

[2017 New Lead2pass Latest AWS Certified Solutions Architect - Associate PDF Guarantee 100% Pass AWS Certified Solutions Architect–Associate (426-450)]

[2017 August Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed!](#) Are you struggling for the AWS Certified Solutions Architect - Associate exam? Good news, Lead2pass Amazon technical experts have collected all the questions and answers which are updated to cover the knowledge points and enhance candidates' abilities. We offer the latest AWS Certified Solutions Architect - Associate PDF and VCE dumps with new version VCE player for free download, and the new AWS Certified Solutions Architect - Associate dump ensures your AWS Certified Solutions Architect - Associate exam 100% pass. Following questions and answers are all new published by Amazon Official Exam Center: <https://www.lead2pass.com/aws-certified-solutions-architect-associate.html>

QUESTION 426 Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as _____. A. snapshots B. images C. instance backups D. mirrors

Answer: A Explanation: Amazon allows you to make backups of the data stored in your EBS volumes through snapshots that can later be used to create a new EBS volume. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

QUESTION 427 To specify a resource in a policy statement, in Amazon EC2, can you use its Amazon Resource Name (ARN)? A. Yes, you can. B. No, you can't because EC2 is not related to ARN. C. No, you can't because you can't specify a particular Amazon EC2 resource in an IAM policy. D. Yes, you can but only for the resources that are not affected by the action.

Answer: A Explanation: Some Amazon EC2 API actions allow you to include specific resources in your policy that can be created or modified by the action. To specify a resource in the statement, you need to use its Amazon Resource Name (ARN). Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-ug.pdf>

QUESTION 428 After you recommend Amazon Redshift to a client as an alternative solution to paying data warehouses to analyze his data, your client asks you to explain why you are recommending Redshift. Which of the following would be a reasonable response to his request? A. It has high performance at scale as data and query complexity grows. B. It prevents reporting and analytic processing from interfering with the performance of OLTP workloads. C. You don't have the administrative burden of running your own data warehouse and dealing with setup, durability, monitoring, scaling, and patching. D. All answers listed are a reasonable response to his question

Answer: D Explanation: Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis or any SSH-enabled host.

AWS recommends Amazon Redshift for customers who have a combination of needs, such as: High performance at scale as data and query complexity grows Desire to prevent reporting and analytic processing from interfering with the performance of OLTP workloads Large volumes of structured data to persist and query using standard SQL and existing BI tools Desire to the administrative burden of running one's own data warehouse and dealing with setup, durability, monitoring, scaling and patching

Reference: https://aws.amazon.com/running_databases/#redshift_anchor

QUESTION 429 One of the criteria for a new deployment is that the customer wants to use AWS Storage Gateway. However you are not sure whether you should use gateway-cached volumes or gateway-stored volumes or even what the differences are. Which statement below best describes those differences? A. Gateway-cached lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-stored enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. B. Gateway-cached is free whilst gateway-stored is not. C. Gateway-cached is up to 10 times faster than gateway-stored. D. Gateway-stored lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.

Answer: A Explanation: Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations: Gateway-cached volumes ?You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data. Gateway-stored volumes ?If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive off-site backups that you can recover to your local data center or Amazon EC2. For example, if

you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2. Reference: <http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateway.html> QUESTION 430 A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone? A. Always select the AZ while launching an instance B. Always select the US-East-1-a zone for HA C. Do not select the AZ; instead let AWS select the AZ D. The user can never select the availability zone while launching an instance Answer: C Explanation: When launching an instance with EC2, AWS recommends not to select the availability zone (AZ). AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html> QUESTION 431 A user is storing a large number of objects on AWS S3. The user wants to implement the search functionality among the objects. How can the user achieve this? A. Use the indexing feature of S3. B. Tag the objects with the metadata to search on that. C. Use the query functionality of S3. D. Make your own DB system which stores the S3 metadata for the search functionality. Answer: D Explanation: In Amazon Web Services, AWS S3 does not provide any query facility. To retrieve a specific object the user needs to know the exact bucket / object key. In this case it is recommended to have an own DB system which manages the S3 metadata and key mapping. Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf QUESTION 432 After setting up a Virtual Private Cloud (VPC) network, a more experienced cloud engineer suggests that to achieve low network latency and high network throughput you should look into setting up a placement group. You know nothing about this, but begin to do some research about it and are especially curious about its limitations. Which of the below statements is wrong in describing the limitations of a placement group? A. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. B. A placement group can span multiple Availability Zones. C. You can't move an existing instance into a placement group. D. A placement group can span peered VPCs Answer: B Explanation: A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. Placement groups have the following limitations: The name you specify for a placement group a name must be unique within your AWS account. A placement group can't span multiple Availability Zones. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. We recommend using the same instance type for all instances in a placement group. You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group. A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs. For more information about VPC peering connections, see VPC Peering in the Amazon VPC User Guide. You can't move an existing instance into a placement group. You can create an AMI from your existing instance, and then launch a new instance from the AMI into a placement group. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html> QUESTION 433

What is a placement group in Amazon EC2? A. It is a group of EC2 instances within a single Availability Zone. B. It the edge location of your web content. C. It is the AWS region where you run the EC2 instance of your web content. D. It is a group used to span multiple Availability Zones. Answer: A Explanation: A placement group is a logical grouping of instances within a single Availability Zone. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html> QUESTION 434 You are migrating an internal server on your DC to an EC2 instance with EBS volume. Your server disk usage is around 500GB so you just copied all your data to a 2TB disk to be used with AWS Import/Export. Where will the data be imported once it arrives at Amazon? A. to a 2TB EBS volume B. to an S3 bucket with 2 objects of 1TB C. to an 500GB EBS volume D. to an S3 bucket as a 2TB snapshot Answer: B Explanation: An import to Amazon EBS will have different results depending on whether the capacity of your storage device is less than or equal to 1 TB or greater than 1 TB. The maximum size of an Amazon EBS snapshot is 1 TB, so if the device image is larger than 1 TB, the image is chunked and stored on Amazon S3. The target location is determined based on the total capacity of the device, not the amount of data on the device. Reference:

<http://docs.aws.amazon.com/AWSImportExport/latest/DG/Concepts.html> QUESTION 435 A client needs you to import some existing infrastructure from a dedicated hosting provider to AWS to try and save on the cost of running his current website. He also needs an automated process that manages backups, software patching, automatic failure detection, and recovery. You are aware that his existing set up currently uses an Oracle database. Which of the following AWS databases would be best for accomplishing this

task? A. Amazon RDS B. Amazon Redshift C. Amazon SimpleDB D. Amazon ElastiCache Answer: A Explanation: Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery. Reference:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html> QUESTION 436 True or false: A VPC contains multiple subnets, where each subnet can span multiple Availability Zones. A. This is true only if requested during the set-up of VPC. B. This is true. C. This is false. D. This is true only for US regions. Answer: C Explanation: A VPC can span several Availability Zones. In contrast, a subnet must reside within a single Availability Zone. Reference: <https://aws.amazon.com/vpc/faqs/> QUESTION 437 An edge location refers to which Amazon Web Service? A. An edge location is referred to the network configured within a Zone or Region B. An edge location is an AWS Region C. An edge location is the location of the data center used for Amazon CloudFront. D. An edge location is a Zone within an AWS Region Answer: C Explanation: Amazon CloudFront is a content distribution network. A content delivery network or content distribution network (CDN) is a large distributed system of servers deployed in multiple data centers across the world. The location of the data center used for CDN is called edge location. Amazon CloudFront can cache static content at each edge location. This means that your popular static content (e.g., your site's logo, navigational images, cascading style sheets, JavaScript code, etc.) will be available at a nearby edge location for the browsers to download with low latency and improved performance for viewers. Caching popular static content with Amazon CloudFront also helps you offload requests for such files from your origin sever ?CloudFront serves the cached copy when available and only makes a request to your origin server if the edge location receiving the browser's request does not have a copy of the file. Reference:

<http://aws.amazon.com/cloudfront/> QUESTION 438 You are looking at ways to improve some existing infrastructure as it seems a lot of engineering resources are being taken up with basic management and monitoring tasks and the costs seem to be excessive. You are thinking of deploying Amazon ElastiCache to help. Which of the following statements is true in regards to ElastiCache? A. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will be more. B. You can't improve load and response times to user actions and queries but you can reduce the cost associated with scaling web applications. C. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will remain the same. D. You can improve load and response times to user actions and queries and also reduce the cost associated with scaling web applications. Answer: D Explanation: Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications. Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications. Reference: <https://aws.amazon.com/elasticache/faqs/> QUESTION 439 Do Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance? A. Yes, they do but only if they are detached from the instance. B. No, you cannot attach EBS volumes to an instance. C. No, they are dependent. D. Yes, they do. Answer: D Explanation: An Amazon EBS volume behaves like a raw, unformatted, external block device that you can attach to a single instance. The volume persists independently from the running life of an Amazon EC2 instance. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html> QUESTION 440 Your supervisor has asked you to build a simple file synchronization service for your department. He doesn't want to spend too much money and he wants to be notified of any changes to files by email. What do you think would be the best Amazon service to use for the email solution? A. Amazon SES B. Amazon CloudSearch C. Amazon SWF D. Amazon AppStream Answer: A Explanation: File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution. Reference:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_filesync_08.pdf QUESTION 441 Does DynamoDB support in-place atomic updates? A. Yes B. No C. It does support in-place non-atomic updates D. It is not defined Answer: A Explanation: DynamoDB supports in-place atomic updates. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithItems.html#WorkingWithItems.AtomicCounters> QUESTION 442 Your manager has just given you access to multiple VPN connections that someone else has recently set up between all your company's offices. She needs you to make sure that the communication between the VPNs is secure. Which of the following services would be best for providing a low-cost hub-and-spoke model for primary or backup connectivity between these

remote offices? A. Amazon CloudFront B. AWS Direct Connect C. AWS CloudHSM D. AWS VPN CloudHub Answer: D Explanation: If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices. Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CloudHub.html QUESTION 443 Amazon EC2 provides a ____.

It is an HTTP or HTTPS request that uses the HTTP verbs GET or POST. A. web database B. .net framework C. Query API D.

C library Answer: C Explanation: Amazon EC2 provides a Query API. These requests are HTTP or HTTPS requests that use the HTTP verbs GET or POST and a Query parameter named Action. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/APIReference/making-api-requests.html> QUESTION 444 In Amazon AWS, which of the following statements is true of key pairs? A. Key pairs are used only for Amazon SDKs. B. Key pairs are used only for Amazon EC2 and Amazon CloudFront. C. Key pairs are used only for Elastic Load Balancing and AWS IAM. D. Key pairs are used for all Amazon services. Answer: B Explanation: Key pairs consist of a public and private key, where you use the private key to create a digital signature, and then AWS uses the corresponding public key to validate the signature. Key pairs are used only for Amazon EC2 and Amazon CloudFront. Reference: <http://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html>

QUESTION 445 Does Amazon DynamoDB support both increment and decrement atomic operations? A. Only increment, since decrement are inherently impossible with DynamoDB's data model. B. No, neither increment nor decrement operations. C. Yes, both increment and decrement operations. D. Only decrement, since increment are inherently impossible with DynamoDB's data model. Answer: C Explanation: Amazon DynamoDB supports increment and decrement atomic operations. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/APISummary.html> QUESTION 446 An organization has three separate AWS accounts, one each for development, testing, and production. The organization wants the testing team to have access to certain AWS resources in the production account. How can the organization achieve this? A. It is not possible to access resources of one account with another account. B. Create the IAM roles with cross account access. C. Create the IAM user in a test account, and allow it access to the production environment with the IAM policy. D. Create the IAM users with cross account access. Answer: B Explanation: An organization has multiple AWS accounts to isolate a development environment from a testing or production environment. At times the users from one account need to access resources in the other account, such as promoting an update from the development environment to the production environment. In this case the IAM role with cross account access will provide a solution. Cross account access lets one account share access to their resources with users in the other AWS accounts.

Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf QUESTION 447 You need to import several hundred megabytes of data from a local Oracle database to an Amazon RDS DB instance. What does AWS recommend you use to accomplish this? A. Oracle export/import utilities B. Oracle SQL Developer C. Oracle Data Pump D.

DBMS_FILE_TRANSFER Answer: C Explanation: How you import data into an Amazon RDS DB instance depends on the amount of data you have and the number and variety of database objects in your database. For example, you can use Oracle SQL Developer to import a simple, 20 MB database; you want to use Oracle Data Pump to import complex databases or databases that are several hundred megabytes or several terabytes in size. Reference:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Oracle.Procedural.Importing.html> QUESTION 448 A user has created an EBS volume with 1000 IOPS. What is the average IOPS that the user will get for most of the year as per EC2 SLA if the instance is attached to the EBS optimized instance? A. 950 B. 990 C. 1000 D. 900 Answer: D Explanation: As per AWS SLA if the instance is attached to an EBS-Optimized instance, then the Provisioned IOPS volumes are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time in a given year. Thus, if the user has created a volume of 1000 IOPS, the user will get a minimum 900 IOPS 99.9% time of the year. Reference: <http://aws.amazon.com/ec2/faqs/>

QUESTION 449 You need to migrate a large amount of data into the cloud that you have stored on a hard disk and you decide that the best way to accomplish this is with AWS Import/Export and you mail the hard disk to AWS. Which of the following statements is incorrect in regards to AWS Import/Export? A. It can export from Amazon S3 B. It can Import to Amazon Glacier C. It can export from Amazon Glacier. D. It can Import to Amazon EBS Answer: C Explanation: AWS Import/Export supports: Import to Amazon S3 Export from Amazon S3 Import to Amazon EBS Import to Amazon Glacier AWS Import/Export does not currently support export from Amazon EBS or Amazon Glacier. Reference: <https://docs.aws.amazon.com/AWSImportExport/latest/DG/whatisdisk.html>

QUESTION 450 You are in the process of creating a Route 53 DNS failover to direct traffic to two EC2 zones. Obviously, if one fails, you would like Route 53 to direct traffic to the other region. Each region has an ELB with some instances being distributed. What is the best way for you to configure the Route 53 health check? A. Route 53 doesn't support ELB with an internal health check. You need to create

your own Route 53 health check of the ELB B. Route 53 natively supports ELB with an internal health check. Turn "Evaluate target health" off and "Associate with Health Check" on and R53 will use the ELB's internal health check. C. Route 53 doesn't support ELB with an internal health check. You need to associate your resource record set for the ELB with your own health check D. Route 53 natively supports ELB with an internal health check. Turn "Evaluate target health" on and "Associate with Health Check" off and R53 will use the ELB's internal health check. Answer: D Explanation: With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly. When you enable this feature, Route 53 uses health checks--regularly making Internet requests to your application's endpoints from multiple locations around the world--to determine whether each endpoint of your application is up or down. To enable DNS Failover for an ELB endpoint, create an Alias record pointing to the ELB and set the "Evaluate Target Health" parameter to true. Route 53 creates and manages the health checks for your ELB automatically. You do not need to create your own Route 53 health check of the ELB. You also do not need to associate your resource record set for the ELB with your own health check, because Route 53 automatically associates it with the health checks that Route 53 manages on your behalf. The ELB health check will also inherit the health of your backend instances behind that ELB. Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/05/30/amazon-route-53-adds-elb-integration-for-dns-failover/> We ensure our new version AWS Certified Solutions Architect - Associate PDF and VCE dumps are 100% valid for passing exam, because Lead2pass is the top IT certification study training materials vendor. Many candidates have passed exam with the help of Lead2pass's VCE or PDF dumps. Lead2pass will update the study materials timely to make them be consistent with the current exam. Download the free demo on Lead2pass, you can pass the exam easily. AWS Certified Solutions Architect ? Associate new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDR1h2VU4tOHhDcW8> **2017 Amazon AWS Certified Solutions Architect ? Associate** exam dumps (All 680 Q&As) from Lead2pass:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> [100% Exam Pass Guaranteed]