

[PDF&VCE AWS Certified Solutions Architect - Associate New Questions For Passing The AWS Certified Solutions Architect - Associate Certification Exam (326-350)]

[2016 October Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed!](#) Test your preparation for Amazon AWS CERTIFIED SOLUTIONS ARCHITECT - ASSOCIATE with these actual AWS CERTIFIED SOLUTIONS ARCHITECT - ASSOCIATE new questions below. Exam questions are a sure method to validate one's preparation for actual certification exam. Following questions and answers are all new published by Amazon Official Exam Center: <http://www.lead2pass.com/aws-certified-solutions-architect-associate.html>

QUESTION 326 A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user. Which two approaches can satisfy these objectives? (Choose 2 answers) A. Develop an identity broker that authenticates against IAM security Token service to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket. B. The application authenticates against LOAP and retrieves the name of an IAMrole associated with the user. The application then calls the IAM Security Token Service to assume that IAM role The application can use the temporary credentials to access the appropriate S3 bucket. C. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credentials The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket. D. The application authenticates against LDAP the application then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials the application can use the IAM temporary credentials to access the appropriate S3 bucket. E. The application authenticates against IAM Security Token Service using the LDAP credentials the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Answer: CD Explanation: User should authenticate against LDAP first; A and E are ruled out. B: Application cannot get IAM role from LDAP; ruled out. QUESTION 327 You are designing a multi-platform web application for AWS The application will run on EC2 instances and will be accessed from PCs, tablets and smart phones Supported accessing platforms are Windows, MACOS, IOS and Android Separate sticky session and SSL certificate setups are required for different platform types which of the following describes the most cost effective and performance efficient architecture setup? A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC. B. Set up one ELB for all platforms to distribute load among multiple instance under it Each EC2 instance implements ail functionality for a particular platform. C. Set up two ELBs The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms for each ELB run separate EC2 instance groups to handle the web application for each platform. D. Assign multiple ELBS to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type Session stickiness and SSL termination are done at the ELBs. Answer: D Explanation: One ELB cannot handle different SSL certificates but since we are using sticky sessions it must be handled at the ELB level. SSL could be handled on the EC2 instances only with TCP configured ELB, ELB supports sticky sessions only in HTTP/HTTPS configurations. The way the Elastic Load Balancer does session stickiness is on a HTTP/HTTPS listener is by utilizing an HTTP cookie. If SSL traffic is not terminated on the Elastic Load Balancer and is terminated on the back-end instance, the Elastic Load Balancer has no visibility into the HTTP headers and therefore can not set or read any of the HTTP headers being passed back and forth.

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-sticky-sessions.html> QUESTION 328 Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst In web traffic due to a company announcement Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructures ability to handle unexpected increases in traffic. The application currently consists of 2 tiers a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required? A. Offload traffic from on-premises environment Setup a CloudFront distribution and configure CloudFront to cache objects from a custom origin Choose to customize your object cache behavior, and select a TTL that objects should exist in cache. B. Migrate to AWS Use VM import `Export to quickly convert an on-premises web server to an AMI create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffic Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to

migrate the database. C. Failover environment: Create an S3 bucket and configure it for website hosting Migrate your DNS to Route53 using zone (file import and leverage Route53 DNS failover to failover to the S3 hosted website. D. Hybrid environment Create an AMI which can be used to launch web servers in EC2 Create an Auto Scaling group which uses the * AMI to scale the web tier based on incoming traffic Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted in AWS. Answer: A QUESTION 329 Your company produces custom commissioned one-of-a-kind skiing helmets combining night vision with custom technical enhancements. Customers can show off their Individuality on the ski slopes and have access to head-up-displays, GPS rear-view cams and any other technical innovation they wish to embed in the helmet. The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards. Assessments are a mixture of human and automated assessments. You need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUD across a cluster of servers with low latency networking. What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time? A. Use AWS Data Pipeline to manage movement of data & meta-data and assessments Use an auto-scaling group of G2 instances in a placement group. B. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data Use an auto-scaling group of G2 instances in a placement group. C. Use Amazon Simple Workflow (SWF) to manage assessments movement of data & meta-data Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O Virtualization). D. Use AWS data Pipeline to manage movement of data & meta-data and assessments use auto-scaling group of C3 with SR-IOV (Single Root I/O virtualization). Answer: B QUESTION 330 You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case? A. Use Storage Gateway and configure it to use Gateway Cached volumes. B. Configure your backup software to use S3 as the target for your data backups. C. Configure your backup software to use Glacier as the target for your data backups. D. Use Storage Gateway and configure it to use Gateway Stored volumes. Answer: C Explanation: A,D is wrong because Storage Gateway has a max limit on how much you can store:

<https://aws.amazon.com/storagegateway/faqs/> Q: What is the maximum size of a volume? Each gateway-cached volume can store up to 32 TB of data. Data written to the volume is cached on your on-premises hardware and asynchronously uploaded to AWS for durable storage. Each gateway-stored volume can store up to 16 TB of data. Data written to the volume is stored on your on-premises hardware and asynchronously backed up to AWS for point-in-time snapshots. B is wrong because it says "Your backup application is only able to write to POSIX-compatible block-based storage." S3 is object based storage. QUESTION 331 You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8x large Instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job? A. Create more smaller files on Amazon S3. B. Add additional cc2 8x large instances by introducing a task group. C. Use smaller instances that have higher aggregate I/O performance. D. Create fewer, larger files on Amazon S3. Answer: C Explanation: <https://aws.amazon.com/elasticmapreduce/faqs/> QUESTION 332 Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse. Your CFO requests that you optimize the cost structure for this system. Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data? A. Use reduced redundancy storage (RRS) for PDF and csv data in Amazon S3. Add Spot instances to Amazon EMR jobs Use Reserved Instances for Amazon Redshift. B. Use reduced redundancy storage (RRS) for all data in S3. Use a combination of Spot instances and Reserved Instances for Amazon EMR jobs use Reserved instances for Amazon Redshift. C. Use reduced redundancy storage (RRS) for all data in Amazon S3 Add Spot Instances to Amazon EMR jobs Use Reserved Instances for Amazon Redshift. D. Use reduced redundancy storage (RRS) for PDF and csv data in S3 Add Spot Instances to EMR jobs Use Spot Instances for Amazon Redshift. Answer: A Explanation: Reserved Instances (a.k.a. Reserved Nodes) are appropriate for steady-state production workloads, and offer significant discounts over On-Demand pricing. <https://aws.amazon.com/redshift> Q: What are some EMR best practices? If you are running EMR in production you should specify an AMI version, Hive version, Pig version, etc. to make sure the version does not change unexpectedly (e.g. when EMR later adds support for a newer version). If your cluster is mission critical, only use Spot instances for task nodes because if the Spot price increases you may lose the instances. In development, use logging and enable debugging to spot and correct errors faster. If you are using GZIP, keep your file size to 1?2 GB because GZIP files cannot be split. Click here to download the white paper on Amazon EMR best practices.

<https://aws.amazon.com/elasticmapreduce/faqs/> QUESTION 333 You are the new IT architect in a company that operates a mobile sleep tracking application. When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend. The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table. Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SMS mobile push notifications that new data is available, which is parsed and visualized by (The mobile app). Currently, you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? (Choose 2 answers) A. Create a new Amazon DynamoDB table each day and drop the one for the previous day after its data is on Amazon S3. B. Have the mobile app access Amazon DynamoDB directly instead of JSON files stored on Amazon S3. C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput. D. Introduce Amazon ElastiCache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput. E. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3. Answer: CD Explanation:

<https://d0.awsstatic.com/whitepapers/performance-at-scale-with-amazon-elasticache.pdf> QUESTION 334 Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally, often on the move, and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant. How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery? A. Elastic Transcoder to transcode original high-resolution MP4 videos to HLS, S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days, CloudFront to serve HLS transcoded videos from S3. B. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue, S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days, CloudFront to serve HLS transcoding videos from Glacier. C. Elastic Transcoder to transcode original high-resolution MP4 videos to HLS, EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days, CloudFront to serve HLS transcoded videos from EC2. D.

A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue, EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days, CloudFront to serve HLS transcoded videos from EC2. Answer: A QUESTION 335 You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoDB for their dynamic data and then archiving nightly into S3 for further processing with EMR. They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access. Which approach provides a cost effective, scalable mitigation to this kind of attack? A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC. They would then establish Internet connectivity into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC. B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet. C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 instances running a host-based WAF. They would redirect Route 53 to resolve to the new WAF tier ELB. The WAF tier would then pass the traffic to the current web tier. The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group. D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering. This will enable the ELB itself to perform WAF functionality. Answer: C QUESTION 336 You currently operate a web application in the AWS US-East region. The application runs on an auto-scaled layer of EC2 instances and an RDS Multi-AZ database. Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2, IAM, and RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend? A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles, S3 bucket policies, and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs. B. Create a new CloudTrail with one new S3 bucket to store the logs. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs. C. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs. D. Create three new CloudTrail trails with three new S3 buckets to store the logs: one for the AWS Management console, one for AWS SDKs, and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs. Answer: A QUESTION 337 An enterprise wants to use a third-party SaaS application. The SaaS application

needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.

B. Create an IAM user within the enterprise account, assign a user policy to the IAM user that allows only the actions required by the SaaS application, create a new access and secret key for the user, and provide these credentials to the SaaS provider.

C. Create an IAM role for cross-account access that allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.

D. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required for the SaaS application to work, provide the role ARN to the SaaS provider to use when launching their application instances.

Answer: C
Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

QUESTION 338 You are designing a data leak prevention solution for your VPC environment. You want your VPC instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party URLs by their URLs. You want to explicitly deny any other outbound connections from your VPC instances to hosts on the internet. Which of the following options would you consider?

A. Configure a web proxy server in your VPC and enforce URL-based rules for outbound access. Remove default routes.

B. Implement security groups and configure outbound rules to only permit traffic to software depots.

C. Move all your instances into private VPC subnets, remove default routes from all routing tables, and add specific routes to the software depots and distributions only.

D. Implement network access control lists to all specific destinations, with an implicit deny as a rule.

Answer: A
Explanation: Organizations usually implement proxy solutions to provide URL and web content filtering, IDS/IPS, data loss prevention, monitoring, and advanced threat protection.

https://d0.awsstatic.com/aws-answers/Controlling_VPC_Egress_Traffic.pdf

QUESTION 339 An administrator is using Amazon CloudFormation to deploy a three-tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage. When creating the CloudFormation template, which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.

B. Use the Parameter section in the CloudFormation template to save the user input Access and Secret Keys from an already created IAM user that has the permissions required to read and write from the required DynamoDB table.

C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.

D. Create an Identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and Secret Keys and pass them to the application instance through user-data.

Answer: C

QUESTION 340 An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 instances. The customer's security policy requires that every outbound connection from these instances to any other service within the customer's Virtual Private Cloud must be authenticated using a unique x.509 certificate that contains the specific instance-id. In addition, an x.509 certificate must be designed by the customer's Key Management Service in order to be trusted for authentication. Which of the following configurations will support these requirements?

A. Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure the AutoScaling group to launch instances with this role. Have the instances bootstrap to get the certificate from Amazon S3 upon first boot.

B. Embed a certificate into the Amazon Machine Image that is used by the AutoScaling group. Have the launched instances generate a certificate signature request with the instance's assigned instance-id to the Key Management Service for signature.

C. Configure the AutoScaling group to send an SNS notification of the launch of a new instance to the trusted Key Management Service. Have the Key Management Service generate a signed certificate and send it directly to the newly launched instance.

D. Configure the launched instances to generate a new certificate upon first boot. Have the Key Management Service poll the AutoScaling group for associated instances and send new instances a certificate signature (that contains the specific instance-id).

Answer: D
Explanation: Because a unique x.509 certificate that contains the specific instance-id.

QUESTION 341 Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed. Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary. You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console. Which option below will meet the needs for your NOC members?

A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.

B. Use Web Identity Federation to retrieve AWS temporary security credentials to enable your NOC

members to sign in to the AWS Management Console. C. Use your on-premises SAML 2 O-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint. D. Use your on-premises SAML2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console. Answer: C Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_enable-console-saml.html QUESTION 342 You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient. Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers) A. Configure ELB with TCP listeners on TCP/4d3. And place the Web servers behind it. B. Configure your Web servers with EIPS Place the Web servers in a Route53 Record Set and configure health checks against all Web servers. C. Configure ELB with HTTPS listeners, and place the Web servers behind it. D. Configure your web servers as the origins for a CloudFront distribution. Use custom SSL certificates on your CloudFront distribution. Answer: CD Explanation:

TCP/443 or HTTPS listener either way you can configure, but you can only upload ssl certificate on HTTPS listener. QUESTION 343 You are designing a connectivity solution between on-premises infrastructure and Amazon VPC Your server's on-premises will be communicating with your VPC instances You will be establishing IPsec tunnels over the internet You will be using VPN gateways and terminating the IPsec tunnels on AWS-supported customer gateways. Which of the following objectives would you achieve by implementing an IPsec tunnel as outlined above? (Choose 4 answers) A. End-to-end protection of data in transit B. End-to-end Identity authentication C. Data encryption across the Internet D. Protection of data in transit over the Internet E. Peer identity authentication between VPN gateway and customer gateway F. Data integrity protection across the Internet Answer: CDEF

QUESTION 344 You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet. Which of the following options would you consider? (Choose 2 answers) A. Implement IDS/IPS agents on each Instance running In VPC B. Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic. C. Implement Elastic Load Balancing with SSL listeners In front of the web applications D. Implement a reverse proxy layer in front of web servers and configure IDS/IPS agents on each reverse proxy server. Answer: AD Explanation: EC2 does not allow promiscuous mode, and you cannot put something in between the ELB and the web server (like a listener or IDP)

QUESTION 345 You are designing a photo sharing mobile app the application will store all pictures in a single Amazon S3 bucket. Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3. You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photo-sharing mobile application? A. Create a set of long-term credentials using AWS Security Token Service with appropriate permissions Store these credentials in the mobile app and use them to access Amazon S3. B. Record the user's Information in Amazon RDS and create a role in IAM with appropriate permissions. When the user uses their mobile app create temporary credentials using the AWS Security Token Service 'AssumeRole' function Store these credentials in the mobile app's memory and use them to access Amazon S3 Generate new credentials the next time the user runs the mobile app. C. Record the user's Information In Amazon DynamoDB. When the user uses their mobile app create temporary credentials using AWS Security Token Service with appropriate permissions Store these credentials in the mobile app's memory and use them to access Amazon S3 Generate new credentials the next time the user runs the mobile app. D. Create IAM user. Assign appropriate permissions to the IAM user Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3. E. Create an IAM user. Update the bucket policy with appropriate permissions for the IAM user Generate an access Key and secret Key for the IAM user, store them In the mobile app and use these credentials to access Amazon S3. Answer: B Explanation: We can use either RDS or DynamoDB, however in our given answers, IAM role is mentioned only with RDS, so I would go with Answer B. Question was explicitly focused on security, so IAM with RDS is the best choice.

QUESTION 346 You have an application running on an EC2 Instance which will allow users to download files from a private S3 bucket using a pre-assigned URL. Before generating the URL the application should verify the existence of the file in S3. How should the application use AWS credentials to access the S3 bucket securely? A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application. B. Create a IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the IAM user and retrieve the IAM user's credentials from the EC2 instance user data. C. Create an IAM role for EC2 that allows list access to objects in the S3 bucket. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata D. Create an IAM user for the application with permissions that allow list access to the S3 bucket. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user. Answer: C Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html> QUESTION 347 You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTPS connections to specific domains from their EC2-hosted applications you deploy a single EC2 instance running proxy software and configure It to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration You have a nightly maintenance window of 10 minutes where all instances fetch new software updates. Each update is about 200MB in size and there are 500 instances in the VPC that routinely fetch updates After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? (Choose 2 answers) A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time. B. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail C. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW) D. You are running the proxy on a file-sized EC2 instance in a private subnet and its network throughput is being throttled by a NAT running on an undersized EC2 instance E. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy. Answer: AB
Explanation: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ec2-config.html> QUESTION 348 To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity As a result, your company purchases two C3.2xlarge medium utilization RIs. You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively? A. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin B. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand ml large instances when triggered by Cloudwatch shut off c3 2xlarge instances C. Route traffic to EC2 ml large and c3 2xlarge instances directly using Route 53 latency based routing and health checks shut off ELB D. Configure ELB with two c3 2xlarge Instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances Shut on ml large instances. Answer: A
Explanation: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html> QUESTION 349 A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements? A. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch. And RDS with read replicas B. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and RDS with read replicas C. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch. And multi-AZ RDS D. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch and multi-AZ RDS Answer: A
Explanation: ? A read only reporting site? ? so stateless and read-replicas can be used to scale. Multi-AZ will not provide the scaling requirements. QUESTION 350 You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDS extra large DB instance with 10,000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range. Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds. How would you improve page load times for your users? (Choose 3 answers) A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively. B. Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries C. Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site D. Switch Amazon RDS database to the high memory extra large Instance type E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region. Answer: BCD These Amazon AWS CERTIFIED SOLUTIONS ARCHITECT - ASSOCIATE exam questions are all a small selection of questions. If you want to practice more questions for actual AWS CERTIFIED SOLUTIONS ARCHITECT - ASSOCIATE exam, use the links at the end of this document. Also you can find links for AWS CERTIFIED SOLUTIONS ARCHITECT - ASSOCIATE VCE software that is great for preparation and self-assessment for Amazon AWS CERTIFIED SOLUTIONS ARCHITECT - ASSOCIATE exam. AWS Certified

Solutions Architect ? Associate new questions on Google Drive:

<https://drive.google.com/open?id=0B3Syig5i8gpDNIBGazRSTENUQW8> **2016 Amazon AWS Certified Solutions Architect ?**

Associate exam dumps (All 423 Q&As) from Lead2pass:

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