SAP-C01 ^{Q&As}

AWS Certified Solutions Architect - Professional (SAP-C01)

Pass Amazon SAP-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.lead4pass.com/aws-solution-architect-professional.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.

B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.

C. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric.

D. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Correct Answer: D

QUESTION 2

IAM Secure and Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too.

Which of the below mentioned options suits their requirements?

A. Since ELB is internet facing, it is recommended to setup HAProxy as the Load balancer within the VPC.

B. Create an Internet facing ELB with VPC and configure all the App servers with it.

C. The user should make ELB with EC2-CLASSIC and enable SSH with it for security.

D. Create an Internal Load balancer with VPC and register all the App servers with it.

Correct Answer: D

Explanation: The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.

Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/vpc-loadbalancertypes.html

---

QUESTION 3

A company operates a group of imaging satellites. The satellites stream data to one of the company\\\'s ground stations where processing creates about 5 GB of images per minute. This data is added to network-attached storage, where 2 PB of data are already stored.

The company runs a website that allows its customers to access and purchase the images over the Internet. This website is also running in the ground station. Usage analysis shows that customers are most likely to access images that have been captured in the last 24 hours.

The company would like to migrate the image storage and distribution system to AWS to reduce costs and increase the number of customers that can be served.

Which AWS architecture and migration strategy will meet these requirements?

A. Use multiple AWS Snowball appliances to migrate the existing imagery to Amazon S3. Create a 1-Gb AWS Direct Connect connection from the ground station to AWS, and upload new data to Amazon S3 through the Direct Connect connection. Migrate the data distribution website to Amazon EC2 instances. By using Amazon S3 as an origin, have this website serve the data through Amazon CloudFront by creating signed URLs.

B. Create a 1-Gb Direct Connect connection from the ground station to AWS. Use the AWS Command Line Interface to copy the existing data and upload new data to Amazon S3 over the Direct Connect connection. Migrate the data distribution website to EC2 instances. By using Amazon S3 as an origin, have this website serve the data through CloudFront by creating signed URLs.

C. Use multiple Snowball appliances to migrate the existing images to Amazon S3. Upload new data by regularly using Snowball appliances to upload data from the network-attached storage. Migrate the data distribution website to EC2 instances. By using Amazon S3 as an origin, have this website serve the data through CloudFront by creating signed URLs.

D. Use multiple Snowball appliances to migrate the existing images to an Amazon EFS file system. Create a 1-Gb Direct Connect connection from the ground station to AWS, and upload new data by mounting the EFS file system over the Direct Connect connection. Migrate the data distribution website to EC2 instances. By using webservers in EC2 that mount the EFS file system as the origin, have this website serve the data through CloudFront by creating signed URLs.

Correct Answer: B

QUESTION 4

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.

What AWS architecture would you recommend?

A. ASK their customers to use an S3 client instead of an FTP client. Create a single S3 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the \\\'username\\\' Policy variable.

B. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.

C. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.

D. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

Correct Answer: A

Reference:

https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-folders- in-anamazon-s3-bucket/

QUESTION 5

A company has a single AWS master billing account, which is the root of the AWS Organizations hierarchy.

The company has multiple AWS accounts within this hierarchy, all organized into organization units (OUs). More OUs and AWS accounts will continue to be created as other parts of the business migrate applications to AWS. These business units may need to use different AWS services. The Security team is implementing the following requirements for all current and future AWS accounts:

1.

Control policies must be applied across all accounts to prohibit AWS servers.

2.

Exceptions to the control policies are allowed based on valid use cases.

Which solution will meet these requirements with minimal optional overhead?

A. Use an SCP in Organizations to implement a deny list of AWS servers. Apply this SCP at the level. For any specific exceptions for an OU, create a new SCP for that OU and add the required AWS services to the allow list.

B. Use an SCP in Organizations to implement a deny list of AWS service. Apply this SCP at the root level and each OU. Remove the default AWS managed SCP from the root level and all OU levels. For any specific exceptions, modify the SCP attached to that OU, and add the required AWS services to the allow list.

C. Use an SCP in Organizations to implement a deny list of AWS service. Apply this SCP at each OU level. Leave the default AWS managed SCP at the root level. For any specific executions for an OU, create a new SCP for that OU.

D. Use an SCP in Organizations to implement an allow list of AWS services. Apply this SCP at the root level. Remove the default AWS managed SCP from the root level and all OU levels. For any specific exceptions for an OU, modify the SCP attached to that OU, and add the required AWS services to the allow list.

Correct Answer: B

QUESTION 6

An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate SSLs for the separate modules although it is already using VPC.

How can the organization achieve this with a single instance?

A. You have to launch two instances each in a separate subnet and allow VPC peering for a single IP.

B. Create a VPC instance which will have multiple network interfaces with multiple elastic IP addresses.

C. Create a VPC instance which will have both the ACL and the security group attached to it and have separate rules for each IP address.

D. Create a VPC instance which will have multiple subnets attached to it and each will have a separate IP address.

Correct Answer: B

Explanation: A Virtual Private Cloud (VPC) is a virtual network dedicated to the user\\'s AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances. The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. With each network interface the organization can assign an EIP. This scenario helps when the user wants to host multiple websites on a single EC2 instance by using multiple SSL certificates on a single server and associating each certificate with a specific EIP address. It also helps in scenarios for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.

Reference:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html

---

QUESTION 7

A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a REST API that the vendor hosts in its AWS environment. The vendor offers multiple options for connectivity to the API and is working with the company to find the best way to connect.

The company\\'s AWS account does not allow outbound internet access from its AWS environment. The vendor\\'s services run on AWS in the same Region as the company\\'s applications.

A solutions architect must implement connectivity to the vendor\\'s API so that the API is highly available in the company\\'s VPC.

Which solution will meet these requirements?

A. Connect to the vendor\\'s public API address for the data service

B. Connect to the vendor by way of a VPC peering connection between the vendor\\'s VPC and the company\\'s VPC

C. Connect to the vendor by way of a VPC endpoint service that uses AWS PrivateLink

D. Connect to a public bastion host that the vendor provides. Tunnel the API traffic

Correct Answer: D

Reference: https://docs.oracle.com/en-us/iaas/big-data/doc/use-bastion-host-connect-your-service.html

QUESTION 8

You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups.

Which RAID setup is not recommended for Amazon EBS?

A. RAID 1 only

B. RAID 5 only

C. RAID 5 and RAID 6 D. RAID 0 only

Correct Answer: C

Explanation: With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.

Reference:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 9

Which EC2 functionality allows the user to place the Cluster Compute instances in clusters?

A. Cluster group

B. Cluster security group

C. GPU units

D. Cluster placement group

Correct Answer: D

Explanation:

The Amazon EC2 cluster placement group functionality allows users to group cluster compute instances in

clusters.

Reference:

https://aws.amazon.com/ec2/faqs/

QUESTION 10

Select the correct statement about Amazon ElastiCache.

A. It makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud.

B. It allows you to quickly deploy your cache environment only if you install software.

C. It does not integrate with other Amazon Web Services.

D. It cannot run in the Amazon Virtual Private Cloud (Amazon VPC) environment.

Correct Answer: A

Explanation: ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in memory cache environment in the cloud. It provides a high-performance, scalable, and cost- effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.

Reference:

http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

QUESTION 11

What is a possible reason you would need to edit claims issued in a SAML token?

A. The NameIdentifier claim cannot be the same as the username stored in AD.

B. Authentication fails consistently.

C. The NameIdentifier claim cannot be the same as the claim URI.

D. The NameIdentifier claim must be the same as the username stored in AD.

Correct Answer: A

Explanation:

The two reasons you would need to edit claims issued in a SAML token are:

The NameIdentifier claim cannot be the same as the username stored in AD, and The app requires a

different set of claim URIs.

Reference:

https://azure.microsoft.com/en-us/documentation/articles/active-directory-saml-claims-customization/

QUESTION 12

The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the "504 Gateway Timeout" error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels.

Which of the following steps would be optimal for debugging these application issues? (Choose two.)

A. Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.

B. Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.

C. Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.

D. Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.

E. Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.

Correct Answer: BD

Explanation: Firstly "A 504 Gateway Timeout Error means your web server didn\\'t receive a timely response from another server upstream when it attempted to load one of your web pages. Put simply, your web servers aren\\'t communicating with each other fast enough". This specific issue is addressed in the AWS article "Tracing, Logging and Monitoring an API Gateway API".

Reference:

https://docs.amazonaws.cn/en_us/apigateway/latest/developerguide/monitoring_overview.html

Latest SAP-C01 Dumps   SAP-C01 Practice Test   SAP-C01 Exam Questions

To Read the Whole Q&As, please purchase the Complete Version from Our website.

Try our product !

100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
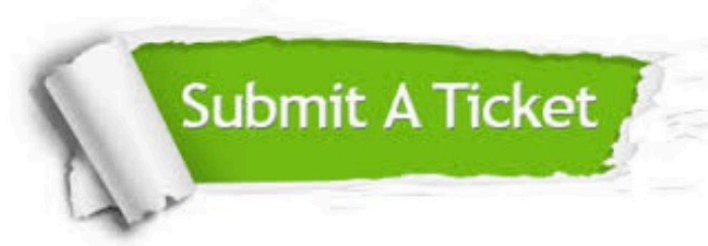Multi-Platform capabilities - Windows, Mac, Android, iPhone, iPod, iPad, Kindle

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

https://www.lead4pass.com/allproducts

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket: