

[New Exams!Free AI-100 PDF and VCE Offered by Braindump2go[Q45-Q55

July/2019 Braindump2go AI-100 Dumps with PDF and VCE New Updated Today! Following are some new AI-100 Exam Questions: 1.|2019 Latest AI-100 Exam Dumps (PDF & VCE) Instant

Download:<https://www.braindump2go.com/ai-100.html>2.|2019 Latest AI-100 Exam Questions & Answers Instant

Download:<https://drive.google.com/drive/folders/16bPnYGUoXhAsx9eAI8URN71n7ufNMWaM?usp=sharing>QUESTION 45You are configuring data persistence for a Microsoft Bot Framework application. The application requires a structured NoSQL cloud data store.You need to identify a storage solution for the application. The solution must minimize costs.What should you identify?A.

Azure Blob storageB. Azure Cosmos DBC. Azure HDInsightD. Azure Table storageAnswer: DExplanation:Table Storage is a NoSQL key-value store for rapid development using massive semi-structured datasetsYou can develop applications on Cosmos DB using popular NoSQL APIs. Both services have a different scenario and pricing model. While Azure Storage Tables is aimed at high capacity on a single region (optional secondary read only region but no failover), indexing by PK/RK and storage-optimized pricing; Azure Cosmos DB Tables aims for high throughput (single-digit millisecond latency), global distribution (multiple failover), SLA-backed predictive performance with automatic indexing of each attribute/property and a pricing model focused on throughput.

References:<https://db-engines.com/en/system/Microsoft+Azure+Cosmos+DB%3BMicrosoft+Azure+Table+Storage>

QUESTION 46You have an Azure Machine Learning model that is deployed to a web service. You plan to publish the web service by using the name ml.contoso.com. You need to recommend a solution to ensure that access to the web service is encrypted. Which three actions should you recommend? Each correct answer presents part of the solution.NOTE: Each correct selection is worth one point.A. Generate a shared access signature (SAS)B. Obtain an SSL certificateC. Add a deployment slotD. Update the web serviceE. Update DNSF. Create an Azure Key VaultAnswer: BDEExplanation:The process of securing a new web service or an existing one is as follows:1. Get a domain name.2. Get a digital certificate.3. Deploy or update the web service with the SSL setting enabled.4. Update your DNS to point to the web service.Note: To deploy (or re-deploy) the service with SSL enabled, set the ssl_enabled parameter to True, wherever applicable. Set the ssl_certificate parameter to the value of the certificate file and the ssl_key to the value of the key file.References:

<https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-secure-web-service>QUESTION 47You plan to design an application that will use data from Azure Data Lake and perform sentiment analysis by using Azure Machine Learning algorithms.The developers of the application use a mix of Windows- and Linux-based environments. The developerscontribute to shared GitHub repositories.You need all the developers to use the same tool to develop the application. What is the best tool to use? More than one answer choice may achieve the goal.A. Microsoft Visual Studio CodeB. Azure NotebooksC. Azure Machine Learning StudioD. Microsoft Visual StudioAnswer: CExplanation:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/machine-learning/studio/algorithm-choice.md>QUESTION

48You have several AI applications that use an Azure Kubernetes Service (AKS) cluster. The cluster supports a maximum of 32 nodes.You discover that occasionally and unpredictably, the application requires more than 32 nodes. You need to recommend a solution to handle the unpredictable application load.Which scaling method should you recommend?A. horizontal pod autoscaler B. cluster autoscalerC. manual scalingD. Azure Container InstancesAnswer: BExplanation:To keep up with application demands in Azure Kubernetes Service (AKS), you may need to adjust the number of nodes that run your workloads. The cluster autoscaler component can watch for pods in your cluster that can't be scheduled because of resource constraints. When issues are detected, the number of nodes is increased to meet the application demand. Nodes are also regularly checked for a lack of running pods, with the number of nodes then decreased as needed. This ability to automatically scale up or down the number of nodes in your AKS cluster lets you run an efficient, cost-effective cluster.References:

<https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler>QUESTION 49You deploy an infrastructure for a big data workload. You need to run Azure HDInsight and Microsoft Machine Learning Server. You plan to set the RevoScaleR compute contexts to run rx function calls in parallel. What are three compute contexts that you can use for Machine Learning Server? Each correct answer presents a complete solution.NOTE: Each correct selection is worth one point.A. SQLB. SparkC. local parallel D. HBaseE. local sequentialAnswer: ABCEExplanation:Remote computing is available for specific data sources on selected platforms. The following tables document the supported combinations.RxInSqlServer, sqlserver: Remote compute context. Target server is a single database node (SQL Server2016 R Services or SQL Server 2017 Machine Learning Services). Computation is parallel, but not distributed.RxSpark, spark: Remote compute context. Target is a Spark cluster on Hadoop. RxLocalParallel, localpar: Compute context is often used to enable controlled, distributed computationsrelying on instructions you provide rather than a built-in scheduler on Hadoop. You can use compute context for manual distributed computing.References:

<https://docs.microsoft.com/en-us/machine-learning-server/r/concept-what-is-compute-context>QUESTION 50Your company has 1,000 AI developers who are responsible for provisioning environments in Azure. You need to control the type, size, and location of the resources that the developers can provision. What should you use? A. Azure Key Vault B. Azure service principals C. Azure managed identities D. Azure Security Center E. Azure Policy Answer: B Explanation: When an application needs access to deploy or configure resources through Azure Resource Manager in Azure Stack, you create a service principal, which is a credential for your application. You can then delegate only the necessary permissions to that service principal. References:

<https://docs.microsoft.com/en-us/azure/azure-stack/azure-stack-create-service-principals>QUESTION 51You are designing an AI solution in Azure that will perform image classification. You need to identify which processing platform will provide you with the ability to update the logic over time. The solution must have the lowest latency for inferencing without having to batch. Which compute target should you identify? A. graphics processing units (GPUs) B. field-programmable gate arrays (FPGAs) C. central processing units (CPUs) D. application-specific integrated circuits (ASICs) Answer: B Explanation: FPGAs, such as those available on Azure, provide performance close to ASICs. They are also flexible and reconfigurable over time, to implement new logic. Incorrect Answers: D: ASICs are custom circuits, such as Google's TensorFlow Processor Units (TPU), provide the highest efficiency. They can't be reconfigured as your needs change. References:

<https://docs.microsoft.com/en-us/azure/machine-learning/service/concept-accelerate-with-fpgas>QUESTION 52You have a solution that runs on a five-node Azure Kubernetes Service (AKS) cluster. The cluster uses an Nseries virtual machine. An Azure Batch AI process runs once a day and rarely on demand. You need to recommend a solution to maintain the cluster configuration when the cluster is not in use. The solution must not incur any compute costs. What should you include in the recommendation? A. Downscale the cluster to one node B. Downscale the cluster to zero nodes C. Delete the cluster Answer: A Explanation: An AKS cluster has one or more nodes. References: <https://docs.microsoft.com/en-us/azure/aks/concepts-clusters-workloads>QUESTION 53

Your company has recently deployed 5,000 Internet-connected sensors for a planned AI solution. You need to recommend a computing solution to perform a real-time analysis of the data generated by the sensors. Which computing solution should you recommend? A. an Azure HDInsight Storm cluster B. Azure Notification Hubs C. an Azure HDInsight Hadoop cluster D. an Azure HDInsight R cluster Answer: C Explanation: Azure HDInsight makes it easy, fast, and cost-effective to process massive amounts of data. You can use HDInsight to process streaming data that's received in real time from a variety of devices. References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-introduction>QUESTION 54Note: 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, 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 uses an Azure Kubernetes Service (AKS) cluster. You are troubleshooting a node issue. You need to connect to an AKS node by using SSH. Solution: You add an SSH key to the node, and then you create an SSH connection. Does this meet the goal? A. Yes B. No Answer: A Explanation: By default, SSH keys are generated when you create an AKS cluster. If you did not specify your own SSH keys when you created your AKS cluster, add your public SSH keys to the AKS nodes. You also need to create an SSH connection to the AKS node. References: <https://docs.microsoft.com/en-us/azure/aks/ssh>QUESTION 55

You are developing a Computer Vision application. You plan to use a workflow that will load data from an on-premises database to Azure Blob storage, and then connect to an Azure Machine Learning service. What should you use to orchestrate the workflow? A. Azure Kubernetes Service (AKS) B. Azure Pipelines C. Azure Data Factory D. Azure Container Instances Answer: C Explanation: With Azure Data Factory you can use workflows to orchestrate data integration and data transformation processes at scale. Build data integration, and easily transform and integrate big data processing and machine learning with the visual interface. References:

<https://azure.microsoft.com/en-us/services/data-factory/>!!!RECOMMEND!!!1.2019 Latest AI-100 Exam Dumps (PDF & VCE) Instant Download: <https://www.braindump2go.com/ai-100.html>2.2019 Latest AI-100 Study Guide Video Instant Download: YouTube Video: [YouTube.com/watch?v=WUM5fSuomxQ](https://www.youtube.com/watch?v=WUM5fSuomxQ)