files are shown below. Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

CaseFileActor.cs

CA01

CA02 internal class CaseFileActor : Actor , ICaseFileActor

CA03 {

CA04 private ISubscriber _subscriber;

CA05

CA06 public CaseFileActor(ActorService actorService, ActorId actorId)

CA07 : base (actorService, actorId)

CA08 {

CA09 _subscriber = ConnectionMultiplexer.Connect("_").GetSubscriber();

CA10 }

CA11 public async Task MarkAsCompleted(string caseFileId)

CA12 {

CA13 var client = new DocumentClient(ConnInfo.Uri, ConnInfo.Key);

CA14 var caseFileUri = UriFactory.CreateDocumentUri(ConnInfo.DatabaseId, ConnInfo.CollectionId, caseFileId);

CA15 var caseFile = await client.ReadDocumentAsync<CaseFile>(caseFileUri);

CA16 caseFile.Document.InProgress = true;

CA17

CA18 var collectionUri = UriFactory.CreateDocumentCollectionUri(ConnInfo.DatabaseId, ConnInfo.CollectionId);

CA19 foreach(var result in client.CreateDocumentQuery<CaseFile>(<collectionUri, query>.Where(r => r.Expires != null

CA20 {

CA21 caseFile.Document.RelatedCaseFiles.Remove(result.Id);

CA22 }

CA23 await client.UpsertDocumentAsync(caseFileUri, caseFile);

CA24 await _subscriber.PublishAsync("CaseFileUpdate", caseFileId);

CA25 }

CA26 private List <string > AlreadyProcessed;

CA27 public async Task UploadCaseFile(string location)

CA28 {

CA29 var contents = await ReadFile(location);

CA30

CA31 if (AlreadyProcessed.Contains(check))

CaseFileSwagger.json

Relevant portions of the app files are shown below. Line numbers in the code segments are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

CaseFileSwagger.json

```
CS01 {
CS02   "swagger" : "2.0",
CS03   "info" : {
CS04     "version" : "v1",
CS05     "title" : "IntelligentAssist"
CS06   },
CS07   "basePath" : "/",
CS08   "paths" : {
CS09     "/api/casefiles" : {
CS10       "post" : {
CS11         "tags" : [ "Values" ],
CS12         "operationId" : "ApiCasefilesPost",
CS13         "consumes" : [ "application/json", "text/json", "application/json-patch+json" ],
CS14         "produces" : [],
CS15         "parameters" : [
CS16           {
CS17             "name" : "value",
CS18             "in" : "body",
CS19             "required" : false,
CS20             "schema" : {
CS21               "$ref" : "#/definitions/CaseFile"
CS22             }
CS23           }
CS24         ],
CS25         "responses" : { "200": { "description" : "Success"}}
CS26       }
CS27     }
CS28   },
CS29   "definitions" : {
CS30
CS31   },
CS32   "securityDefinitions" : {}
CS33 }
```

QUESTION 1

Note: This 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 need to ensure that users can upload transcripts and notes.

Solution: You use Azure Resource Manager (ARM) to provision a storage account, and then use ARM to provision an Azure Cosmos DB database.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Resource Manager should be used.

The infrastructure for your application is typically made up of many components – maybe a virtual machine, storage account, and virtual network, or a web app, database, database server, and 3rd party services. You do not see these components as separate entities, instead you see them as related and interdependent parts of a single entity. You want to deploy, manage, and monitor them as a group. Azure Resource Manager enables you to work with the resources in your solution as a group. You can deploy, update, or delete all the resources for your solution in a single, coordinated operation. You use a template for deployment and that template can work for different environments such as testing, staging, and production. Resource Manager provides security, auditing, and tagging features to help you manage your resources after deployment.

Scenario: The transcripts and notes created during a trial are uploaded from machines that do not have consistent Internet connections. Documents may be uploaded multiple times with different file names. This action must not result in multiple case files being generated. The processing of uploaded documents must be started by an Azure function.

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-overview>

QUESTION 2

Note: This 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 need to ensure that users can upload transcripts and notes.

Solution: You use Azure Resource Manager (ARM) to provision a storage account, and then use the Azure portal to create an Azure File share.

Does the solution meet the goal?



<https://www.gratisexam.com/>

A. Yes

B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Resource Manager should be used.

The infrastructure for your application is typically made up of many components – maybe a virtual machine, storage account, and virtual network, or a web app, database, database server, and 3rd party services. You do not see these components as separate entities, instead you see them as related and interdependent parts of a single entity. You want to deploy, manage, and monitor them as a group. Azure Resource Manager enables you to work with the resources in your solution as a group. You can deploy, update, or delete all the resources for your solution in a single, coordinated operation. You use a template for deployment and that template can work for different environments such as testing, staging, and production. Resource Manager provides security, auditing, and tagging features to help you manage your resources after deployment.

Scenario: The transcripts and notes created during a trial are uploaded from machines that do not have consistent Internet connections. Documents may be uploaded multiple times with different file names. This action must not result in multiple case files being generated. The processing of uploaded documents must be started by an Azure function.

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-overview>

QUESTION 3

<https://www.gratisexam.com/>

Note: This 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 need to ensure that users can upload transcripts and notes.

Solution: You use Azure Resource Manager (ARM) to provision a storage account, and then use Azure PowerShell to provision an Azure File share.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The infrastructure for your application is typically made up of many components – maybe a virtual machine, storage account, and virtual network, or a web app, database, database server, and 3rd party services. You do not see these components as separate entities, instead you see them as related and interdependent parts of a single entity. You want to deploy, manage, and monitor them as a group. Azure Resource Manager enables you to work with the resources in your solution as a group. You can deploy, update, or delete all the resources for your solution in a single, coordinated operation. You use a template for deployment and that template can work for different environments such as testing, staging, and production. Resource Manager provides security, auditing, and tagging features to help you manage your resources after deployment.

Scenario: The transcripts and notes created during a trial are uploaded from machines that do not have consistent Internet connections. Documents may be uploaded multiple times with different file names. This action must not result in multiple case files being generated. The processing of uploaded documents must be started by an Azure function.

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-overview>

QUESTION 4

DRAG DROP

You need to ensure that third-party data providers can upload case files using the API.

What should you insert at line CS30? To answer, drag the appropriate JSON fragment to the correct location. Each JSON fragment 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:

JSON segments

CaseFile

/api/casefile

array

list

date-time

epoch-time

boolean

int

Answer area

```
{  
  "RelatedCaseFiles": {  
    "type": "[input]",  
    "items": {  
      "type": "string"  
    }  
  },  
  "expires": {  
    "format": "[input]"  
    "type": "string"  
  },  
  "InProgress": {  
    "type": "[input]"  
  }  
}
```

Correct Answer:

JSON segments

/api/casefile

list

date-time

boolean

int

Answer area

```
{
  "RelatedCaseFiles": {
    "type": "epoch-time",
    "items": {
      "type": "string"
    }
  },
  "expires": {
    "format": "array",
    "type": "string"
  },
  "InProgress": {
    "type": "CaseFile"
  }
}
```

Section: [none]

Explanation

Explanation/Reference:

QUESTION 5

Note: This 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 need to ensure that users can upload transcripts and notes.

Solution: You create a virtual machine (VM) image that contains a Server Message Block (SMB) share, and then use Azure Resource Manager (ARM) to provision a VM based on the image.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Resource Manager should be used.

The infrastructure for your application is typically made up of many components – maybe a virtual machine, storage account, and virtual network, or a web app, database, database server, and 3rd party services. You do not see these components as separate entities, instead you see them as related and interdependent parts of a single entity. You want to deploy, manage, and monitor them as a group. Azure Resource Manager enables you to work with the resources in your solution as a group. You can deploy, update, or delete all the resources for your solution in a single, coordinated operation. You use a template for deployment and that template can work for different environments such as testing, staging, and production. Resource Manager provides security, auditing, and tagging features to help you manage your resources after deployment.

Scenario: The transcripts and notes created during a trial are uploaded from machines that do not have consistent Internet connections. Documents may be uploaded multiple times with different file names. This action must not result in multiple case files being generated. The processing of uploaded documents must be started by an Azure function.

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-overview>

QUESTION 6

DRAG DROP

You need to ensure that operations are billed correctly.

Which code should you insert at line AP09? To answer, drag the appropriate code fragment to the correct location. Each code fragment 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:

Code fragments

ActorServiceSettings

ActorGarbageCollectionSettings

ReentrancyMode

ActorMethodContext

ActorReentrancyMode.Disallowed

ActorReentrancyMode.LogicalCallContext

ActorCallType.ActorInterfaceMethod

ActorCallType.ReminderMethod

Answer Area

Code fragment

= new

Co

{

Code fragment

=

Code

}

Correct Answer:

Code fragments

ActorServiceSettings

ActorGarbageCollectionSettings

ReentrancyMode

ActorMethodContext

ActorReentrancyMode.Disallowed

ActorReentrancyMode.LogicalCallContext

ActorCallType.ActorInterfaceMethod

ActorCallType.ReminderMethod

Answer Area

ActorServiceSettings

= new

ActorServiceSettings

{

ReentrancyMode

=

ActorReentrancyMode

}

Section: [none]
Explanation

Explanation/Reference:

Explanation:

Scenario: All operations on case files must be billed when an operation is performed. Processing of an uploaded file is considered a single operation.

References: <https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicefabric.actors.runtime.actorservicesettings?view=azure-dotnet>

QUESTION 7

DRAG DROP

You need to add code at line CA17 to ensure that third-party licenses are handled correctly when a case is closed.

How should you complete the code segment? To answer, drag the appropriate code segment to the correct location. Each code segment 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:

Code segments

CaseFiles

Trey

RelatedCaseFiles

RelatedCaseFiles _rid

COUNT

CONTAINS

ARRAY_CONTAINS

IS_DEFINED

Answer Area

```
var query = $ "SELECT * FROM Code segment WHERE '  
Code segment ( Code segment ,  
Code segment , ' {caseFileId} ' ) " ;
```

Correct Answer:

Code segments

Trey

COUNT

CONTAINS

IS_DEFINED

Answer Area

```
var query = $ "SELECT * FROM CaseFiles WHERE '  
ARRAY_CONTAINS ( RelatedCaseFiles  
RelatedCaseFiles _rid , ' {caseFileId} ' ) " ;
```

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Scenario: You plan to build an application named IntelligentAssist that lawyers will use during trials. IntelligentAssist uses several third-party servers to automatically generate a set of case files using licensed data. These third-party case files have an expiration data, after which they must not be accessed.

Testlet 1

Background

Contoso, Ltd. designs and manufactures medical vaccines. The company has sales and manufacturing facilities in locations all over the world. The company must permanently retain all data from clinical testing of vaccines. Contoso is subject to penalties and lawsuits if any data is lost

You maintain an application that allows medical professionals to manage data about patients, vaccine results, and possible side effects of the vaccines. You host the application in Microsoft Azure and manage the application by using the Azure portal.

Current environment

Applications

You have two on-premises applications that connect to Azure resources. The first application is named Contoso1. Contoso1 uses Windows identity Foundation (WIF) and includes a secured area of the application that is used for single-sign on with credentials from Azure Active Directory (AD).

The second application is an auditing application, named ClinicalSearch. It allows you to search for clinical test results. The application resides on an Azure Virtual Machine (VM) and must meet Azure Service Level Agreements (SLAs) for availability. Auditors will be given access to the application in 24 hour intervals.

You have a data Collection and distribution service named DataCollect. The service runs Node.js in an Azure cloud service. The service must use a minimum of 2048-bit encryption.

Infrastructure

You provision the following objects:

- Premium storage
- Two SQL Azure databases named CONTOSO-SQL1 and CONTOSO-SQL2. CONTOSO-SQL1 will store current product data and will be scaled up to optimize performance during peak hours, CONTOSO-SQL2 will store archived product data,
- Table storage

Deployment

You deploy resources as a resource group by using Azure Resource Manager (ARM). You use templates to deploy resources to different environments. The cloud service is deployed using standard methods. You provide medical professionals and other clients access to the storage resources.

Problem statements

Users are reporting that during peak hours, it takes a long time to save documents to the environment. Researchers report that the DataCollect service takes a long time to respond when data is being sent to the service.

Monitoring

You enable minimal monitoring metrics for the application. You enable boot diagnostics for all VMs.

Requirements

Security

You identify the following security-related requirements:

- You must not make changes to any application itself.
- You must not save any data on the local drive of any device that uses any application.
- You must secure access to data and network connections between any application and Azure resources.
- You must delete all resources that exist in the resource group not specified in the templates.
- Any new code introduced to the environment needs to be submitted for architectural review with senior IT architects.

Monitoring

You must monitor the application at all times. You have two monitoring events, Monitoring event 1 and Monitoring event 2.

- Monitoring event 1 monitors entity group transactions.
- Monitoring event 2 monitors entity versioning.

System administrators must include graphical views of monitoring information in a daily report. The graphs need to be obtained using minimal effort.

QUESTION 1

You need to save the diagnostics information.

What should you use?



<https://www.gratisexam.com/>

- A. the Check Health feature
- B. a Standard storage account for the diagnostics
- C. a custom script to store the diagnostics in the storage account
- D. a Remote Desktop connection to access the VM and save the log files

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Scenario: You enable minimal monitoring metrics for the application. You enable boot diagnostics for all VMs.

Enable diagnostics on a new virtual machine

1. When creating a new virtual machine from the Azure Portal, select the Azure Resource Manager from the deployment model dropdown:
2. In Settings, enable the Boot diagnostics, and then select a storage account that you would like to place these diagnostic files.
3. If you are deploying from an Azure Resource Manager template, navigate to your virtual machine resource and append the diagnostics profile section.
4. If you are deploying from an Azure Resource Manager template, navigate to your virtual machine resource and append the diagnostics profile section.

References: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/boot-diagnostics>

QUESTION 2

HOTSPOT

You need to configure monitoring.

For each event, which property should you use? To answer, select the appropriate property from each list in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Event	Property								
Monitoring event 1	<table><tr><td></td><td>▼</td></tr><tr><td colspan="2">PartitionKey</td></tr><tr><td colspan="2">RowKey</td></tr><tr><td colspan="2">Timestamp</td></tr></table>		▼	PartitionKey		RowKey		Timestamp	
	▼								
PartitionKey									
RowKey									
Timestamp									
Monitoring event 2	<table><tr><td></td><td>▼</td></tr><tr><td colspan="2">PartitionKey</td></tr><tr><td colspan="2">RowKey</td></tr><tr><td colspan="2">Timestamp</td></tr></table>		▼	PartitionKey		RowKey		Timestamp	
	▼								
PartitionKey									
RowKey									
Timestamp									

Correct Answer:

Answer Area

Event	Property								
Monitoring event 1	<table><tr><td></td><td>▼</td></tr><tr><td>PartitionKey</td><td></td></tr><tr><td>RowKey</td><td></td></tr><tr><td>Timestamp</td><td></td></tr></table>		▼	PartitionKey		RowKey		Timestamp	
	▼								
PartitionKey									
RowKey									
Timestamp									
Monitoring event 2	<table><tr><td></td><td>▼</td></tr><tr><td>PartitionKey</td><td></td></tr><tr><td>RowKey</td><td></td></tr><tr><td>Timestamp</td><td></td></tr></table>		▼	PartitionKey		RowKey		Timestamp	
	▼								
PartitionKey									
RowKey									
Timestamp									

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Monitoring event 1: PartitionKey

An entity group transaction must meet the following requirement:

All entities subject to operations as part of the transaction must have the same PartitionKey value.

Monitoring event 2: RowKey

Scenario: You must monitor the application at all times. You have two monitoring events, Monitoring event 1 and Monitoring event 2.

- Monitoring event 1 monitors entity group transactions.
- Monitoring event 2 monitors entity versioning.

References: <https://docs.microsoft.com/en-us/rest/api/storageservices/performing-entity-group-transactions>

Question Set 1

QUESTION 1

You deploy a website to Azure. When the website starts, it loads and caches common data.

Updates to the website must occur without downtime or performance degradation that is noticeable to users.

You need to upgrade to a new version of website code.

What should you do?

- ☐ A. Create a staging slot for the new version of the website. Run the following Windows PowerShell command, and then deploy the new code.

```
Switch-AzureWebsiteSlot -Name "MyWebsiteName"
```

- ☐ B. Create a staging slot for the new version of the website. Deploy the new code to that slot. Then run the following Windows PowerShell command:

```
Switch-AzureWebsiteSlot -Name "MyWebsiteName"
```

- ☐ C. Run the following Windows PowerShell command:

```
New-AzureWebsite -Name "Staging" -Location "East US"
```

Deploy the new code to the staging site. Then run the following Windows PowerShell command:

```
Switch-AzureWebsiteSlot -Name "MyWebsiteName"
```

- ☐ D. Create a staging slot for the new version of the website. Run the following Windows PowerShell command:

```
Switch-AzureWebsiteSlot -Name "MyWebsiteName"
```

Then deploy the new code to the staging slot.

- A. Option A
B. Option B

- C. Option C
- D. Option D

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 2

You deploy a stateless ASP.NET application to an Azure website. You scale out the application by adding website instances.

Only newly signed in users are routed to the recently added website instances. Users must be evenly distributed among all of the instances.

You need to configure the environment to ensure that the load balancer evenly distributes requests.

What should you do?

- ☐ A. Add the following markup to the web.config file for the application:

```
<system.webServer>
  <httpProtocol>
    <customHeaders>
      <add name="Arr-Disable-Session-Affinity" value="False" />
    </customHeaders>
  </httpProtocol>
</system.webServer>
```

- ☐ B. Configure autoscaling rules based on metrics.

- ☐ C. Add the following markup to the web.config file for the application:

```
<system.webServer>
  <httpProtocol>
    <customHeaders>
      <add name="Arr-Disable-Session-Affinity" value="True" />
    </customHeaders>
  </httpProtocol>
</system.webServer>
```

- ☐ D. Enable Always On support.

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: C
Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 3

You maintain an application that is used by local food delivery companies. When a customer requests a delivery, the application sends a message to all of the delivery companies. One company accepts the request and fulfills the order.

The application currently supports orders of 100 products or fewer. Some of the delivery companies can now deliver large orders that contain up to 500 products.

You must modify the application so that it supports both small orders and large orders. Messages about large orders should be sent to only delivery companies that can fulfill them. Messages about small orders should be sent to all delivery companies.

Which service should you use?

- A. Azure Service Bus Queue
- B. Azure Service Bus Relay
- C. Azure Service Bus Topics
- D. Azure Service Bus Namespace

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 4

DRAG DROP

You are developing a web application that integrates with Azure Active Directory (AD). The application uses the OAuth 2.0 protocol to authorize secure connections to a web service that is at <https://service.adatum.com>.

The application must request an access token to invoke the web service methods.



<https://www.gratisexam.com/>

You need to submit an HTTP request to the Azure AD endpoint.

How should you complete the request? To answer, drag the appropriate code segments to the correct locations. Each code segment 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:

HTTP Request Segments	Answer Area
<input type="text" value="adatum.com/oauth2/token"/>	POST <input type="text"/> HTTP/1.1
<input type="text" value="common/oauth2/token"/>	...
<input type="text" value="grant_type"/>	resource=https%3A%2F%2Fservice.adatum.com%2F&
<input type="text" value="response_type"/>	<input type="text"/> = <input type="text"/> &
<input type="text" value="client_credentials"/>	<input type="text"/> =F92FC9B0-F515-433E-BB72-482CC2303E62
<input type="text" value="client_id"/>	

Correct Answer:

HTTP Request Segments	Answer Area
<input type="text"/>	POST <input type="text" value="adatum.com/oauth2/token"/> HTTP/1.1
<input type="text" value="common/oauth2/token"/>	...
<input type="text"/>	resource=https%3A%2F%2Fservice.adatum.com%2F&
<input type="text" value="response_type"/>	<input type="text" value="grant_type"/> = <input type="text" value="client_credentials"/> &
<input type="text"/>	<input type="text" value="client_id"/> = F92FC9B0-F515-433E-BB72-482CC2303E62
<input type="text"/>	

Section: [none]

Explanation

Explanation/Reference:

QUESTION 5

A company plans to increase its virtual network capacity by adding virtual network subscriptions.

You must increase the number of subscriptions from 3 to 15.

You need to configure the virtual networks.

What should you do?

- A. Export and modify the network configuration file. Then import the modified file.
- B. Export and modify the service definition file. Then import the modified file.
- C. Create and import a new network configuration file.
- D. Create a multi-site virtual network.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 6

You are maintaining an application that uses the Azure Content Delivery Network (CDN) to serve terabytes of content that is stored in page blobs.

Your bill for CDN services is higher than you expect.

You need to monitor the application to find issues that increase costs.

Which two operations should you monitor? Each correct answer presents part of the solution.

- A. The Time-To-Live (TTL) of the blobs.
- B. The country of origin for the client computer and the CDN region.
- C. The number of requests that result in an HTTP status code over 400.
- D. The allocated size of page blobs.
- E. The expiration date of the blobs.

Correct Answer: BD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 7

You create a software-as-a-service (SaaS) application. Websites, cloud services, and virtual machines (VMs) read common data values from the database for the application.

The application does not scale efficiently. All VMs, websites, and cloud services must read from the same data source.

You need to design a cache solution for the SaaS application.

What should you do?

- A. Deploy a cache by using Azure Redis Cache. Access the cache from the websites, cloud services, and VMs.
- B. Configure a cache by using ASP.NET. Access the cache from the websites, cloud services, and VMs.
- C. Use Azure Redis Cache to deploy one cache for each website, one cache for each cloud service, and one cache for each VM. Configure each cache to ensure that data is consistent in all the cache instances.
- D. Deploy a cache by using Azure Redis Cache. Configure the cache to use database connection strings.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 8

You are modifying a web application so that it uses Azure Active Directory to manage users. You create a security group named Users and a security group named Administrators. The Administrators security group is a member of the Users security group.

You create the following code segment. Line numbers are included for reference only.

```
01 function canAccessUserResources(userId) {  
02  
03 }  
04 function getGroupId(groupName) {  
05 ...  
06 }  
07 function domain() {  
08 ...  
09 }
```

You need to implement the canAccessUserResources function.

Which code segment should you insert at line 02?

- ☐ A.

```
var groupId = getGroupId("Users");
var link = domain().concat("/users/", userId, "/memberOf?api-version=2013-04-05");
var json = $.getJSON(link);
for (entry in json.Value)
    if (entry.objectId == groupId)
        return true;
return false;
```
- ☐ B.

```
var groupId = getGroupId("Users");
var link = domain().concat("/isMemberOf?api-version=2013-04-05");
var json = $.post(link, { groupId: groupId, memberId: userId });
return json.value;
```
- ☐ C.

```
var groupId = getGroupId("User");
var link = domain().concat("/roles/", groupId, "?api-version=2013-04-05");
var json = $.getJSON(link);
return json.value;
```
- ☐ D.

```
var groupId = getGroupId("Users");
var link = domain().concat("/groups/", groupId, "/members?api-version=2013-04-05");
var json = $.getJSON(link);
for (entry in json.Value)
    if (entry.objectId == userId)
        return true;
return false;
```

- A. Option A
B. Option B
C. Option C
D. Option D

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 9

You are developing a messaging solution to integrate two applications named WeatherSummary and WeatherDetails. The WeatherSummary application displays a

summary of weather information for major cities. The WeatherDetails application displays weather details for a specific city.

You need to ensure that the WeatherDetails application displays the weather details for the city that the user selects in the WeatherSummary application.

What should you do?

- A. Create an Azure Service Bus Queue communication. In the WeatherDetails application, implement the PeekLock method.
- B. Create an Azure Service Bus Topics object. In the WeatherDetails application, create a filter.
- C. Create an Azure Service Bus Relay object. In the WeatherDetails application, create a filter.
- D. Create an Azure Service Bus Queue communication. In the WeatherDetails application, implement the ReceiveAndDelete method.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 10

You store data by using table storage in Azure.

The storage analytics logs do not contain any data.

You must configure the Azure storage account to retain logs for the maximum length of time that Azure permits.

In the Azure management portal, what should you do?

- A. Set the monitoring level to Minimal, and set the number of days the data in the logs is retained to 0.
- B. Set the monitoring level to Verbose, and set the number of days the data in the logs is retained to 365.
- C. Set the monitoring level to Minimal, and set the number of days the data in the logs is retained to 99.
- D. Set the monitoring level to Verbose, and set the number of days the data in the logs is retained to 30.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Reference: <http://azure.microsoft.com/en-gb/documentation/articles/storage-monitor-storage-account/>

QUESTION 11

You host an application on an Azure virtual machine (VM) that uses a data disk. The application performs several input and output operations per second.

You need to disable disk caching for the data disk.

Which two actions will achieve the goal? Each answer presents a complete solution.

- A. Use the Azure Resource Manager REST API
- B. Use the Service Management REST API.
- C. Run the following Windows PowerShell cmdlet:Remove-AzureDataDisk
- D. Run the following Windows PowerShell cmdlet:Set-AzureDataDisk

Correct Answer: AD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Reference: <http://msdn.microsoft.com/en-us/library/azure/jj157190.aspx>

QUESTION 12

You are developing a REST API service that provides data about products. The service will be hosted in an Azure virtual machine (VM).

The product data must be stored in Azure tables and replicated to multiple geographic locations.

API calls that use the HTTP GET operation must continue to function when the data tables at the primary Azure datacenter are not accessible.

You need to configure storage for the service.

Which type of replication should you choose?

- A. Locally Redundant Storage replication
- B. Geo-Redundant Storage replication
- C. Zone-Redundant Storage replication
- D. Read-Access Geo-Redundant Storage replication

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 13

You are migrating an existing solution to Azure. The solution includes a user interface tier and a database tier. The user interface tier runs on multiple virtual machines (VMs). The user interface tier has a website that uses Node.js. The user interface tier has a background process that uses Python. This background process runs as a scheduled job. The user interface tier is updated frequently. The database tier uses a self-hosted MySQL database.

The user interface tier requires up to 25 CPU cores. You must be able to revert the user interface tier to a previous version if updates to the website cause technical problems. The database requires up to 50 GB of memory. The database must run in a single VM.

You need to deploy the solution to Azure.

What should you do first?

- A. Deploy the entire solution to an Azure Web App. Use a web job that runs continuously to host the database.
- B. Configure Microsoft Visual Team Services to continuously deploy the user interface tier to the Azure Web App service. Deploy the production builds and the staging builds of the user interface tier to separate slots.
- C. Deploy the entire solution to an Azure Web App. Use a web job that runs continuously to host the user interface tier.
- D. Deploy the user interface tier to a VM. Use multiple availability sets to continuously deploy updates from Microsoft Visual Studio Online.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 14

You store data in an Azure blob. Data accumulates at a rate of 0.10 GB per day.

You must use storage analytics data to verify that the service level agreement (SLA) has been met and to analyze the performance of VHDs, including the pattern of usage.

Analytics data must be deleted when it is older than 100 days or when the total amount of data exceeds 10 GB.

You need to configure storage analytics and access the storage analytics data.

Which two approaches will achieve the goal? Each correct answer presents part of the solution.

- A. Disable the data retention policy.
- B. Access analytics data by using the Service Management REST APL
- C. Access analytics data by using the APIs used to read blob and table data.
- D. Configure a data retention policy of 100 days.

Correct Answer: CD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 15

You develop a web application that will use the Azure Table service. The web application will store entities in the form of XML data within a single table.

The web application must support high traffic throughput.

You need to avoid exceeding the throttle limit for the table.

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

- A. Add additional partition keys to the table.
- B. Batch transactions for entities that are in the same partition group in the table.
- C. Compress the entities before storing them in the table.
- D. Store the entities in JSON format.

Correct Answer: BD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 16

You are managing an application. The application uses data that is stored in an Azure SQL database.

You must be able to reset the application to the state that existed on any day in the previous 35 days.

You need to choose a backup solution.

What should you do?

- A. Run SQL replication on the SQL database once a day.
- B. Use Microsoft Azure SQL Database Point in Time Restore
- C. Use the SQL Server Data-Tier Application Framework to build a data-tier application (DAC) file once a day.
- D. Use the bcp utility to export data to an Azure page blob once a day.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 17

You develop a web application that uses table storage in Azure.

You create a storage account named Contoso that stores a table named CityPopulationData.

The web application stores entities in this table.

You need to query the table data by using OData.

Which URL should you use?

- A. <http://contoso.table.core.windows.net/citypopulationdata>
- B. <http://contoso.table.core.windows.net/odata/citypopulationdata>
- C. <http://azurestorage.table.core.windows.net/contoso>
- D. <http://microsoft.table.core.windows.net/contoso>
- E. <http://azure.table.core.windows.net/contoso/citypopulationdata>

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 18

HOTSPOT

Your company works with trusted partners. These partners upload files into a storage account that you control.

Partners must be able to create, read, and write files. Partners must NOT be allowed to see files from other partners. You generate a shared access signature (SAS) for each partner.



<https://www.gratisexam.com/>

You create the following Windows PowerShell script to create a new container for each partner. Line numbers are included for reference only.

```
01 $containerName = "partner123files"
02 $key = (Get-AzureStorageKey -StorageAccountName $storageAccountName).Primary
03 $context = New-AzureStorageContext -StorageAccountName $storageAccountName `
    -StorageAccountKey $key
05 New-AzureStorageContainer -Name $containerName -Context $context
06 $filepath = "welcome.txt"
07 $blobname = "welcome.txt"
08 Set-AzureStorageBlobContent -Container $containerName -File "$filepath" `
    -Blob $blobname -Context $context -Properties @{"ContentType"="text/plain"}
09 $oneYearFromNow = (Get-Date).AddYears(1)
10 $sasToken = New-AzureStorageContainerSASToken -Name $containerName `
    -Permission 'rwdl' -ExpiryTime $oneYearFromNow -Context $context
11 $sasBlobUri = New-AzureStorageBlobSASToken -Container $containerName `
    -Permission 'r' -ExpiryTime $oneYearFromNow -Context $context `
    -FullUri -Blob $blobname
```

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

Hot Area:

Answer Area

	Yes	No
Running the command at line 10 a second time invalidates the previously generated SAS token.	<input type="radio"/>	<input type="radio"/>
Web browsers can open the welcome.txt file directly by using the full URI and the SAS token for the file.	<input type="radio"/>	<input type="radio"/>
If the primary storage key is regenerated, the SAS token is still valid until its expiration date is reached.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

	Yes	No
Running the command at line 10 a second time invalidates the previously generated SAS token.	<input type="radio"/>	<input checked="" type="radio"/>
Web browsers can open the welcome.txt file directly by using the full URI and the SAS token for the file.	<input checked="" type="radio"/>	<input type="radio"/>
If the primary storage key is regenerated, the SAS token is still valid until its expiration date is reached.	<input type="radio"/>	<input checked="" type="radio"/>

Section: [none]
Explanation

Explanation/Reference:

QUESTION 19

A company maintains an Azure storage account. The storage account uses blobs and tables.

Customers access the storage account by using shared access signatures (SASs).

You need to monitor the usage of the storage services. You need to do the following:

- Understand which storage areas perform operations that incur an Azure fee.
- Understand which requests are denied because of insufficient permissions.
- Validate that the performance of the storage account meets the service level agreement (SLA) for the Azure Storage service.

Which three data analysis tasks should you perform? Each correct answer presents part of the solution.

- A. Use data from the logs of the storage services to find individual storage access attempts that do not comply with the SLA.
- B. Use data from the logs of the storage services to calculate aggregate server latency across individual requests. Determine whether the results of this calculation indicate that the Azure Storage service is in compliance with the SLA.
- C. Analyze the logs of the storage services to determine which storage services were inaccessible because of permissions issues.
- D. Review the Azure documentation to determine which storage operations are billable. Then find records of those operations in the logs of the storage services.
- E. Analyze the logs of the storage services to find records of operations that are marked as billable.
- F. Correlate the data logged from the storage service with the permissions to store data in the individual blobs and containers. Determine which storage services were inaccessible because of permissions issues.

Correct Answer: BCD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 20

You are creating virtual machines (VMs) that are hosted on Azure.

You must be able to change the Remote Desktop access settings for the VMs. You must also be able to change the password for the built-in administrator account on all VMs. You identify the VMAccess VM extensions that have the required capabilities.

You need to enable the VMAccess VM extensions.

Which approach should you use?

- A. Download and install the Microsoft Installer file to enable the VM Agent on each VM.
- B. Use the Azure management portal to restart each VM.
- C. When you configure the new VMs, use the Azure management portal to install the VM Agent.

D. For each VM, use Windows PowerShell cmdlets to enable the VM Agent and the VMAccess VM extensions.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 21

You connect to an existing service over the network by using HTTP. The service listens on HTTP port 80. You plan to create a test environment for this existing service by using an Azure virtual machine (VM) that runs Windows Server.

The service must be accessible from the public Internet over HTTP port 8080.

You need to configure the test environment.

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

- A. Configure an endpoint to route traffic from port 8080 to port 80.
- B. Configure an endpoint to route traffic from port 80 to port 8080.
- C. Ensure that the public IP address is configured as a static IP address.
- D. Configure the Windows Server firewall to allow incoming and outgoing traffic on port 8080.
- E. Configure the Windows Server firewall to allow incoming and outgoing traffic on port 80.

Correct Answer: AE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 22

An application sends Azure push notifications to a client application that runs on Windows Phone, iOS, and Android devices.

Users cannot use the application on some devices. The authentication mechanisms that the application uses are the source of the problem.

You need to monitor the number of notifications that failed because of authentication errors.

Which three metrics should you monitor? Each correct answer presents part of the solution.

- A. Microsoft Push Notification Service (MPNS) authentication errors
- B. External notification system errors
- C. Apple Push Notification Service (APNS) authentication errors
- D. Channel errors
- E. Windows Push Notification Services (WNS) authentication errors
- F. Google Cloud Messaging (GCM) authentication errors

Correct Answer: ACF

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 23

You deploy an application as a cloud service in Azure.

The application consists of five instances of a web role.

You need to move the web role instances to a different subnet.

Which file should you update?

- A. Service definition
- B. Diagnostics configuration
- C. Service configuration
- D. Network configuration

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 24

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

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

You plan to deploy an application as a cloud service. The application uses a virtual network to extend your on-premises network into Azure.

You need to configure a site-to-site VPN for cross-premises network connections.

Which two objects should you configure? Each correct answer presents part of the solution.

- A. Dynamic routing gateway
- B. External-facing IPv6 address
- C. VPN gateway
- D. External-facing IPv4 address

Correct Answer: AD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

A Site-to-Site VPN gateway connection is used to connect your on-premises network to an Azure virtual network over an IPsec/IKE (IKEv1 or IKEv2) VPN tunnel.

Verify that you have an externally facing public IPv4 address for your VPN device. This IP address cannot be located behind a NAT.

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

QUESTION 25

A company creates an API and makes it accessible on an Azure Web App. External partners use the API occasionally. The Web App uses the Standard web hosting plan.

Partners report that the first API call in a sequence of API calls occasionally takes longer than expected to run. Subsequent API calls consistently perform as expected.

You need to ensure that all API calls perform consistently.

What should you do?

- A. Configure the Web App to use the Basic web hosting plan.
- B. Enable Always On support.
- C. Configure the Web App to automatically scale.
- D. Add a trigger to the web.config file for the Web App that causes the website to recycle periodically.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 26

DRAG DROP

You deploy an application as a cloud service to Azure. The application contains a web role to convert temperatures between Celsius and Fahrenheit.

The application does not correctly convert temperatures. You must use Microsoft Visual Studio to determine why the application does not correctly convert temperatures.

You need to debug the source code in 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:

Actions	Answer Area
<div data-bbox="313 274 1084 363">Attach the debugger to the role instance of the cloud service.</div> <div data-bbox="313 386 1084 443">Publish the application.</div> <div data-bbox="313 466 1084 596">In the Microsoft Azure Publish Settings dialog, set the build configuration to Release and enable the remote debugger for all roles.</div> <div data-bbox="313 619 1084 718">In the Windows Azure Publish Settings dialog, set the build configuration to Debug.</div> <div data-bbox="313 740 1084 871">In the Microsoft Azure Publish Settings dialog, enable Remote Desktop for cloud configuration and enable the remote debugger for all roles.</div>	

Correct Answer:

Actions	Answer Area
	<div>Publish the application.</div>
	<div>In the Microsoft Azure Publish Settings dialog, set the build configuration to Release and enable the remote debugger for all roles.</div>
<div>In the Windows Azure Publish Settings dialog, set the build configuration to Debug.</div>	<div>Attach the debugger to the role instance of the cloud service.</div>
<div>In the Microsoft Azure Publish Settings dialog, enable Remote Desktop for cloud configuration and enable the remote debugger for all roles.</div>	

Section: [none]

Explanation

Explanation/Reference:

QUESTION 27

HOTSPOT

You are developing an Azure cloud service for a company. The cloud service monitors a queue for incoming messages and then processes invoices based on the contents of these messages.

Some messages are formed incorrectly and cause exceptions. There is no time limit for how long the service takes to process an individual message.

All messages must be processed at least once by using the ProcessMessage method. Messages must not be processed more than twice by using the ProcessMessage method. Messages that fail normal processing must be processed by using the ProcessPoisonMessage method.

You need to configure message processing.

How should you complete the relevant code? To answer, select the appropriate option or options in the answer area.

Hot Area:

Answer Area

```
private bool ProcessNextQueueMessage(CloudQueue cloudQueue)
{
    var msg = cloudQueue.GetMessage();
```

```
if (msg == null) return false;
if (msg.DequeueCount > 0) return false;
if (msg.PopReceipt == null) return false;
if (msg.ExpirationTime.HasValue) return false;
```

```
if (msg == null)
if (msg.DequeueCount > 0)
if (msg.DequeueCount > 2)
if (msg.PopReceipt == null)
```

```
    ProcessPoisonMessage(msg);
else
    ProcessMessage(msg);
```

```
cloudQueue.Delete();
cloudQueue.DeleteMessage(msg);
cloudQueue.EndAddMessage(null);
cloudQueue.DeleteMessage(null);
```

```
    return true;
```

Correct Answer:

Answer Area

```
private bool ProcessNextQueueMessage(CloudQueue cloudQueue)
{
    var msg = cloudQueue.GetMessage();
```

```
if (msg == null) return false;
if (msg.DequeueCount > 0) return false;
if (msg.PopReceipt == null) return false;
if (msg.ExpirationTime.HasValue) return false;
```

```
if (msg == null)
if (msg.DequeueCount > 0)
if (msg.DequeueCount > 2)
if (msg.PopReceipt == null)
```

```
    ProcessPoisonMessage(msg);
else
    ProcessMessage(msg);
```

```
cloudQueue.Delete();
cloudQueue.DeleteMessage(msg);
cloudQueue.EndAddMessage(null);
cloudQueue.DeleteMessage(null);
```

```
    return true;
```

Section: [none]

Explanation

Explanation/Reference:

QUESTION 28

Which of the following applications would be a good candidate to move to a cloud-based platform?

- A. Mission critical financial data
- B. Customer Relationship Management (CRM)
- C. High-performance computing
- D. Database that requires a low latency for indexing

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The best place to start is with new applications that are customer-, partner- and employee-facing.

Cloud CRM (or CRM cloud) means any customer relationship management (CRM) technology where the CRM software, CRM tools and the organization's customer data resides in the cloud and is delivered to end-users via the Internet.

Cloud CRM typically offers access to the application via Web-based tools (or Web browser) logins where the CRM system administrator has previously defined access levels across the organization. Employees can log in to the CRM system, simultaneously, from any Internet-enabled computer or device. Often, cloud CRM provide users with mobile apps to make it easier to use the CRM on smartphones and tablets.

References:

<https://azure.microsoft.com/en-us/blog/a-key-it-decision-which-apps-to-move-to-the-cloud/>

http://www.webopedia.com/TERM/C/crm_cloud.html

QUESTION 29

Companies that are looking to move from capital expenses to operating expenses benefit from cloud services.

- A. True
- B. False

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

"Capex vs. Opex" refers to the fact that stocking your own data center requires capital expenditure, while using an external cloud service that offers pay-as-you-go service falls into ongoing operating expenditures: thus the contrast of "Capex vs. Opex."

References: <http://www.cio.com/article/2430099/virtualization/capex-vs--opex--most-people-miss-the-point-about-cloud-economics.html>

QUESTION 30

A private cloud is defined as:

- A. A deployment model that uses an external cloud to provide host application services that are Internet accessible.
- B. A deployment model that partners with other industry related companies to provide infrastructure services.
- C. A deployment model that uses virtualization technologies to provide infrastructure on demand within its network.
- D. A deployment model that uses an external cloud provider to provide host infrastructure services that are Internet accessible.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Private cloud is a type of cloud computing that delivers similar advantages to public cloud, including scalability and self-service, but through a proprietary architecture. Unlike public clouds, which deliver services to multiple organizations, a private cloud is dedicated to a single organization.

Private cloud expenses include virtualization, cloud software and cloud management tools.

References: <http://searchcloudcomputing.techtarget.com/definition/private-cloud>

QUESTION 31

Which of the following is the logical progression in internal private cloud adoption?

- A. Virtualize, PaaS, IaaS and SaaS
- B. SaaS, PaaS, IaaS and Virtualize
- C. Virtualize, IaaS, PaaS and SaaS
- D. IaaS, PaaS, Virtualize and SaaS

Correct Answer: C

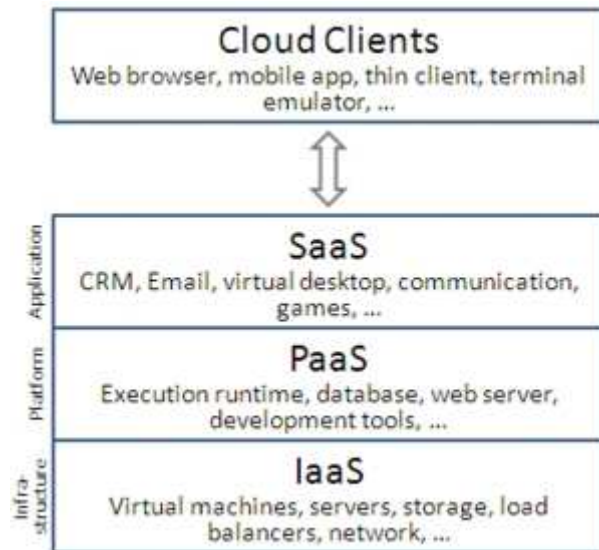
Section: [none]

Explanation

Explanation/Reference:

Explanation:

Cloud computing service models arranged as layers in a stack.



References: https://en.wikipedia.org/wiki/Cloud_computing#Service_models

QUESTION 32

A cloud computing vendor is focusing on delivering applications to customers. The goal is to simplify the deployment of database functionality while removing the need for customers to manage the operation system and application patching. Which of the following types of solution is the vendor offering?

- A. IT as a Service
- B. Infrastructure as a Service
- C. Anything as a Service
- D. Platform as a Service
- E. Software as a Service

Correct Answer: D

Section: [none]

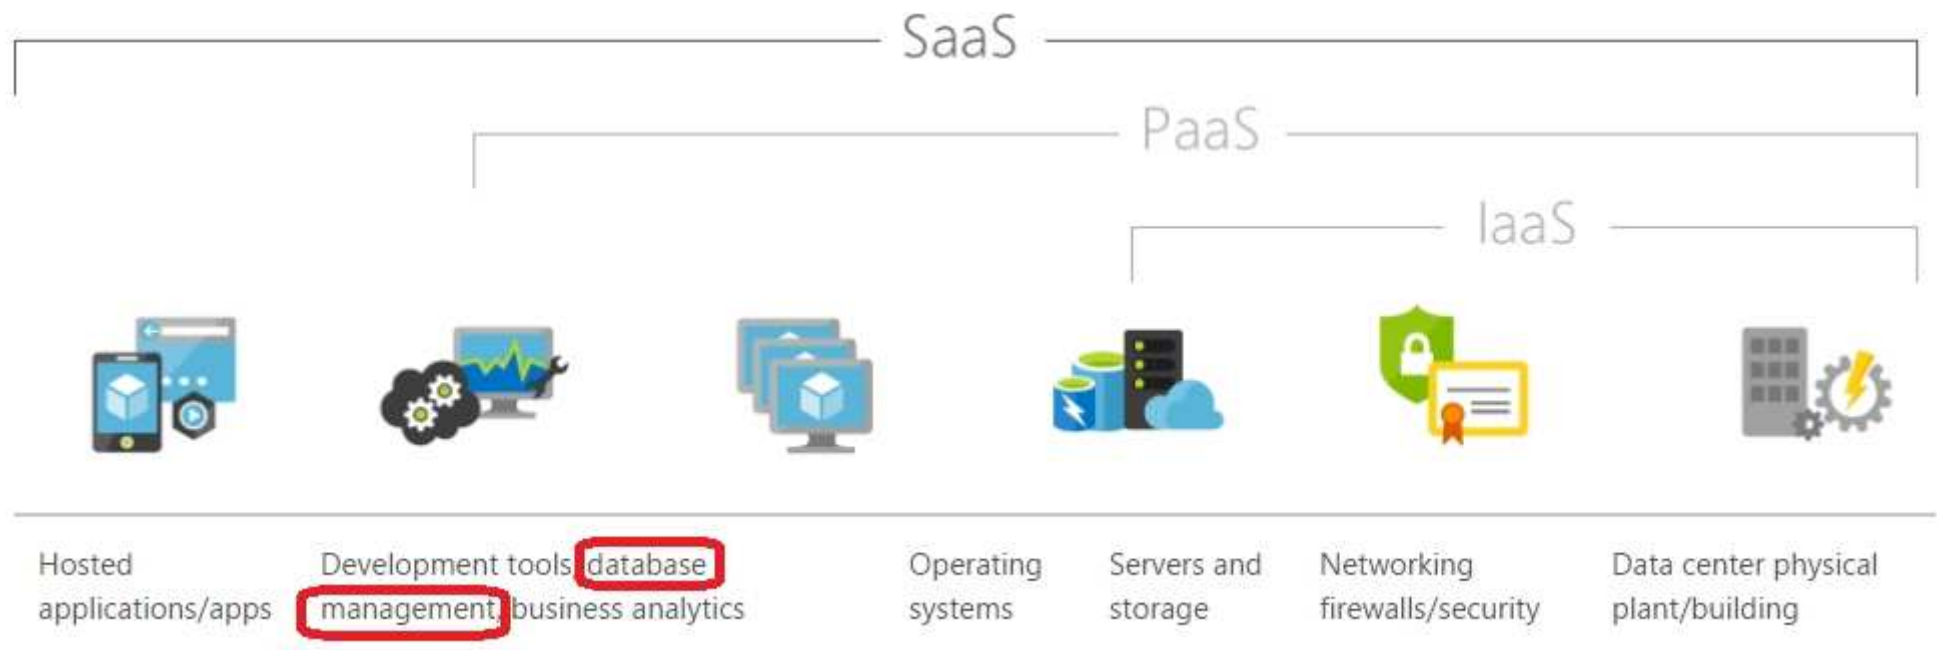
Explanation

Explanation/Reference:

Explanation:

PaaS includes infrastructure — servers, storage, and networking — but also middleware, development tools, business intelligence (BI) services, database management systems, and more.

Note:



References: <https://azure.microsoft.com/en-us/overview/what-is-paas/>

QUESTION 33

Which of the following describes what is meant by the ITIL Service Strategy component?

- A. Defining processes required to manage the solution.
- B. Designing the solution to the ITIL specifications.
- C. Ensuring changes are designed to meet customer expectations.

D. Understanding the intended customer and what services are required.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The objective of ITIL Service Strategy is to decide on a strategy to serve customers. Starting from an assessment of customer needs and the market place, the Service Strategy lifecycle stage determines which services the IT organization is to offer and what capabilities need to be developed. Its ultimate goal is to make the IT organization think and act in a strategic manner.

References: http://wiki.en.it-processmaps.com/index.php/ITIL_Service_Strategy

QUESTION 34

Using https instead of http for accessing a cloud service is considered more secure.

A. True

B. False

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

HTTPS (also called HTTP over TLS, HTTP over SSL, and HTTP Secure) is a protocol for secure communication over a computer network which is widely used on the Internet.

References: <https://en.wikipedia.org/wiki/HTTPS>

QUESTION 35

A company is designing a new web-based software application that must be highly available and resistant. Which of the following is the BEST environment for the application?

A. The primary instance of the application will be locally hosted with a weekly copy of the instance sent to a cloud service provider.

B. The primary instance of the application will be locally hosted with a nightly file-level backup being performed to an off-site location.

C. The primary instance of the application will be running in a cloud service provider's hosted environment with a continuous backup to the company's local infrastructure.

D. The primary instance of the application will be locally hosted with a nightly copy of the instance sent to a client service provider.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 36

Which of the following are the MOST important benefits of a cloud computing solution for an application development provider? (Select two.)

- A. Reduced training time for new developers
- B. Reduced storage requirements.
- C. Reduced complexity for users.
- D. Reduced bandwidth usage.
- E. Reduced cost.
- F. Reduced development timeframe.

Correct Answer: EF

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The biggest promise of Azure-based applications is the ability to write them to scale as needed in real-time. Customers will therefore only use the amount of resources they need, rather than budgeting a set amount of resources that can overtax or underutilize their current setup.

References: <http://searchcloudcomputing.techtarget.com/tutorial/An-introduction-to-developing-for-Microsoft-Azure>

QUESTION 37

Which of the following virtualization characteristics allows the use of different types of physical types or physical servers?

- A. Security
- B. Hardware independence
- C. Scalability
- D. Variable costs

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Virtualization is a conversion process that translates unique IT hardware into emulated and standardized software-based copies. Through hardware independence, virtual servers can easily be moved to another virtualization host, automatically resolving multiple hardware-software incompatibility issues. As a result, cloning and manipulating virtual IT resources is much easier than duplicating physical hardware.

References: http://whatiscloud.com/virtualization_technology/hardware_independence

QUESTION 38

Which of the following cloud computing services requires the MOST involvement from a company's in-house staff?

- A. IaaS
- B. MaaS
- C. PaaS
- D. SaaS

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Infrastructure as a service (IaaS) is an instant computing infrastructure, provisioned and managed over the Internet. Quickly scale up and down with demand, and pay only for what you use.

IaaS helps you avoid the expense and complexity of buying and managing your own physical servers and other datacenter infrastructure. Each resource is offered as a separate service component, and you only need to rent a particular one for as long as you need it. The cloud computing service provider manages the infrastructure, while you purchase, install, configure, and manage your own software—operating systems, middleware, and applications.

References: <https://azure.microsoft.com/en-us/overview/what-is-iaas/>

QUESTION 39

As part of a cloud provider's services, customers can provision a new virtual machine as needed without human interaction with the provider. The scenario is BEST described by which of the following cloud characteristics?

- A. On-demand self-service
- B. Measured service
- C. Broad network access

D. Rapid elasticity

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

On-demand self service refers to the service provided by cloud computing vendors that enables the provision of cloud resources on demand whenever they are required. In on-demand self service, the user accesses cloud services through an online control panel.

On-demand self service resource sourcing is a prime feature of most cloud offerings where the user can scale the required infrastructure up to a substantial level without disrupting the host operations.

References: <https://www.techopedia.com/definition/27915/on-demand-self-service>

QUESTION 40

A business has recently implemented a hybrid cloud federated solution, which will allow it to rapidly and dynamically allocate resources during high demand, and quickly implement its Disaster Recovery Plan (DRP) and Continuity of Operations (COOP). Given this implementation, the IT director is mostly likely concerned about:

- A. maintaining strategic flexibility
- B. reducing OPEX allocations
- C. hiring additional IT staff
- D. eliminating security risks

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Initially, the differences between Disaster Recover and Business Continuity should be understood. In comparison they are very similar in that they are (or should be) detailed plans to prepare an organization for events in which a situation presents itself which can cause internal systems failures, or a disruption of business systems in which they are no longer able to function to meet the requirements to perform day to day tasks. These situations almost always result in loss of revenue, and in some cases, loss of client base. Where these plans differ is in the main concept topic for which they prepare. Business Continuity Plans generally focus on the continuation of business services in the event of any type of interruptions whether its IT based or other. Disaster Recover Plans often refer to a company's strategy if something happens to crucial business data, and how to restore / recover that data (generally in the shortest amount of time possible).

References: <https://stumpj.wordpress.com/2010/10/18/coop-and-drp-what-is-the-difference/>

QUESTION 41

Which of the following is the primary difference between private and public cloud?

- A. Tenancy of the cloud
- B. Management of the cloud
- C. Service model of the cloud
- D. Locations on the cloud

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

A private cloud hosting solution, also known as an internal or enterprise cloud, resides on company's intranet or hosted data center where all of your data is protected behind a firewall.

The main differentiator between public and private clouds is that you aren't responsible for any of the management of a public cloud hosting solution. Your data is stored in the provider's data center and the provider is responsible for the management and maintenance of the data center.

References: <https://www.expedient.com/blog/private-vs-public-cloud-whats-difference/>

QUESTION 42

After migrating the company's entire datacenter infrastructure to a private IaaS solution, while at the same time maintaining the current network and server logical configuration, the IT director eliminated 50% of the IT engineering staff. The remaining staff has now shifted focus from a daily server maintenance and upkeep role, to more of a service provisioning, performance, and reporting role. Which of the following was MOST impacted by this migration?



<https://www.gratisexam.com/>

- A. Service design
- B. Service strategy
- C. Service operation
- D. Service transitions

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 43

Which of the following enables hardware independence?

- A. In-sourcing
- B. Outsourcing
- C. Virtualization
- D. Abstraction

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Virtualization is a conversion process that translates unique IT hardware into emulated and standardized software-based copies. Through hardware independence, virtual servers can easily be moved to another virtualization host, automatically resolving multiple hardware-software incompatibility issues. As a result, cloning and manipulating virtual IT resources is much easier than duplicating physical hardware.

References: http://whatisccloud.com/virtualization_technology/hardware_independence

QUESTION 44

Which of the following are the common elements of platform as a service and software as a service? (Select two.)

- A. Both allow the OS to be patched by the customer.
- B. Both take advantage of incremental scalability.
- C. Both require the customer to maintain the hardware.
- D. Both provide granular access to the backend storage.
- E. Both implement hardware abstraction.

Correct Answer: BE

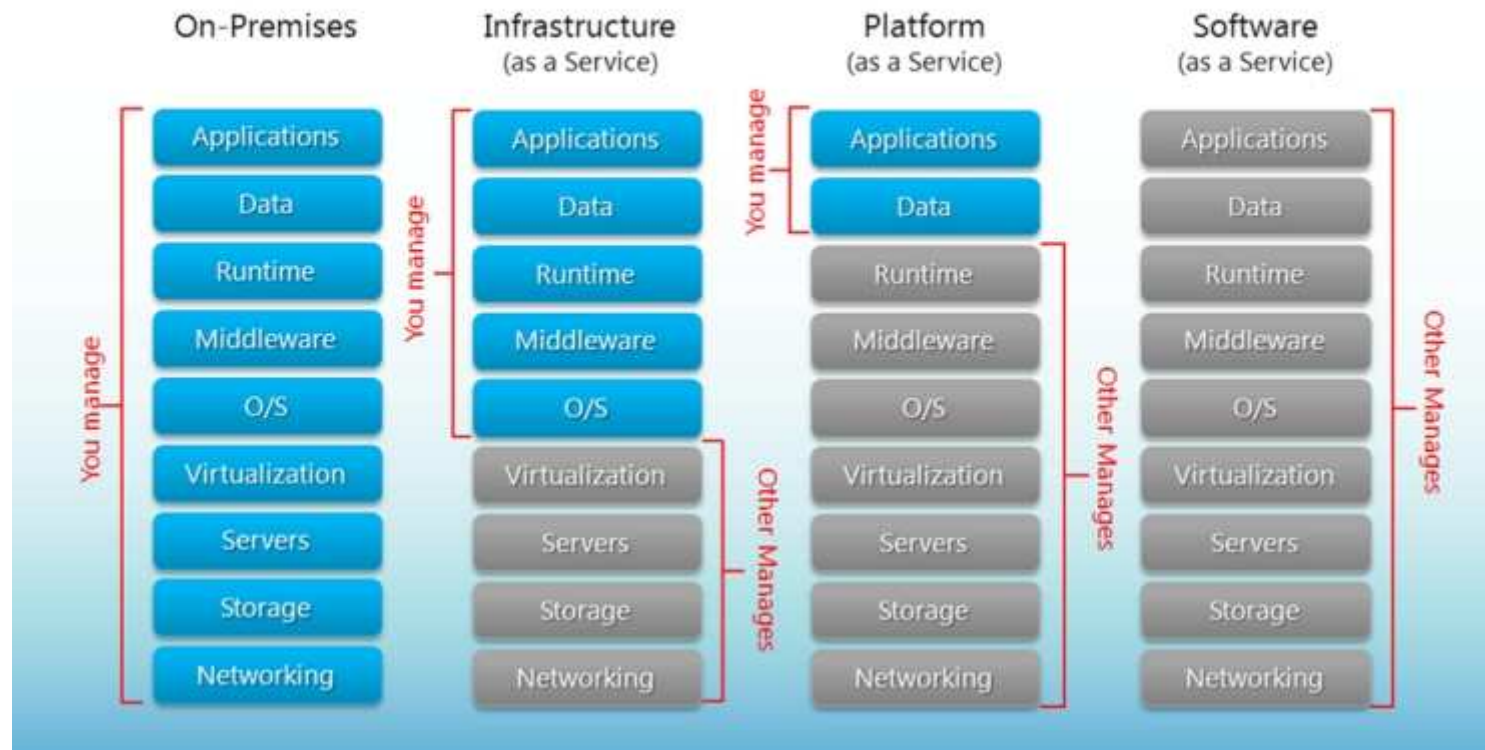
Section: [none]

Explanation

Explanation/Reference:

Explanation:

Separation of Responsibilities



QUESTION 45

When using SaaS, the cloud computing vendor is responsible to maintain which of the following?

- A. Client infrastructure
- B. Client firewall

- C. Updates and licenses.
- D. Workstation OS version.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 46

Which of the following is the MOST significant risk to business continuity when using an external cloud service provider?

- A. Unauthorized access to customer data
- B. Vendor being purchased
- C. Virtual server failure
- D. Vendor going out of business

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

If your application stores and retrieves very sensitive data, you might not be able to maintain it in the cloud. Similarly, compliance requirements could also limit your choices.

References: <http://cloudacademy.com/blog/cloud-migration-benefits-risks/>

QUESTION 47

Which of the following is an example of SaaS?

- A. Offshore help desk support
- B. Hosted database software and development tools
- C. Hosted email software
- D. Hosted network hardware

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

If you've used a web-based email service such as Outlook, Hotmail, or Yahoo! Mail, then you've already used a form of SaaS. With these services, you log into your account over the Internet, often from a web browser. The email software is located on the service provider's network, and your messages are stored there as well. You can access your email and stored messages from a web browser on any computer or Internet-connected device.

References: <https://azure.microsoft.com/en-us/overview/what-is-saas/>

QUESTION 48

Cloud computing relies heavily on which of the following virtualization characteristics? (Select two.)

- A. User federation
- B. Hardware independence
- C. Simplistic setup
- D. Scalable resources
- E. Information sharing

Correct Answer: BD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

B: Virtualization is a conversion process that translates unique IT hardware into emulated and standardized software-based copies. Through hardware independence, virtual servers can easily be moved to another virtualization host, automatically resolving multiple hardware-software incompatibility issues. As a result, cloning and manipulating virtual IT resources is much easier than duplicating physical hardware.

D: Infrastructure as a Service (IaaS) is a form of cloud computing that provides virtualized computing resources over the Internet. IaaS platforms offer highly scalable resources that can be adjusted on-demand.

References:

http://whatisccloud.com/virtualization_technology/hardware_independence

<http://searchcloudcomputing.techtarget.com/definition/Infrastructure-as-a-Service-IaaS>

QUESTION 49

Following an IT Service Management lifecycle approach, a Chief Information Officer would take which of the following paths to implement a cloud solution?

- A. Choose the SaaS provider, Design the application; Choose whether to develop the service application in-house or outsource; Operate the service application in

the cloud.

- B. Decide whether to implement on the cloud; Choose a XaaS provider; Design the application; Choose where to develop the service application; Operate the service application in the cloud.
- C. Decide whether to implement the application on the cloud; Choose an IaaS provider; Choose whether to develop the service in-house; Operate the Service application in the cloud.
- D. Strategize which IaaS provider to use; Design the application; Transition processes to the cloud; Operate the service application in the cloud.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 50

An organization wants to host a critical application on two redundant leased servers located on the ISP's datacenter. Which of the following is this an example of?

- A. PaaS
- B. IaaS
- C. Public cloud
- D. SaaS

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Infrastructure as a service (IaaS) is an instant computing infrastructure, provisioned and managed over the Internet.

IaaS helps you avoid the expense and complexity of buying and managing your own physical servers and other datacenter infrastructure. Each resource is offered as a separate service component, and you only need to rent a particular one for as long as you need it. The cloud computing service provider manages the infrastructure, while you purchase, install, configure, and manage your own software—operating systems, middleware, and applications.



References: <https://azure.microsoft.com/en-us/overview/what-is-iaas/>

QUESTION 51

A cloud usage metering scheme allows for which of the following customer chargeback alternatives?

- A. Cost allocation
- B. Cost amortization
- C. Shared cost
- D. Direct cost

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

CHARGEBACK METHODS

A range of approaches have been developed for implementing chargeback in an organization, as summarized in the figure below. The degree of complexity, degree of difficulty, and cost to implement decreases from the top of the chart [service-based pricing (SBP)], to the bottom [high-level allocation (HLA)]. HLA is the simplest method; it uses a straight division of IT costs based on a generic metric such as headcount. Slightly more effort to implement is low-level allocation (LLA), which bases consumer costs on something more related to IT activity such as the number of users or servers. Direct cost (DC) more closely resembles a time and materials charge but is often tied to headcount as well.

Figure, Methods for chargeback allocation.

METHOD	DESCRIPTION
Service Based Pricing (SBP)	Charges per a specific measured unit of service
Negotiated Flat Rate (NFR)	Charges based on a negotiated and often projected usage of a service
Tiered Flat Rate (TFR)	Charges based on providing access to a service whether the service is being used or not (fliers or bands pricing)
Measured Resource Usage (MRU)	Charges based on actual measured usage of specific IT resources (e.g., kW consumed, network bandwidth consumed, and storage consumed)
Direct Cost (DC)	Charges based on dedicated ownership of the resource (e.g., time and material based costing)
Low-level Allocation (LLA)	Charges based on simpler user metrics (e.g., user counts and server counts)
High-level Allocation (HLA)	Charges based on user size (e.g., number of employees and amount of revenue)

References: <https://journal.uptimeinstitute.com/it-chargeback-drives-efficiency/>

QUESTION 52

From a risk assessment perspective, which of the following is MOST important to acquire and review before a business integrates cloud computing into its existing environment?

- A. Cloud provider DRP and COOP
- B. The time to market expectation
- C. The total cost of ownership
- D. The company original RFP

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

A disaster recovery plan (DRP) - sometimes referred to as a business continuity plan (BCP) or business process contingency plan (BPCP) - describes how an organization is to deal with potential disasters.

QUESTION 53

Which of the following describes the commonality between cloud computing and outsourcing?

- A. Shift from CAPEX to OPEX
- B. Reduced compliance cost
- C. Simplified security management
- D. Reduced system architecture complexity.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

"Capex vs. Opex" refers to the fact that stocking your own data center requires capital expenditure, while using an external cloud service that offers pay-as-you-go service falls into ongoing operating expenditures: thus the contrast of "Capex vs. Opex."

References: <http://www.cio.com/article/2430099/virtualization/capex-vs--opex--most-people-miss-the-point-about-cloud-economics.html>

QUESTION 54

Virtual Desktop Interface (VDI) will present challenges for the network administrator as they move their users to the cloud. Which of the following would be considered a major challenge?

- A. Developing a backup environment for the end user
- B. Troubleshooting the users' applications

- C. Supporting multiple devices (e.g. tablets, thin clients)
- D. Centralizing the applications

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

References: <https://msdn.microsoft.com/en-us/library/dn903170.aspx>

QUESTION 55

A company Chief Information Officer (CIO) who wants to ensure rapid elasticity for the company's cloud solution would MOST likely choose which of the following types of cloud?

- A. Public cloud
- B. Private community cloud
- C. Private cloud
- D. Community cloud

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Rapid elasticity is a cloud computing term for scalable provisioning, or the ability to provide scalable services.

Software that can scale in a private cloud faces two security related issues:

- Although the private cloud infrastructure can enable rapid elasticity in the supply of virtual resources, hosted applications and services must be designed correctly if they are to function securely when they are scaled out.
- Hosted applications and services that initiate scaling requests automatically based on monitored demand or a timetable must perform these operations without impacting their own or other services availability within the cloud.

References: <http://social.technet.microsoft.com/wiki/contents/articles/6810.private-cloud-security-challenges-rapid-elasticity.aspx>

QUESTION 56

Consumption statistics for individual cloud service offerings is used by which of the following ITIL processes?

- A. Supplier management
- B. Continuous service improvement
- C. Service level management
- D. Information security management

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

ITIL defines Service Management as “a set of specialized organisational capabilities for providing value to customers in the form of services”.

The managed service provider will intermediate between the cloud service provider and consumer, aligning the two and ensuring minimal service disruptions.

ITIL Service Level Management aims to negotiate Service Level Agreements with the customers and to design services in accordance with the agreed service level targets. Service Level Management is also responsible for ensuring that all Operational Level Agreements and Underpinning Contracts are appropriate, and to monitor and report on service levels.

References:

<https://blog.kloud.com.au/2016/04/06/consumption-based-service-management/>

http://wiki.en.it-processmaps.com/index.php/Service_Level_Management

QUESTION 57

Which of the following is a potential advantage of using Storage as a Service?

- A. Data is accessible when the Internet is not functioning
- B. In-house IT staff controls all data
- C. Increase in encryption technologies
- D. Decrease in IT management of the platform

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

One advantages of SaaS is that it makes it easy to “mobilize” your workforce because users can access SaaS apps and data from any Internet-connected computer or mobile device. You don’t need to worry about developing apps to run on different types of computers and devices because the service provider has

already done so. In addition, you don't need to bring special expertise onboard to manage the security issues inherent in mobile computing. A carefully chosen service provider will ensure the security of your data, regardless of the type of device consuming it.

References: <https://azure.microsoft.com/en-us/overview/what-is-saas/>

QUESTION 58

Which of the following describes the difference between SaaS and IaaS?

- A. SaaS defines a standard while IaaS implements the standard.
- B. SaaS enables the software developer while IaaS provides the specifications.
- C. SaaS provides applications while IaaS provides equipment.
- D. SaaS implements security while IaaS provides information.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Software as a service (SaaS) allows users to connect to and use cloud-based apps over the Internet.



References: <https://azure.microsoft.com/en-us/overview/what-is-saas/>

QUESTION 59

Locating datacenters close to target markets is the result of which of the following strategic initiatives?

- A. Geoproximity
- B. Geodiversity
- C. Geography
- D. Localization

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure allows specification of geographical regions or affinity groups. Geographical regions are related to the data centers, like North Central US, South Central US,

Anywhere US, East Asia, North Europe, and so on. The list of options will grow as more data centers are added.

Instead of selecting a region, it is possible to specify an affinity group. Affinity groups are hints to Azure that essentially state that everything within the group should be as close in proximity as Azure will allow. That usually means keeping items within the same data center, which besides having the benefit of geo-location, can sometimes offer performance improvements for communication.

References: <https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-overview>

QUESTION 60

A company regularly doubles the number of employees over the summer by hiring on temporary stuff. The company currently pays the same price every month for its email software, equal to the maximum number of employees on staff who have email addresses. The company would like to only pay for the number of active email addresses and is considering migrating to a cloud solution. Which of the following characteristics of cloud computing is the MOST beneficial reason the company should implement a cloud solution?

- A. Scalability
- B. Hardware independence
- C. Variable costs
- D. Security

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 61

Which of the following is a benefit of public cloud computing?

- A. Enhances fixed expenditures for hardware and software
- B. Contributes to the quality of user input data
- C. Reduces OPEX costs for application and databases
- D. Adds flexibility and agility to enterprise architecture

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Public cloud computing offers greater flexibility, agility, and scalability.

QUESTION 62

Which of the following is the cloud characteristic that speeds up development, deployment and overall time of market?

- A. Rapid elasticity
- B. Cloud bursting
- C. Universal access
- D. Network pooling

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Rapid elasticity is a cloud computing term for scalable provisioning, or the ability to provide scalable services. Experts point to this kind of scalable model as one of five fundamental aspects of cloud computing.

Rapid elasticity allows users to automatically request additional space in the cloud or other types of services.

References: <https://www.techopedia.com/definition/29526/rapid-elasticity>

QUESTION 63

A company would like to move an application to the cloud which resides on a single physical server in their datacenter. The server has two drives, one of which hosts operating system, and the other hosts the application data. The operating system has been showing errors recently and the application data was corrupted last Friday at 4:00PM. Data is backed up every day at 1:00PM. Which of the following would be the BEST option for migrating this application to the cloud?

- A. Setup a server in the cloud, install an operating system, install the application and copy the data to the cloud server from last Friday's backup.
- B. Setup a server in the cloud, install an operating system, install and configure the application and copy the data to the cloud server from last Thursday's backup.
- C. Clone or P2V the server with both drivers to the cloud platform.
- D. Clone or P2V the server with the application to the cloud platform and copy the operating system to the cloud server.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Use the latest backup of the application data.

QUESTION 64

An existing capability is being migrated into the cloud. Capacity management issues have been noticed in the past and an exercise is being performed to calculate current and future volumes. In which of the following lifecycle phases is this likely to be performed?

- A. Operation
- B. Design
- C. Transition
- D. Strategy

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 65

An entrepreneur has decided to open an e-commerce site to complement their retail store. After researching their options, they decide that a PaaS solution will be sufficient. To reduce upfront cost, the entrepreneur intends to build the site themselves. Which of the following skill-tests will be needed?



- A. Firewall Administration
- B. Web-Server Administration
- C. Security standard development
- D. Application development

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Platform as a service (PaaS) is a complete development and deployment environment in the cloud, with resources that enable you to deliver everything from simple cloud-based apps to sophisticated, cloud-enabled enterprise applications.

PaaS allows you to avoid the expense and complexity of buying and managing software licenses, the underlying application infrastructure and middleware or the development tools and other resources. You manage the applications and services you develop, and the cloud service provider typically manages everything else.

References: <https://azure.microsoft.com/en-us/overview/what-is-paas/>

QUESTION 66

An organization is planning to host a number of its critical applications in the cloud. Which of the following is the Best way to gain a broad assurance of the cloud provider's security posture?

- A. A review that includes interviewing key security stakeholders and identifying the key controls that they operate.
- B. A review that includes security policies, evidence of the controls, physical site assessments and vulnerability scanning.
- C. A review that includes the right to audit on a yearly basis and review of the security clauses in the contract.
- D. A review that includes security applications, external audits, intrusion detection and firewall policy reviews.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 67

One major impact that cloud computing has had on the application development process is the need for greater:

- A. security
- B. speed
- C. isolation
- D. standardization

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 68

Why is it important to know the physical location for a governmental cloud based storage solution?

- A. Data stored in other countries could be accessed by the local government.
- B. Data stored in other countries could slow down application response.
- C. Data stored in other countries could impact access latency.
- D. Data stored in other countries could reduce revenue for the originating country.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

With Azure Government all data, applications, and hardware reside in the continental United States.

References: <https://azure.microsoft.com/en-us/overview/clouds/government/>

QUESTION 69

Which of the following should be measured with a direct cost chargeback method?

- A. Power and cooling consumed
- B. CPU cycles used
- C. Technical staff
- D. Square footage cost of the facility

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

CHARGEBACK METHODS

A range of approaches have been developed for implementing chargeback in an organization, as summarized in the figure below. The degree of complexity, degree of difficulty, and cost to implement decreases from the top of the chart [service-based pricing (SBP)], to the bottom [high-level allocation (HLA)]. HLA is the simplest method; it uses a straight division of IT costs based on a generic metric such as headcount. Slightly more effort to implement is low-level allocation (LLA), which bases consumer costs on something more related to IT activity such as the number of users or servers. Direct cost (DC) more closely resembles a time and materials charge but is often tied to headcount as well.

References: <https://journal.uptimeinstitute.com/it-chargeback-drives-efficiency/>

QUESTION 70

An organization is moving web server clusters to a public IaaS cloud while keeping database servers in the company owned datacenter. The organization will continue utilizing the internal service desk to manage the application. Which of the following ITIL processes will plan the move?

- A. Release Management
- B. Incident Management
- C. Problem Management
- D. Change Management

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Change Management is an IT service management discipline. The objective of change management in this context is to ensure that standardized methods and procedures are used for efficient and prompt handling of all changes to control IT infrastructure, in order to minimize the number and impact of any related incidents upon service.

References: [https://en.wikipedia.org/wiki/Change_management_\(ITSM\)](https://en.wikipedia.org/wiki/Change_management_(ITSM))

QUESTION 71

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are administrating an Azure environment for your company. You plan to deploy virtual machines (VMs) for a mobile application. You have the following requirements:

- Ensure that all VMs use the **Standard D3** size.
- Ensure that at least two of the four servers must be available at all times.
- Ensure that users of the application do not experience downtime or loss of connection.

You need to configure four VMs for application development.

Solution: Create a Virtual Machine Scale Set (VMSS) that has an instance count of 4.

Does the solution meet the goal?

- A. Yes

B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 72

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are administrating an Azure environment for your company. You plan to deploy virtual machines (VMs) for a mobile application. You have the following requirements:

- Ensure that all VMs use the **Standard D3** size.
- Ensure that at least two of the four servers must be available at all times.
- Ensure that users of the application do not experience downtime or loss of connection.

You need to configure four VMs for application development.

Solution: You create an availability set that has two fault domains and two update domains by using the Azure portal. You create four virtual machines and assign the new availability set to each VM.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 73

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are administrating an Azure environment for your company. You plan to deploy virtual machines (VMs) for a mobile application. You have the following requirements:

- Ensure that all VMs use the **Standard D3** size.
- Ensure that at least two of the four servers must be available at all times.
- Ensure that users of the application do not experience downtime or loss of connection.

You need to configure four VMs for application development.

Solution: Create two resource groups by using the Azure portal. Create four VMs. Assign two VMs to the first resource group and two to the second group.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 74

You are building an ASP.NET Azure Web App that is built from source code on GitHub. Automatic deployment is used for integration testing. The web.config file has settings that are updated during development deployments by using a TransformXml MSBuild task.

The settings in the web.config must be set to specific values during integration testing.

You need to ensure that the web.config is updated when the Web App is deployed to Azure.

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

- A. Add the integration settings and values to the ServiceDefinition.csdef and ServiceConfiguration.csfg files.
- B. In Azure, create a new deployment slot namedIntegration.
- C. Create an XML Document Transform (XDT) file namedweb.Integration.configthat converts the values to the integration test values.
- D. In Azure, add a tag with the keyEnvironmentand the valueIntegration.

Correct Answer: BC

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 75

DRAG DROP

You are developing an ASP.NET Software-as-a-Service (SaaS) application that stores database credentials in Azure Key Vault.

When the application encounters an error, information about the error is shown to the end user for support purposes. Information shown includes:

- stack trace
- application version and release date
- configuration values

Database credentials must not be accessible to end users.

You need to ensure that end users cannot view database credentials.

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

Create a ClientId and Client Secret using the Azure portal

Create a new Azure Web App with the ClientId and Client Secret values in application settings and deploy

Configure Azure Key Vault to use HSM

Create an Azure Web App with the certificate in custom domains and Secure Socket Layer (SSL) and deploy

Create an X509 certificate using by using makecert

Create a new Azure Active Directory (Azure AD) application with the certificate

Create an Azure Web App with the certificate pfx in application settings and deploy

Answer Area



Correct Answer:

Actions

-
-
-
- Create an Azure Web App with the certificate in custom domains and Secure Socket Layer (SSL) and deploy
- Create an X509 certificate using by using makecert
- Create a new Azure Active Directory (Azure AD) application with the certificate
- Create an Azure Web App with the certificate pfx in application settings and deploy

Answer Area

- Configure Azure Key Vault to use HSM
- Create a ClientId and Client Secret using the Azure portal
- ◀ Create a new Azure Web App with the ClientId and Client Secret values in application settings and deploy ▶
- ▶

Section: [none]

Explanation

Explanation/Reference:

QUESTION 76

You have an application that stores data in Azure Cosmos DB. You have the following class:

```
public static class Repository<T> where T : class
{
    public static async Task<T> SaveItem<T>(string id, T item)
    {
        var documentUri = UriFactory.CreateDocumentUri(DatabaseId, CollectionId, id);
        await client.ReplaceDocumentAsync(documentUri, item);
        var document = await client.ReadDocumentAsync(documentUri);

        return (T) (dynamic) document;
    }

    private static readonly string DatabaseId = ConfigurationManager.AppSettings["database"];
    private static readonly string CollectionId = ConfigurationManager.AppSettings["collection"];
    private static DocumentClient client;
}
```

You need to determine which consistency level to apply to this code.

Which three consistency levels can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Bounded Staleness
- B. Session
- C. Consistent Prefix
- D. Strong
- E. Eventual

Correct Answer: BCE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 77

You are creating virtual machines (VMs) that are hosted on Azure.

You must be able to change the Remote Desktop access settings for the VMs. You must also be able to change the password for the built-in administrator account on all VMs. You identify the VMAccess VM extensions that have the required capabilities.

You need to enable the VMAccess VM extensions.

Which approach should you use?

- A. Use Azure PowerShell cmdlets to change the name of the availability set to the same name for all of the VMs.
- B. Use the Azure portal to restart each VM.
- C. In the Azure Portal, create an image from a virtual hard disk (VHD) for each VM by running Sysprep.
- D. For each VM, use Azure PowerShell cmdlets to enable the VM Agent and the VMAccess VM extensions.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 78

You connect to an existing service over the network by using HTTP. The service listens on HTTP port 80. You plan to create a test environment for this existing service by using an Azure virtual machine (VM) that runs Windows Server.

The service must be accessible from the public Internet over HTTP port 8080.

You need to configure the test environment.

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

- A. Configure a Network Security Group to route traffic from port 8080 to port 80
- B. Configure a Network Security Group to route traffic from port 80 to port 8080.
- C. Ensure that the public IP address is configured as a static IP address.
- D. Configure the Windows Server firewall to allow incoming and outgoing traffic on port 8080.
- E. Configure the Windows Server firewall to allow incoming and outgoing traffic on port 80.

Correct Answer: AE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 79

HOTSPOT

You administer an Azure environment that contains multiple virtual machines (VMs).

You need to view and retrieve diagnostic logs for all VMs.

Which storage type should you use for each data source? To answer, select the appropriate options in the answer area.

Hot Area:

Answer Area

Data source	Storage type
Azure logs	<div><div></div><div>Blob</div><div>Table</div><div>Queue</div></div>
IIS 7.0 logs	<div><div></div><div>Blob</div><div>Table</div><div>Queue</div></div>
Windows Event logs	<div><div></div><div>Blob</div><div>Table</div><div>Queue</div></div>

Correct Answer:

Answer Area

Data source	Storage type
Azure logs	<div><div></div><div>Blob</div><div>Table</div><div>Queue</div></div>
IIS 7.0 logs	<div><div></div><div>Blob</div><div>Table</div><div>Queue</div></div>
Windows Event logs	<div><div></div><div>Blob</div><div>Table</div><div>Queue</div></div>

Section: [none]

Explanation

Explanation/Reference:

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/cloud-services/cloud-services-dotnet-diagnostics-storage>

QUESTION 80

You have an existing virtual network with a gateway that is deployed to Azure. You delete all objects that are deployed to the virtual network.

You use the Azure portal to delete the virtual network, but the deletion fails.

You need to determine the cause of the error.

What should you do first?

- A. Delete any local network settings and DNS servers.
- B. Save all settings.
- C. Delete all point-to-site virtual network connections.
- D. Delete the virtual network gateway.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 81

You have an on-premises Windows Identity Foundation (WIF) application. A section of the application uses resources that are hosted in Azure. The application uses Azure Active Directory (Azure AD) to control access to the section of the application that accesses Azure resources. You synchronize all user principals to Azure Active Directory.

The application has the following requirements:

- Use Windows integrated credentials for single sign-on (SSO).
- Use Azure Active Directory as an identity provider.

You need to create an endpoint to use for web sign-in to the secured section of the application.

Which endpoint should you use?

- A. SAML-P
- B. OAuth
- C. Azure AD Graph API
- D. WS-Federation

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 82

DRAG DROP

You need to add code at line CC63 to ensure that the Interaction Agent is invoked.

How should you complete the code? To answer, drag the appropriate code segments to the correct locations. Each code segment 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 point.

Select and Place:

Code fragments

ActorId.CreateRandom()

new ActorId(comment.UserId)

new ActorId(comment.Id)

ActorProxy

ServiceProxy

ActorServiceProxy

Answer area

```
var actorId = Code fragment ;  
var actor = Code fragment  
                .Create<ICommentAgent>(actorId, "...");  
await actor.ModifyCommentText(comment.Id, comment.Body,  
                                comment.Title);
```

Correct Answer:

Code fragments

`new ActorId(comment.UserId)`

`new ActorId(comment.Id)`

`ServiceProxy`

`ActorServiceProxy`

Answer area

```
var actorId = ActorId.CreateRandom();  
  
var actor = ActorProxy  
    .Create<ICommentAgent>(actorId, "...");  
  
await actor.ModifyCommentText(comment.Id, comment.Body,  
    comment.Title);
```

Section: [none]

Explanation

Explanation/Reference:

QUESTION 83

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 deploy a Virtual Machine Scale Set (VMSS) named CorpWebVMSS to Azure by using Azure PowerShell and set the instance count to 1. The VMSS includes a storage account, load balancer, public IP address, and six Standard_A1 Windows virtual machines (VMs) that run Internet Information Services (IIS). All components are deployed to a resource group named CorpWebRG.

You must increase the instance count to support the increased load on IIS.

You need to manually scale out the number of VMs in the scale set to 5.

Solution: You run the following command by using the Azure Command-Line Interface (CLI):

```
azure vmss scale -g CorpWebRG -n CorpWebVMSS -c 5
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 84

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 deploy a Virtual Machine Scale Set (VMSS) named CorpWebVMSS to Azure by using Azure PowerShell and set the instance count to 1. The VMSS includes a storage account, load balancer, public IP address, and six Standard_A1 Windows virtual machines (VMs) that run Internet Information Services (IIS). All components are deployed to a resource group named CorpWebRG.

You must increase the instance count to support the increased load on IIS.

You need to manually scale out the number of VMs in the scale set to 5.

Solution: You deploy the following JSON template by using Azure PowerShell:

```
{
  "$schema": "http://schema.management.azure.com/schemas/2015-01-01-preview/de-
ploymentTemplate.json",
  "contentVersion": "1.0.0.0",
  "resources": [
    {
      "type": "Microsoft.Compute/virtualMachineScaleSets",
      "apiVersion": "2016-03-30",
      "name": "CorpWebVMSS",
      "location": "[resourceGroup().location]",
      "sku": {
        "name": "Standard_A1",
        "tier": "Standard",
        "capacity": "5"
      }
    }
  ]
}
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-autoscale-overview>

QUESTION 85

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 deploy a Virtual Machine Scale Set (VMSS) named CorpWebVMSS to Azure by using Azure PowerShell and set the instance count to 1. The VMSS includes a storage account, load balancer, public IP address. and six Standard_A1 Windows virtual machines (VMs) that run Internet Information Services (IIS). All components are deployed to a resource group named CorpWebRG.

You must increase the instance count to support the increased load on IIS.

You need to manually scale out the number of VMs in the scale set to 5.

Solution: You run the following Azure PowerShell commands:

```
$vmss = Get-AzureRmVmss -ResourceGroupName CorpWebRG -VMScalesSetName CorpWebVMSS  
$vmss.Sku.Capacity = 5  
Update-AzureRmVmss -ResourceGroupName CorpWebRG -Name CorpWebVMSS -VirtualMachineScaleSet $vmss
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 86

You manage API management policies in Azure.

You attempt to add a policy that is marked as unavailable.

You need to ensure that you can add the desired policy.

What should you do?

- A. Modify the API Management policy definition.
- B. Enable custom caching for the API Management service.
- C. Modify the scope of the API policy.

D. Integrate the API Management service with the Azure Event Hub service.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 87

You manage an on-premises server that runs Windows Server 2016. The server has a disk that contains 4 terabytes (TB) of data and thousands of files. None of the individual files are larger than 1 TB. You plan to create a virtual machine (VM) in Azure to process the workload currently handled by the on-premises server.

You need to create a storage location for the data.

What should you do?



<https://www.gratisexam.com/>

- A. Create premium storage account. Use a D-series VM.
- B. Configure a StorSimple virtual array. Configure the VM to use the array with the SMB protocol.
- C. Add a new table storage account. Update the VM workload to use the table storage.
- D. Add a singlefile share to the VM. In the VM operating system, assign a drive letter.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 88

You administer an Azure-based solution that performs image processing. You have four Standard D3 Azure Resource Manager (ARM) virtual machines (VMs). All VMs are deployed in a Virtual Machine Scale Set (VMSS).

The servers must scale up or down as the workload increases or decreases.

You need to configure auto-scaling to scale the VMSS when the server workload is above 95 percent or below 5 percent.

What should you do?

- A. Navigate to the VM's Size panel and increase the instance count.
- B. Navigate to the VMSS Metric panel and add a new alert for the CPU Percentage Metric. Configure the alert to notify Via email.
- C. Navigate to the VM's Metric panel and enable diagnostics for basic metrics.
- D. Navigate to the VMSS Metric panel and add a new alert for the CPU Percentage Metric. Configure the alert to notify via webhook.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-vertical-scale-reprovision>

QUESTION 89

HOTSPOT

You deploy an application that uses a secure data storage solution to Azure. You use Redis Cache and select the Premium tier.

You have the following requirements:

- Create point-in-time snapshots of the dataset at specific intervals.
- Limit specific clients from using the cache.
- Use primary/replica cache pairs.

You need to configure the environment.

Which feature should you implement for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer area

Requirement	Feature
Create point-in-time snapshots of the dataset at specific intervals.	<div><div></div><div>Redis cluster</div><div>Redis persistence</div><div>Azure Virtual Network</div></div>
Limit specific clients from using the cache.	<div><div></div><div>Redis cluster</div><div>Redis persistence</div><div>Azure Virtual Network</div></div>
Use primary/replica cache pairs.	<div><div></div><div>Redis cluster</div><div>Redis persistence</div><div>Azure Virtual Network</div></div>

Correct Answer:

Answer area

Requirement	Feature
Create point-in-time snapshots of the dataset at specific intervals.	<div><div></div><div>Redis cluster</div><div>Redis persistence</div><div>Azure Virtual Network</div></div>
Limit specific clients from using the cache.	<div><div></div><div>Redis cluster</div><div>Redis persistence</div><div>Azure Virtual Network</div></div>
Use primary/replica cache pairs.	<div><div></div><div>Redis cluster</div><div>Redis persistence</div><div>Azure Virtual Network</div></div>

Section: [none]

Explanation

Explanation/Reference:

QUESTION 90

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing a new Azure Logic App. The Logic App requires a custom action to evaluate data from an internal, proprietary system. You create a custom ASP.NET Web API to retrieve data from the system and update the Logic App to use the API.

The Logic App generates a timeout error when it requests data from the API.

You need to eliminate the timeout error and allow the Logic App to retrieve data by using the API.

What should you do?

- A. Update the API to immediately return an HTTP '102 PROCESSING' response when a request is received and an HTTP '205 RESET CONTENT' response when the data is returned from the system.
- B. Update the Logic App to use a new HTTPWebhook trigger to call out to the API's newly-created subscribe and unsubscribe methods.
- C. Update the API to immediately return an HTTP '202 ACCEPTED' response when a request is received and an '200 OK' response when the data is returned from the system.
- D. Update the Logic App adding a wait action to include the interval object's unit and count properties set to valid values.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 91

You develop an Azure App Service Mobile App.

The Azure App Service must use Twitter as an authentication provider. You start by registering your application with Twitter.

You need to update your app's authentication and authorization in the Azure Portal.

Which two values should you provide? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. API Key
- B. Azure Active Directory (Azure AD) Bearer Token
- C. JSON Web Token (JWT)
- D. API Secret
- E. Mobile App gateway URL

Correct Answer: AD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 92

You are developing an Azure-hosted application that processes request messages for multiple office locations. You create an Azure Service Bus topic named Requests. The topic has a maximum size of 5 gigabytes (GB) and a default message time to live (TTL) of 5 minutes. You also create subscriptions named PriorityRequest and StandardRequest and include appropriate logic to route the messages.

Users report that the application has not processed messages from PriorityRequest in several days.

You need to retrieve the number of messages in the PriorityRequest subscription.

Which metric Should you use?

- A. Subscription Length
- B. Subscription Incoming Requests
- C. Topic Incoming Messages
- D. Topic Size

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 93

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an enterprise application that will be used only by the employees of a company. The application is not Internet-facing. You deploy instances of the application to Azure datacenters on two continents.

You must implement a load balancing solution that meets the following requirements:

- Provide network-level distribution of traffic across all instances of the application.
- Support HTTP and HTTPS protocols.
- Manage all inbound and outbound connections.

Any back-end virtual machine (VM) must be able to service requests from the same user or client session.

Solution: You implement Traffic Manager.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

A Traffic Manager works at the DNS level. It uses DNS responses to direct end-user traffic to globally distributed endpoints. Clients then connect to those endpoints directly.

An application manager, which works at the Application level (Layer 7), is also required.

References: <https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-introduction>

QUESTION 94

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an enterprise application that will be used only by the employees of a company. The application is not Internet-facing. You deploy instances of the application to Azure datacenters on two continents.

You must implement a load balancing solution that meets the following requirements:

- Provide network-level distribution of traffic across all instances of the application.
- Support HTTP and HTTPS protocols.
- Manage all inbound and outbound connections.

Any back-end virtual machine (VM) must be able to service requests from the same user or client session.

Solution: You implement Traffic Manager and Application Gateway.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Application Gateway works at the application layer (Layer 7 in the OSI network reference stack). It acts as a reverse-proxy service, terminating the client connection and forwarding requests to back-end endpoints. It supports the HTTP, HTTPS, and WebSockets protocols.

Application Gateway is useful for applications that require requests from the same user/client session to reach the same back-end virtual machine. Examples of these applications would be shopping cart applications and web mail servers.

Traffic Manager works at the DNS level. It uses DNS responses to direct end-user traffic to globally distributed endpoints. Clients then connect to those endpoints directly.

Microsoft Azure Traffic Manager allows you to control the distribution of user traffic for service endpoints in different datacenters

References: <https://docs.microsoft.com/en-us/azure/traffic-manager/traffic-manager-overview>

QUESTION 95

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an enterprise application that will be used only by the employees of a company. The application is not Internet-facing. You deploy instances of the application to Azure datacenters on two continents.

You must implement a load balancing solution that meets the following requirements:

- Provide network-level distribution of traffic across all instances of the application.
- Support HTTP and HTTPS protocols.
- Manage all inbound and outbound connections.

Any back-end virtual machine (VM) must be able to service requests from the same user or client session.

Solution: You implement Application Gateway.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Traffic Manager is also needed. It allows you to control the distribution of user traffic for service endpoints in different datacenters

References: <https://docs.microsoft.com/en-us/azure/traffic-manager/traffic-manager-overview>

QUESTION 96

You administer an Azure environment that includes six Azure Resource Manager (ARM) virtual machines (VMs) that support development. The development team uses Azure SQL databases and Azure Queues for application storage. All Azure resources are grouped within a single subscription and resource group. You need to reduce the recurring monthly Azure costs without degrading server performance. You must minimize the administrative effort involved.

What should you do?

- A. Configure an auto-shutdown schedule for each VM by using the Azure Portal.
- B. Update the development environment to use Azure Table storage.
- C. Create an Azure Automation runbook that compresses unused virtual hard disk (VHD) files daily.
- D. Create an Azure PowerShell script that backs up and deprovisions all Azure SQL databases daily.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Table storage is cheaper than SQL Databases.

It's tempting to go with answer A. However, this would degrade server performance (although shutting the servers down when not in use would save money).

There is another version of this question that does not include shutting down the servers as an answer option.

QUESTION 97

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a web app that is deployed to Azure.

You need to download a compressed collection of the diagnostic logs.

What should you use?

- A. Azure PowerShell
- B. File Transfer Protocol (FTP)
- C. Application Insights
- D. Microsoft Visual Studio

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Diagnostic information stored to the web app file system can be accessed directly using FTP. It can also be downloaded as a Zip archive using Azure PowerShell or the Azure Command-Line Interface.

References: <https://docs.microsoft.com/en-us/azure/app-service/web-sites-enable-diagnostic-log#download>

QUESTION 98

You are planning to move streaming media content to Windows Azure Storage.

You need to recommend an approach for providing worldwide users the fastest possible access to the content.

Which two actions should you recommend? (Choose two.)

- A. Use a Shared Access Signature.
- B. Use Windows Azure page blob storage.
- C. Use Windows Azure block blob storage.
- D. Use the Windows Azure Content Delivery Network (CDN).

Correct Answer: CD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

You can store text and binary data in either of two types of blobs: Block blobs, which are optimized for streaming. Page blobs, which are optimized for random read/write operations and which provide the ability to write to a range of bytes in a blob. Windows Azure provides the Windows Azure Content Delivery Network (CDN) to deliver Windows Azure Blob content. Windows Azure CDN offers developers a global solution for delivering high-bandwidth content. The benefit of using a CDN is better performance and user experience for users who are farther from the source of the content stored in the Windows Azure blob storage.

QUESTION 99

You are designing a plan for migrating Virtual Hard Disks (VHDs) and video files to Windows Azure Storage.

The VHDs must be optimized for random read/write operation.

The video files must be optimized for sequential access.

You need to recommend storage types for storing the VHDs and video files.

Which two storage types should you recommend? (Each correct answer presents part of the solution. Choose two.)

- A. Store VHDs in Windows Azure page blob storage.
- B. Store VHDs in Windows Azure block blob storage.
- C. Store video files in Windows Azure page blob storage.
- D. Store video files in Windows Azure block blob storage.

Correct Answer: AD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

You can store text and binary data in either of two types of blobs: Block blobs, which are optimized for streaming. Page blobs, which are optimized for random read/write operations and which provide the ability to write to a range of bytes in a blob. After you create or change the server image, you are ready to upload the .vhd file that contains the image data to Windows Azure. There are two opportunities for uploading VHDs to Windows Azure. When you initially create a VM role, you upload a base VHD to Windows Azure, which is used as a template to create VM role instances.

QUESTION 100

You are designing a Windows Azure application that will use Windows Azure Table storage.

The application will allow teams of users to collaborate on projects.

Each user is a member of only one team. You have the following requirements:

- Ensure that each user can efficiently query records related to his or her team's projects.
- Minimize data access latency.

You need to recommend an approach for partitioning table storage entities.

What should you recommend?

- A. Partition by user.

- B. Partition by team.
- C. Partition by project.
- D. Partition by the current date.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 101

You are designing an upgrade strategy for a Windows Azure application that includes one web role with one instance.

You have the following requirements:

- Test the application on the Windows Azure platform.
- Ensure that application upgrades can be rolled back.
- Ensure that upgrade and rollback processes do not cause downtime.

You need to recommend an approach for upgrading the application.

What should you recommend?

- A. Deploy to the Production slot. Test the application, and then perform a VIP swap.
- B. Deploy to the Staging slot. Test the application, and then perform a VIP swap.
- C. Deploy to the Staging slot. Test the application, and then perform a manual in-place upgrade to the Production slot.
- D. Deploy to the Staging slot. Test the application, and then perform an automatic in-place upgrade to the Production slot.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Run Set-AzureDeploymentSlot from the service directory to set the deployment environment for the current service to either Staging or Production.

This updates the DeploymentSettings.json file for the service.

A hosted service is a service that runs your code in the Windows Azure environment.

It has two separate deployment slots: staging and production. The staging deployment slot allows you to test your service in the Windows Azure environment before you deploy it to production. You can upgrade your service by deploying a new package to the staging environment and then swapping the staging and production deployments. This type of upgrade is called a Virtual IP or VIP swap, as it swaps the addresses of the two deployments.

Both deployments remain online during the swap process. You can swap VIPs using the Windows Azure Platform Management Portal, or by using the Service Management API. If you are upgrading your service with a new service definition file, you must swap VIPs; you cannot perform an in-place upgrade. However, you can swap VIPs only if the number of endpoints specified by the service definition is identical for both deployments. For example, if you add an HTTPS endpoint to a web role that previously exposed only an HTTP endpoint, you cannot upgrade your service using a VIP swap; you'll need to delete your production deployment and redeploy instead

QUESTION 102

You are migrating a solution to Windows Azure.

The solution includes a web application and a business logic layer.

The web application runs on three dual-core servers.

The business logic layer runs on two quad-core servers.

The Windows Azure application must match or exceed the current hardware specifications.

You need to recommend role instance sizes that minimize cost.

What should you recommend?

- A. Small for the web application and Medium for the business logic layer
- B. Small for the web application and Large for the business logic layer
- C. Medium for the web application and Large for the business logic layer
- D. Large for the web application and Extra Large for the business logic layer

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Virtual Machine Size CPU Cores Memory Disk Space for LSR in Web and Worker Roles Disk Space for LSR in a VM Role Allocated Bandwidth (Mbps)

Extra Small Shared 768 MB 19,480 MB (6,144 MB is reserved for system files) 20 GB 5

Small 1 1.75 GB 229,400 MB (6,144 MB is reserved for system files) 165 GB 100

Medium 2 3.5 GB 500,760 MB (6,144 MB is reserved for system files) 340 GB 200

Large 4 7 GB 1,023,000 MB (6,144 MB is reserved for system files) 850 GB 400

Extra Large 8 14 GB 2,087,960 MB (6,144 MB is reserved for system files) 1890 GB 800

Pricing and Metering for Compute:

Each compute instance is a virtual server.

There are 5 compute virtual server sizes you can choose from.
The table below summarizes the resources provided by each Compute instance size.
Virtual Machine Size CPU Cores Memory Cost Per Hour
Extra Small Shared 768 MB \$0.02
Small 1 1.75 GB \$0.12
Medium 2 3.5 GB \$0.24
Large 4 7 GB \$0.48
Extra Large 8 14 GB \$0.96

References: <http://msdn.microsoft.com/en-us/library/windowsazure/ee814754.aspx>

QUESTION 103

You are designing an automated deployment process for a Windows Azure application.

The process must deploy the application to Windows Azure without any user interaction.

You need to recommend a deployment strategy.

What should you recommend?

- A. Use the Service Management API to deploy the application package.
- B. Use the cspack and csrun command-line utilities and pass the cloud project as an argument.
- C. Publish the cloud project to a local directory and upload the application package to Windows Azure Blob storage.
- D. Publish the cloud project to a local directory and use the Windows Azure Developer Portal to upload the application.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 104

A Windows Azure application is running in the development fabric.

You need to recommend an approach for deploying the application to Windows Azure.

What should you recommend?

- A. Use XCopy deployment.

- B. Use the Windows Azure AppFabric.
- C. Use the Windows Azure Storage Services REST API.
- D. Use the Windows Azure Tools for Microsoft Visual Studio 2010.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 105

You are planning an upgrade strategy for a Windows Azure application.

You need to identify changes that will require application downtime.

Which change will always require downtime?

- A. Changing the virtual machine size
- B. Adding an HTTPS endpoint to a web role
- C. Changing the value of a configuration setting
- D. Upgrading the hosted service by deploying a new package

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 106

You are designing a Windows Azure application.

The application contains one web role and three worker roles.

You need to recommend an approach for updating only one role without interrupting the other roles.

What should you recommend?

- A. Perform a VIP swap.

- B. Perform an in-place upgrade.
- C. Delete the current deployment and then redeploy the application.
- D. Copy the cloud package to blob storage and then restart the service.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

QUESTION 107

You are designing a Windows Azure application that will generate events for multiple clients.

Client web services might be behind NAT gateways.

You need to recommend an approach that will allow you to broadcast the events to clients.

What should you recommend?

- A. Use ADO.NET Data Services and provide a shared key to clients.
- B. Use Windows Azure Queues and provide a shared key to clients.
- C. Use Windows Azure Table storage and provide a shared key to clients.
- D. Use the Windows Azure AppFabric Service Bus and provide a shared secret to clients.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The Windows Azure Service Bus provides a hosted, secure, and widely available infrastructure for widespread communication, large-scale event distribution, naming, and service publishing. The Service Bus provides connectivity options for Windows Communication Foundation (WCF) and other service endpoints including REST endpoints -- that would otherwise be difficult or impossible to reach. Endpoints can be located behind network address translation (NAT) boundaries, or bound to frequently-changing, dynamically-assigned IP addresses, or both. The Service Bus provides both "relayed" and "brokered" messaging capabilities. In the relayed messaging pattern, the relay service supports direct one-way messaging, request/response messaging, and peer-to-peer messaging. Brokered messaging provides durable, asynchronous messaging components such as Queues, Topics, and Subscriptions, with features that support publish-subscribe and temporal decoupling:

- senders and receivers do not have to be online at the same time;
- the messaging infrastructure reliably stores messages until the receiving party is ready to receive them.

QUESTION 108

You are designing a Windows Azure application.

The application includes a web role and a worker role that communicate by using a Windows Azure Queue. The worker role processes each message within 10 seconds of retrieving it from the queue. The worker role must process each message exactly one time.

If a process does not complete, the worker role must reprocess the message.

You need to recommend an approach for the worker role to manage messages in the queue.

What should you recommend?

- A. Process the message and then delete it from the queue.
- B. Delete the message from the queue when retrieving the message.
- C. Set the visibility timeout of the message to 1 when retrieving the message.
- D. Process the message and then set the visibility timeout of the message to the maximum value.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 109

You are designing a Windows Azure application.

The application includes two web roles and three instances of a worker role. The web roles will send requests to the worker role through one or more Windows Azure Queues.

You have the following requirements:

- Ensure that each request is processed exactly one time.
- Minimize the idle time of each worker role instance.
- Maximize the reliability of request processing.

You need to recommend a queue design for sending requests to the worker role.

What should you recommend?

- A. Create a single queue. Send requests on the single queue.

- B. Create a queue for each web role. Send requests on all queues at the same time.
- C. Create a queue for each worker role instance. Send requests on each worker queue in a round robin.
- D. Create a queue for each combination of web roles and worker role instances. Send requests to all worker role instances based on the sending web role.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 110

You are designing a Windows Azure application that will process images. The maximum size of an image is 10 MB.

The application includes a web role that allows users to upload images and a worker role with multiple instances that processes the images.

The web role communicates with the worker role by using a Windows Azure Queue.

You need to recommend an approach for storing images that minimizes storage transactions.

What should you recommend?

- A. Store images in the queue.
- B. Store images in Windows Azure Blob storage. Store references to the images in the queue.
- C. Store images in local storage on the web role instance. Store references to the images in the queue.
- D. Store images in Windows Azure Drives attached to the worker role instances. Store references to the images in the queue.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

You can store text and binary data in either of two types of blobs: Block blobs, which are optimized for streaming. Page blobs, which are optimized for random read/write operations and which provide the ability to write to a range of bytes in a blob.

QUESTION 111

You are designing a Windows Azure application.

The application will store data in Windows Azure Blob storage.

Many of the application services will be interdependent.

You need to recommend an approach for optimizing the performance of the application.

What should you recommend?

- A. Create one affinity group. Associate only the storage services with the affinity group.
- B. Create one affinity group. Associate only the compute services with the affinity group.
- C. Create one affinity group. Associate the compute services and storage services with the affinity group.
- D. Create two affinity groups. Associate the compute services with one group and the storage services with the other group.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Use the following procedures to create an affinity group, which can be used to direct Windows Azure storage accounts and hosted services to the same geographical grouping within a specified region. Each affinity group is associated with a Windows Azure subscription, and can be used by multiple storage accounts and hosted services for that subscription.

Affinity groups can be created and managed by the service administrator and co-administrators for a subscription.

QUESTION 112

Which of the following are valid differences between page blobs and block blobs? (Choose two.)

- A. Page blobs are much faster for all operations.
- B. Block blobs allow files to be uploaded and assembled later. Blocks can be resubmitted individually.
- C. Page blobs are good for all sorts of files, like video and images.
- D. Block blobs have a max size of 200 GB. Page blobs can be 1 terabyte.



<https://www.gratisexam.com/>

Correct Answer: BD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 113

Which of the following are valid options for scaling queues? (Choose three.)

- A. Distributing messages across multiple queues
- B. Automatically scaling websites based on queue metrics
- C. Automatically scaling VMs based on queue metrics
- D. Automatically scaling cloud services based on queue metrics

Correct Answer: ACD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 114

Which of the following is not true about metadata? (Choose two.)

- A. Both containers and blobs have writable system properties.
- B. Blob user-defined metadata is accessed as a key value pair.
- C. System metadata can influence how the blob is stored and accessed in Azure Storage.
- D. Only blobs have metadata; containers do not.

Correct Answer: AD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 115

Which of the following are true regarding supported operations granted with an SAS token? (Choose three.)

- A. You can grant read access to existing blobs.
- B. You can create new blob containers.
- C. You can add, update, and delete queue messages.
- D. You can add, update, and delete table entities.
- E. You can query table entities.

Correct Answer: ADE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

You can add or update but not delete queue messages.

References:

<https://blogs.msdn.microsoft.com/windowsazurestorage/2012/06/12/introducing-table-sas-shared-access-signature-queue-sas-and-update-to-blob-sas/>

QUESTION 116

Which of the following are valid options for processing queue messages? (Choose Two.)

- A. A single compute instance can process only one message at a time.
- B. A single compute instance can process up to 31 messages at a time.
- C. A single compute instance can retrieve up to 32 messages at a time.
- D. Messages can be read one at a time or in batches of up to 32 messages at a time.
- E. Messages are deleted as soon as they are read.

Correct Answer: CD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 117

Which of the following statements are true of stored access policies? (Choose two.)

- A. You can modify the start or expiration date for access.
- B. You can revoke access at any point in time.

- C. You can modify permissions to remove or add supported operations.
- D. You can add to the list of resources accessible by an SAS token.

Correct Answer: AB

Section: [none]

Explanation

Explanation/Reference:

Explanation:

You can use a stored access policy to change the start time, expiry time, or permissions for a signature, or to revoke it after it has been issued.

References: <https://docs.microsoft.com/en-us/rest/api/storageservices/establishing-a-stored-access-policy>

QUESTION 118

How should you choose a good partition key for a Table storage implementation? (Choose two.)

- A. They should always be unique, like a primary key in a SQL table.
- B. You should always use the same partition key for all records.
- C. Think about how you're likely to update the data using batch transactions.
- D. Find an even way to split them so that you have relatively even partition sizes.

Correct Answer: CD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 119

Which of the following is not a method for replicating a Table storage account?

- A. Transactional replication
- B. Zone redundant storage
- C. Read access geo-redundant storage
- D. Geo-redundant storage

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 120

Which of the following statements are correct for submitting operations in a batch? (Choose three.)

- A. All operations have to be in the same partition.
- B. Total batch size can't be greater than 4 MB.
- C. Max operation count is 100.
- D. Minimum operation count is three

Correct Answer: ABC

Section: [none]

Explanation

Explanation/Reference:

QUESTION 121

Which of the following statements are true of CORS support for storage? (Choose Two.)

- A. It is recommended you enable CORS so that browsers can access blobs.
- B. To protect CORS access to blobs from the browser, you should generate SAS tokens to secure blob requests.
- C. CORS is supported only for Blob storage.
- D. CORS is disabled by default.

Correct Answer: BD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 122

Which of the following is not a requirement for creating an online secondary for SQL Database?

- A. The secondary database must have the same name as the primary.
- B. They must be on separate servers.

- C. They both must be on the different subscription.
- D. The secondary server cannot be a lower performance tier than the primary.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

QUESTION 123

Which statement is true of Storage Analytics Metrics?

- A. Capacity metrics are recorded only for blobs.
- B. You can set hourly or by minute metrics through the management portal.
- C. By default, metrics are retained for one year.
- D. If you disable metrics, existing metrics are deleted from storage.

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 124

Which metrics should you add to monitoring that will help you select the appropriate level of SQL Database? (Choose three.)

- A. CPU Processor Count
- B. CPU Percentage
- C. Physical Data Reads Percentage
- D. Log Writes Percentage

Correct Answer: BCD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 125

Which of the following are captured by Storage Analytics Logging? (Choose two.)

- A. Successful requests for authenticated calls only
- B. Failed requests for authenticated calls only
- C. Server errors
- D. Requests using SAS URLs.

Correct Answer: CD

Section: [none]

Explanation

Explanation/Reference:

QUESTION 126

Which of the following statement are true about queuing messages?

- A. Storage queue messages have no size restrictions. The reason for using smaller messages sizes is to increase throughput to the queue.
- B. Storage queue messages are limited to 64 KB.
- C. Storage queue messages are durable.
- D. The client application should save the message identifier returned after adding a message to a queue for later use.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 127

You administer an Azure environment that includes six Azure Resource Manager (ARM) virtual machines (VMs) that support development. The development team uses Azure SQL databases and Azure Queues for application storage. All Azure resources are grouped within a single subscription and resource group.

You need to reduce the recurring monthly Azure costs without degrading server performance. You must minimize the administrative effort involved.

What should you do?

- A. Remove the development team role from the resource group daily.
- B. Create an Azure Automation runbook that cycles the VMs daily.
- C. Update the development environment to use Azure Table storage.
- D. Create an Azure PowerShell script that updates the VM size to Standard_A0 daily.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 128

You are developing a Windows console application that uses a third-party C++ library. The console application is designed to be run as an Azure WebJob that has authentication and authorization enabled.

You need to ensure that the console application can determine the current user identity.

What should you do?

- A. Perform an HTTP request to the **/.auth/me** endpoint.
- B. Call the **System.Security.Principal.WindowsIdentity.GetCurrent()** method.
- C. Read the **X-MS-CLIENT-PRINCIPAL-NAME** header.
- D. Read the identity from the **UserName** environment variable.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

QUESTION 129

You are developing an Azure Logic App that summarizes and translates Microsoft Word documents. Clients upload Word documents to an Azure File share, and all documents are processed once per day.

Each document takes more than 10 minutes to process, and each client processes thousands of documents per day.

You need to ensure that customers can process documents on demand while minimizing costs.

Which two types of Logic App triggers can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. HTTP
- B. Request
- C. HTTPWebhook
- D. ApiConnection
- E. ApiConnectionWebhook

Correct Answer: BE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 130

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 are developing an application that will run as an Azure API App. The application tracks flights between airports, including duration of flight, if the flight was on time, the capacity of the airplane, and the number of seats sold. Queries can be performed to show multiple routes, multi-leg journeys, and filtering based on the attributes of the flight.

Flight information will be used by customers to perform data mining, drive interactive display, perform airspace tracking, and other applications.

Customers require that the response time of the API be as low as possible, both for retrieving information for a single flight, and for queries across flights. To achieve the required level of performance, each API invocation must be satisfied by a single operation against the data store containing flight information.

You need to implement the data store for this application.

Solution: You use Azure Search.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 131

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 are developing an application that will run as an Azure API App. The application tracks flights between airports, including duration of flight, if the flight was on time, the capacity of the airplane, and the number of seats sold. Queries can be performed to show multiple routes, multi-leg journeys, and filtering based on the attributes of the flight.

Flight information will be used by customers to perform data mining, drive interactive display, perform airspace tracking, and other applications.

Customers require that the response time of the API be as low as possible, both for retrieving information for a single flight, and for queries across flights. To achieve the required level of performance, each API invocation must be satisfied by a single operation against the data store containing flight information.

You need to implement the data store for this application.

Solution: You use Azure Blob Storage.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 132

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 are developing an application that will run as an Azure API App. The application tracks flights between airports, including duration of flight, if the flight was on time, the capacity of the airplane, and the number of seats sold. Queries can be performed to show multiple routes, multi-leg journeys, and filtering based on the attributes of the flight.

Flight information will be used by customers to perform data mining, drive interactive display, perform airspace tracking, and other applications.

Customers require that the response time of the API be as low as possible, both for retrieving information for a single flight, and for queries across flights. To achieve the required level of performance, each API invocation must be satisfied by a single operation against the data store containing flight information.

You need to implement the data store for this application.

Solution: You use Cosmos DB.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 133

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 are a system administrator at your company. Your company recently acquired two of its competitors, as well as their IT infrastructure. The acquired companies have applications that are written in Java, .NET, Ruby, php, Node.js, and other languages. The applications run on Linux and Windows Server in Amazon Web Services, Azure, and SAP Cloud Platform.

The applications require access to the Azure Service Broker, and must be managed by the PCF Ops Manager.

You need to consolidate the applications onto a single cloud provider in Azure.

Solution: Deploy the open-source Cloud Foundry packages by setting up a BOSH director.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Use the Azure Marketplace Pivotal Cloud Foundry offer.

References: <https://azuremarketplace.microsoft.com/en-us/marketplace/apps/pivotal.pivotal-cloud-foundry>

QUESTION 134

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 are a system administrator at your company. Your company recently acquired two of its competitors, as well as their IT infrastructure. The acquired companies have applications that are written in Java, .NET, Ruby, php, Node.js, and other languages. The applications run on Linux and Windows Server in Amazon Web Services, Azure, and SAP Cloud Platform.

The applications require access to the Azure Service Broker, and must be managed by the PCF Ops Manager.

You need to consolidate the applications onto a single cloud provider in Azure.

Solution: Create a customized environment by deploying Pivotal Cloud Foundry manually.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 135

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 are a system administrator at your company. Your company recently acquired two of its competitors, as well as their IT infrastructure. The acquired companies have applications that are written in Java, .NET, Ruby, php, Node.js, and other languages. The applications run on Linux and Windows Server in Amazon Web Services, Azure, and SAP Cloud Platform.

The applications require access to the Azure Service Broker, and must be managed by the PCF Ops Manager.

You need to consolidate the applications onto a single cloud provider in Azure.

Solution: Use the Azure Cloud Shell to install the Cloud Foundry CLI and connect to the Cloud Controller.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Use Azure Marketplace Pivotal Cloud Foundry.

References: <https://azuremarketplace.microsoft.com/en-us/marketplace/apps/pivotal.pivotal-cloud-foundry>

QUESTION 136

You are administering an Azure environment for your company. You plan to deploy virtual machines (VMs) for a mobile application. You have the following requirements:

- Ensure that all VMs use the **Standard D3** size.
- Ensure that at least two of the four servers must be available at all times.
- Ensure that users of the application do not experience downtime or loss of connection.

You need to configure four VMs for application development.

Solution: Create a resource group by using the Azure portal. Create four VMs and assign all VMs to the new resource group.

Does the solution meet the goal?



<https://www.gratisexam.com/>

A. Yes

B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 137

You need to implement authentication.

What should you use?

A. Active Directory Federation Services (AD FS)

B. Azure Active Directory (Azure AD) business to-consumer (B2C) stand-alone

C. Enterprise State Roaming

D. Azure Active Directory (Azure AD) self-service signup

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 138

You need to create an Azure Function to run the Daily Sponsor Report.

<https://www.gratisexam.com/>

Which function template should you use?

- A. TimerTrigger
- B. HTTPTrigger
- C. ServiceBusQueueTrigger
- D. EventHubTrigger

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 139

You need to ensure that transcripts and notes are processed.

Which trigger or binding should you use?

- A. Storage tables
- B. Schedule
- C. No-SQL DB
- D. Blob storage

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

QUESTION 140

You are creating virtual machines (VMs) that are hosted on Azure.

You must be able to change the Remote Desktop access settings for the VMs. You must also be able to change the password for the built-in administrator account on all VMs. You identify the VMAccess VM extensions that have the required capabilities.

You need to enable the VMAccess VM extensions.

Which approach should you use?

- A. Use the Azure portal to restart each VM.
- B. Download and install the Microsoft Installer file to enable the VM Agent on each VM.
- C. Create an application that updates the configuration of each VM.
- D. Use Azure PowerShell cmdlets to change the name of the availability set to the same name for all of the VMs.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Several PowerShell commands exist for running individual extensions.

In the following example, the VM Access extension is used to reset the administrative password of a Windows VM to a temporary password. For more information on the VM Access extension, see [Reset Remote Desktop service in a Windows VM](#). Once you have run this, you should reset the password at first login:

```
$cred=Get-Credential
```

```
Set-AzureRmVMAccessExtension -ResourceGroupName "myResourceGroup" -VMName "myVM" -Name "myVMAccess" `
  -Location WestUS -UserName $cred.GetNetworkCredential().Username `
  -Password $cred.GetNetworkCredential().Password -typeHandlerVersion "2.0"
```

Incorrect Answers:

A: VM extensions can be applied to an existing VM through the Azure portal. Select the VM in the portal, choose Extensions, then select Add. Choose the extension you want from the list of available extensions and follow the instructions in the wizard.

Reference: <https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/features-windows>

QUESTION 141

You administer an Azure environment that includes six Azure Resource Manager (ARM) virtual machines (VMs) that support development. The development team uses Azure SQL databases and Azure Queues for application storage. All Azure resources are grouped within a single subscription and resource group.

You need to reduce the recurring monthly Azure costs without degrading server performance. You must minimize the administrative effort involved.

What should you do?

- A. Remove the development team role from the resource group daily.
- B. Create an Azure Automation runbook that compresses unused virtual hard disk (VHD) files daily.

- C. Create an Azure PowerShell script that backs up and deprovisions all Azure SQL databases daily.
- D. Create an Azure Automation runbook that cycles the VMs daily.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Schedule startup and shutdown of your virtual machines using Azure Automation. The processes you want to automate can be done through runbooks

Reference:

<https://azure.microsoft.com/en-us/blog/azure-automation-runbook-management/>

<https://ppolyzos.com/2017/08/15/schedule-startup-and-shutdown-of-your-virtual-machines-using-azure-automation/>

QUESTION 142

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 are a system administrator at your company. Your company recently acquired two of its competitors, as well as their IT infrastructure. The acquired companies have applications that are written in Java, .NET, Ruby, php, Node.js, and other languages. The applications run on Linux and Windows Server in Amazon Web Services, Azure, and SAP Cloud Platform.

The applications require access to the Azure Service Broker, and must be managed by the PCF Ops Manager.

You need to consolidate the applications onto a single cloud provider in Azure.

Solution: Use the Azure Marketplace Pivotal Cloud Foundry offer to create a standard environment that includes PCF Ops Manager and the Azure Service Broker.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Use a customized, not a standard, of the Azure Marketplace Pivotal Cloud Foundry.

References: <https://azuremarketplace.microsoft.com/en-us/marketplace/apps/pivotal.pivotal-cloud-foundry>

QUESTION 143

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 a virtual machine scale set (VMSS) with three virtual machines (VMs). You define rules based on performance metrics and application response. You must define the performance metrics based on the data collected.

You need to configure the autoscale rules.

Solution: Use Azure PowerShell to create rules to automatically scale out the VMSS.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

To create autoscale rules that use more detailed performance metrics, you can install and configure the Azure diagnostics extension on VM instances, or configure your application use App Insights.

Autoscale rules that use host-based metrics can be created with one of the following tools:

- Azure portal
- Azure PowerShell
- Azure CLI 2.0
- Azure template

Note: You can create autoscale rules that built-in host metrics available from your VM instances. Host metrics give you visibility into the performance of the VM instances in a scale set without the need to install or configure additional agents and data collections. Autoscale rules that use these metrics can scale out or in the number of VM instances in response to CPU usage, memory demand, or disk access.

References: <https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-autoscale-overview>

QUESTION 144

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 a virtual machine scale set (VMSS) with three virtual machines (VMs). You define rules based on performance metrics and application response. You must define the performance metrics based on the data collected.

You need to configure the autoscale rules.

Solution: Use the Visual Studio Cloud Explorer to autoscale VMSS.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

To create autoscale rules that use more detailed performance metrics, you can install and configure the Azure diagnostics extension on VM instances, or configure your application use App Insights.

Autoscale rules that use host-based metrics can be created with one of the following tools:

- Azure portal
- Azure PowerShell
- Azure CLI 2.0
- Azure template

Note: You can create autoscale rules that built-in host metrics available from your VM instances. Host metrics give you visibility into the performance of the VM instances in a scale set without the need to install or configure additional agents and data collections. Autoscale rules that use these metrics can scale out or in the number of VM instances in response to CPU usage, memory demand, or disk access.

References: <https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-autoscale-overview>

