

[2017 New Lead2pass Amazon AWS Certified Solutions Architect - Associate Exam Dumps Free Download (451-475)]

2017 August Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! AWS Certified Solutions Architect - Associate exam questions and answers provided by Lead2pass will guarantee you pass AWS Certified Solutions Architect - Associate exam, because Lead2pass is the top IT Certification study training materials vendor. Many candidates have passed exam with the help of Lead2pass. We offer the latest AWS Certified Solutions Architect - Associate PDF and VCE dumps with new version VCE player for free download, you can pass the exam beyond any doubt. Following questions and answers are all new published by Amazon Official Exam Center:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> QUESTION 451A user wants to use an EBS-backed Amazon EC2 instance for a temporary job. Based on the input data, the job is most likely to finish within a week. Which of the following steps should be followed to terminate the instance automatically once the job is finished? A. Configure the EC2 instance with a stop instance to terminate it.B. Configure the EC2 instance with ELB to terminate the instance when it remains idle.C. Configure the CloudWatch alarm on the instance that should perform the termination action once the instance is idle.D. Configure the Auto Scaling schedule activity that terminates the instance after 7 days.Answer: CExplanation:Auto Scaling can start and stop the instance at a pre-defined time. Here, the total running time is unknown. Thus, the user has to use the CloudWatch alarm, which monitors the CPU utilization. The user can create an alarm that is triggered when the average CPU utilization percentage has been lower than 10 percent for 24 hours, signaling that it is idle and no longer in use. When the utilization is below the threshold limit, it will terminate the instance as a part of the instance action.Reference:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html> QUESTION 452Which of the following is true of Amazon EC2 security group? A. You can modify the outbound rules for EC2-Classic.B. You can modify the rules for a security group only if the security group controls the traffic for just one instance.C. You can modify the rules for a security group only when a new instance is created.D. You can modify the rules for a security group at any time. Answer: D Explanation:A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group.Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html> QUESTION 453An Elastic IP address (EIP) is a static IP address designed for dynamic cloud computing. With an EIP, you can mask the failure of an instance or software by rapidly remapping the address to another instance in your account. Your EIP is associated with your AWS account, not a particular EC2 instance, and it remains associated with your account until you choose to explicitly release it. By default how many EIPs is each AWS account limited to on a per region basis? A. 1B. 5C. UnlimitedD. 10 Answer: BExplanation:By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication.If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses.Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html#using-instance-addressing-limit> QUESTION 454In Amazon EC2, partial instance-hours are billed _____. A. per second used in the hourB. per minute usedC. by combining partial segments into full hoursD. as full hours Answer: DExplanation:Partial instance-hours are billed to the next hour.Reference:

<http://aws.amazon.com/ec2/faqs/> QUESTION 455In EC2, what happens to the data in an instance store if an instance reboots (either intentionally or unintentionally)? A. Data is deleted from the instance store for security reasons.B. Data persists in the instance store.C. Data is partially present in the instance store.D. Data in the instance store will be lost. Answer: BExplanation:The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data on instance store volumes is lost under the following circumstances.Failure of an underlying driveStopping an Amazon EBS-backed instanceTerminating an instanceReference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/InstanceStorage.html> QUESTION 456You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet? A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.B. Subnet's traffic is routed to an internet gateway.C. Subnet's traffic is not routed to an internet gateway.D. None of these answers can be considered a public subnet. Answer: BExplanation:A virtual private cloud (VPC) is a virtual network dedicated to your AWS

account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC: you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings. A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the internet, and a private subnet for resources that won't be connected to the Internet. If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet. If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet. If a subnet doesn't have a route to the internet gateway, but has its traffic routed to a virtual private gateway, the subnet is known as a VPN-only subnet. Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 457 Can you specify the security group that you created for a VPC when you launch an instance in EC2-Classic? A. No, you can specify the security group created for EC2-Classic when you launch a VPC instance. B. No. C. Yes. D. No, you can specify the security group created for EC2-Classic to a non-VPC based instance only. Answer: B Explanation: If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. When you launch an instance in EC2-Classic, you must specify a security group in the same region as the instance. You can't specify a security group that you created for a VPC when you launch an instance in EC2-Classic. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#ec2-classic-security-groups>

QUESTION 458 While using the EC2 GET requests as URLs, the _____ is the URL that serves as the entry point for the web service. A. token B. endpoint C. action D. None of these Answer: B Explanation: The endpoint is the URL that serves as the entry point for the web service. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-query-api.html>

QUESTION 459 You have been asked to build a database warehouse using Amazon Redshift. You know a little about it, including that it is a SQL data warehouse solution, and uses industry standard ODBC and JDBC connections and PostgreSQL drivers. However you are not sure about what sort of storage it uses for database tables. What sort of storage does Amazon Redshift use for database tables? A. InnoDB Tables B. NDB data storage C. Columnar data storage D. NDB CLUSTER Storage Answer: C Explanation: Amazon Redshift achieves efficient storage and optimum query performance through a combination of massively parallel processing, columnar data storage, and very efficient, targeted data compression encoding schemes. Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O requirements and reduces the amount of data you need to load from disk. Reference: http://docs.aws.amazon.com/redshift/latest/dg/c_columnar_storage_disk_mem_mgmt.html

QUESTION 460 You are checking the workload on some of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes and it seems that the I/O latency is higher than you require. You should probably check the _____ to make sure that your application is not trying to drive more IOPS than you have provisioned. A. Amount of IOPS that are available B. Acknowledgement from the storage subsystem C. Average queue length D. Time it takes for the I/O operation to complete Answer: C Explanation: In EBS workload demand plays an important role in getting the most out of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes. In order for your volumes to deliver the amount of IOPS that are available, they need to have enough I/O requests sent to them. There is a relationship between the demand on the volumes, the amount of IOPS that are available to them, and the latency of the request (the amount of time it takes for the I/O operation to complete). Latency is the true end-to-end client time of an I/O operation; in other words, when the client sends a IO, how long does it take to get an acknowledgement from the storage subsystem that the IO read or write is complete. If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume). Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.html>

QUESTION 461 Which of the below mentioned options is not available when an instance is launched by Auto Scaling with EC2 Classic? A. Public IP B. Elastic IP C. Private DNS D. Private IP Answer: B Explanation: Auto Scaling supports both EC2 classic and EC2-VPC. When an instance is launched as a part of EC2 classic, it will have the public IP and DNS as well as the private IP and DNS. Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html>

QUESTION 462 You have been given a scope to deploy some AWS infrastructure for a large organisation. The requirements are that you will have a lot of EC2 instances but may need to add more when the average utilization of your Amazon EC2 fleet is high and conversely remove them when CPU utilization is low. Which AWS services would be best to use to accomplish this? A. Auto Scaling, Amazon CloudWatch and AWS Elastic Beanstalk B. Auto Scaling, Amazon CloudWatch and Elastic Load Balancing C. Amazon CloudFront, Amazon CloudWatch and Elastic Load Balancing D. AWS Elastic Beanstalk, Amazon CloudWatch and Elastic Load Balancing. Answer: B Explanation: Auto Scaling enables you to follow the demand curve for your applications closely, reducing the need to manually provision Amazon EC2 capacity in advance. For example, you can set a condition to add new Amazon EC2 instances in increments

to the Auto Scaling group when the average utilization of your Amazon EC2 fleet is high; and similarly, you can set a condition to remove instances in the same increments when CPU utilization is low. If you have predictable load changes, you can set a schedule through Auto Scaling to plan your scaling activities. You can use Amazon CloudWatch to send alarms to trigger scaling activities and Elastic Load Balancing to help distribute traffic to your instances within Auto Scaling groups. Auto Scaling enables you to run your Amazon EC2 fleet at optimal utilization. Reference: <http://aws.amazon.com/autoscaling/> QUESTION 463 You are building infrastructure for a data warehousing solution and an extra request has come through that there will be a lot of business reporting queries running all the time and you are not sure if your current DB instance will be able to handle it. What would be the best solution for this? A. DB Parameter Groups B. Read Replicas C. Multi-AZ DB Instance deployment D. Database Snapshots Answer: B Explanation: Read Replicas make it easy to take advantage of MySQL's built-in replication functionality to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. There are a variety of scenarios where deploying one or more Read Replicas for a given source DB Instance may make sense. Common reasons for deploying a Read Replica include: Scaling beyond the compute or I/O capacity of a single DB Instance for read-heavy database workloads. This excess read traffic can be directed to one or more Read Replicas. Serving read traffic while the source DB Instance is unavailable. If your source DB Instance cannot take I/O requests (e.g. due to I/O suspension for backups or scheduled maintenance), you can direct read traffic to your Read Replica(s). For this use case, keep in mind that the data on the Read Replica may be "stale" since the source DB Instance is unavailable. Business reporting or data warehousing scenarios; you may want business reporting queries to run against a Read Replica, rather than your primary, production DB Instance. Reference: <https://aws.amazon.com/rds/faqs/> QUESTION 464 In DynamoDB, could you use IAM to grant access to Amazon DynamoDB resources and API actions? A. In DynamoDB there is no need to grant access B. Depended to the type of access C. No D. Yes Answer: D Explanation: Amazon DynamoDB integrates with AWS Identity and Access Management (IAM). You can use AWS IAM to grant access to Amazon DynamoDB resources and API actions. To do this, you first write an AWS IAM policy, which is a document that explicitly lists the permissions you want to grant. You then attach that policy to an AWS IAM user or role. Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/UsingIAMWithDDB.html> QUESTION 465 Much of your company's data does not need to be accessed often, and can take several hours for retrieval time, so it's stored on Amazon Glacier. However someone within your organization has expressed concerns that his data is more sensitive than the other data, and is wondering whether the high level of encryption that he knows is on S3 is also used on the much cheaper Glacier service. Which of the following statements would be most applicable in regards to this concern? A. There is no encryption on Amazon Glacier, that's why it is cheaper. B. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3 but you can change it to AES-256 if you are willing to pay more. C. Amazon Glacier automatically encrypts the data using AES-256, the same as Amazon S3. D. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3. Answer: C Explanation: Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often, and for which retrieval times of several hours are suitable. Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks, and is built to be automatically self-healing. Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf> QUESTION 466 Your EBS volumes do not seem to be performing as expected and your team leader has requested you look into improving their performance. Which of the following is not a true statement relating to the performance of your EBS volumes? A. Frequent snapshots provide a higher level of data durability and they will not degrade the performance of your application while the snapshot is in progress. B. General Purpose (SSD) and Provisioned IOPS (SSD) volumes have a throughput limit of 128 MB/s per volume. C. There is a relationship between the maximum performance of your EBS volumes, the amount of I/O you are driving to them, and the amount of time it takes for each transaction to complete. D. There is a 5 to 50 percent reduction in IOPS when you first access each block of data on a newly created or restored EBS volume Answer: A Explanation: Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration. Frequent snapshots provide a higher level of data durability, but they may slightly degrade the performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.html> QUESTION 467 You've created your first load balancer and have registered your EC2 instances with the load balancer. Elastic Load Balancing routinely performs health checks on

all the registered EC2 instances and automatically distributes all incoming requests to the DNS name of your load balancer across your registered, healthy EC2 instances. By default, the load balancer uses the ___ protocol for checking the health of your instances.

A. HTTPS B. HTTP C. ICMP D. IPv6

Answer: B

Explanation: In Elastic Load Balancing a health configuration uses information such as protocol, ping port, ping path (URL), response timeout period, and health check interval to determine the health state of the instances registered with the load balancer. Currently, HTTP on port 80 is the default health check.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

QUESTION 468

A major finance organisation has engaged your company to set up a large data mining application. Using AWS you decide the best service for this is Amazon Elastic MapReduce (EMR) which you know uses Hadoop. Which of the following statements best describes Hadoop?

A. Hadoop is 3rd Party software which can be installed using AMIB. Hadoop is an open source python web framework C. Hadoop is an open source Java software framework D. Hadoop is an open source javascript framework

Answer: C

Explanation: Amazon EMR uses Apache Hadoop as its distributed data processing engine. Hadoop is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware. Hadoop implements a programming model named "MapReduce," where the data is divided into many small fragments of work, each of which may be executed on any node in the cluster. This framework has been widely used by developers, enterprises and startups and has proven to be a reliable software platform for processing up to petabytes of data on clusters of thousands of commodity machines.

Reference: <http://aws.amazon.com/elasticmapreduce/faqs/>

QUESTION 469

In Amazon EC2 Container Service, are other container types supported?

A. Yes, EC2 Container Service supports any container service you need. B. Yes, EC2 Container Service also supports Microsoft container service. C. No, Docker is the only container platform supported by EC2 Container Service presently. D. Yes, EC2 Container Service supports Microsoft container service and Openstack.

Answer: C

Explanation: In Amazon EC2 Container Service, Docker is the only container platform supported by EC2 Container Service presently.

Reference: <http://aws.amazon.com/ecs/faqs/>

QUESTION 470

_____ is a fast, flexible, fully managed push messaging service.

A. Amazon SNS B. Amazon SES C. Amazon SQS D. Amazon FPS

Answer: A

Explanation: Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, fully managed push messaging service. Amazon SNS makes it simple and cost-effective to push to mobile devices such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services.

Reference: http://aws.amazon.com/sns/?nc1=h_l2_as

QUESTION 471

As AWS grows, most of your clients' main concerns seem to be about security, especially when all of their competitors also seem to be using AWS. One of your clients asks you whether having a competitor who hosts their EC2 instances on the same physical host would make it easier for the competitor to hack into the client's data. Which of the following statements would be the best choice to put your client's mind at rest?

A. Different instances running on the same physical machine are isolated from each other via a 256-bit Advanced Encryption Standard (AES-256). B. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor and via a 256-bit Advanced Encryption Standard (AES-256). C. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor. D. Different instances running on the same physical machine are isolated from each other via IAM permissions.

Answer: C

Explanation: Amazon Elastic Compute Cloud (EC2) is a key component in Amazon's Infrastructure as a Service (IaaS), providing resizable computing capacity using server instances in AWS's data centers. Amazon EC2 is designed to make web-scale computing easier by enabling you to obtain and configure capacity with minimal friction. You create and launch instances, which are collections of platform hardware and software. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor. Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.

Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

QUESTION 472

In Amazon RDS, security groups are ideally used to:

A. Define maintenance period for database engines B. Launch Amazon RDS instances in a subnet C. Create, describe, modify, and delete DB instances D. Control what IP addresses or EC2 instances can connect to your databases on a DB instance

Answer: D

Explanation: In Amazon RDS, security groups are used to control what IP addresses or EC2 instances can connect to your databases on a DB instance. When you first create a DB instance, its firewall prevents any database access except through rules specified by an associated security group.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.html>

QUESTION 473

You need to set up a complex network infrastructure for your organization that will be reasonably easy to deploy, replicate, control, and track changes on. Which AWS service would be best to use to help you accomplish this?

A. AWS Import/Export B. AWS CloudFormation C. Amazon

Route 53D. Amazon CloudWatch Answer: BExplanation:AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. You don't need to individually create and configure AWS resources and figure out what's dependent on what. AWS CloudFormation handles all of that.Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html> QUESTION 474 You have just been given a scope for a new client who has an enormous amount of data(petabytes) that he constantly needs analysed. Currently he is paying a huge amount of money for a data warehousing company to do this for him and is wondering if AWS can provide a cheaper solution. Do you think AWS has a solution for this? A. Yes. Amazon SimpleDB. No. Not presentlyC. Yes. Amazon RedshiftD. Yes. Your choice of relational AMIs on Amazon EC2 and EBS Answer: CExplanation: Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small for just \$0.25 per hour with no commitments or upfront costs and scale to a petabyte or more for \$1,000 per terabyte per year, less than a tenth of most other data warehousing solutions. 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.Reference: https://aws.amazon.com/running_databases/#redshift_anchor QUESTION 475In an experiment, if the minimum size for an Auto Scaling group is 1 instance, which of the following statements holds true when you terminate the running instance? A. Auto Scaling must launch a new instance to replace it.B. Auto Scaling will raise an alarm and send a notification to the user for action. C. Auto Scaling must configure the schedule activity that terminates the instance after 5 days.D. Auto Scaling will terminate the experiment. Answer: AExplanation:If the minimum size for an Auto Scaling group is 1 instance, when you terminate the running instance, Auto Scaling must launch a new instance to replace it.Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html Lead2pass is the leader in supplying candidates with current and up-to-date training materials for Amazon certification and exam preparation. Comparing with others, our AWS Certified Solutions Architect - Associate exam questions are more authoritative and complete. 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 exam 100% pass. 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]