


FreeDumps



WHAT PEOPLE SAY

Disclaimer Policy: The site does not guarantee the content of the comments. Because of the different time and the changes in the scope of the exam, it can produce different effect. Before you purchase the dump, please carefully read the product introduction from the page. In addition, please be advised the site will not be responsible for the content of the comments and contradictions between users.

“ Can not believe C2040-440! it is really same with the exam



ATWOOD

“ If you want to pass exam casually I advise you to purchase study guide. A2010-578 study guide have a part of questions with real test.



MILES

<http://www.freedumps.top>

Latest valid dumps torrent and free demo for certification exam prep

Exam : **SOA-C02**

Title : AWS Certified SysOps
Administrator - Associate
(SOA-C02)

Vendor : Amazon

Version : DEMO

NO.1 A SysOps administrator is testing an application that is hosted on five Amazon EC2 instances. The instances run in an Auto Scaling group behind an Application Load Balancer (ALB). High CPU utilization during load testing is causing the Auto Scaling group to scale out. The SysOps administrator must troubleshoot to find the root cause of the high CPU utilization before the Auto Scaling group scales out.

Which action should the SysOps administrator take to meet these requirements?

- A.** Enable instance scale-in protection.
- B.** Place the instance into the Standby state.
- C.** Remove the listener from the ALB
- D.** Suspend the Launch and Terminate process types.

Answer: D

Explanation:

You can put an instance that is in the InService state into the Standby state, update or troubleshoot the instance, and then return the instance to service. Instances that are on standby are still part of the Auto Scaling group, but they do not actively handle load balancer traffic.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-enter-exit-standby.html>

NO.2 A company hosts its website on Amazon EC2 instances behind an Application Load Balancer. The company manages its DNS with Amazon Route 53 and wants to point its domain's zone apex to the website.

Which type of record should be used to meet these requirements?

- A.** An AAAA record for the domain's zone apex
- B.** An A record for the domain's zone apex
- C.** A CNAME record for the domain's zone apex
- D.** An alias record for the domain's zone apex

Answer: D

Explanation:

Route 53 supports redirection of zone apex to the ALB via alias.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html>

NO.3 A company operates compute resources in a VPC and in the company's on-premises data center. The company already has an AWS Direct Connect connection between the VPC and the on-premises data center. A SysOps administrator needs to ensure that Amazon EC2 instances in the VPC can resolve DNS names for hosts in the on-premises data center.

Which solution will meet this requirement with the LEAST amount of ongoing maintenance?

- A.** Create an Amazon Route 53 private hosted zone. Populate the zone with the hostnames and IP addresses of the hosts in the on-premises data center.
- B.** Create an Amazon Route 53 Resolver outbound endpoint. Add the IP addresses of an on-premises DNS server for the domain names that need to be forwarded.
- C.** Set up a forwarding rule for reverse DNS queries in Amazon Route 53 Resolver. Set the `enableDnsHostnames` attribute to true for the VPC.
- D.** Add the hostnames and IP addresses for the on-premises hosts to the `/etc/hosts` file of each EC2

instance.

Answer: B

Explanation:

By creating an Amazon Route 53 Resolver outbound endpoint and configuring forwarding rules to send DNS queries to the on-premises DNS servers, EC2 instances in the VPC can dynamically resolve hostnames in the on-premises data center. This solution leverages AWS managed DNS forwarding and minimizes ongoing maintenance compared to manually maintaining DNS records or host file configurations.

NO.4 A team of on-call engineers frequently needs to connect to Amazon EC2 instances in a private subnet to troubleshoot and run commands.

The instances use either the latest AWS-provided Windows Amazon Machine Images (AMIs) or Amazon Linux AMIs.

The team has an existing IAM role for authorization. A SysOps administrator must provide the team with access to the instances by granting IAM permissions to this role.

Which solution will meet this requirement?

A. Add a statement to the IAM role policy to allow the `ssm:StartSession` action on the instances. Instruct the team to use AWS Systems Manager Session Manager to connect to the instances by using the assumed IAM role.

B. Associate an Elastic IP address and a security group with each instance.

Add the engineers' IP addresses to the security group inbound rules.

Add a statement to the IAM role policy to allow the `ec2:AuthorizeSecurityGroupIngress` action so that the team can connect to the instances.

C. Create a bastion host with an EC2 instance, and associate the bastion host with the VPC.

Add a statement to the IAM role policy to allow the `ec2:CreateVpnConnection` action on the bastion host.

Instruct the team to use the bastion host endpoint to connect to the instances.

D. Create an internet-facing Network Load Balancer.

Use two listeners. Forward port 22 to a target group of Linux instances.

Forward port 3389 to a target group of Windows instances.

Add a statement to the IAM role policy to allow the `ec2:CreateRoute` action so that the team can connect to the instances.

Answer: A

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NO.5 A SysOps administrator needs to provision a new fleet of Amazon EC2 Spot Instances in an Amazon EC2 Auto Scaling group. The Auto Scaling group will use a wide range of instance types. The configured fleet must come from pools that have the most availability for the number of instances that are launched.

Which solution will meet these requirements?

A. Launch the Spot Instances up to the maximum capacity of the Auto Scaling group.

B. Launch the Spot Instances by using the diversified strategy.

C. Launch the Spot Instances by using the capacity optimized strategy.

D. Use the Spot Instance advisor to help determine the best Spot allocation strategy.

Answer: D

Explanation:

The Spot Instance advisor helps you determine pools with the least chance of interruption and provides the savings you get over on-demand rates.

<https://aws.amazon.com/ec2/spot/instance-advisor/>

NO.6 A company is running a website on Amazon EC2 instances that are in an Auto Scaling group. When the website traffic increases, additional instances take several minutes to become available because of a longrunning user data script that installs software.

A SysOps administrator must decrease the time that is required for new instances to become available.

Which action should the SysOps administrator take to meet this requirement?

- A.** Reduce the scaling thresholds so that instances are added before traffic increases.
- B.** Purchase Reserved Instances to cover 100% of the maximum capacity of the Auto Scaling group.
- C.** Update the Auto Scaling group to launch instances that have a storage optimized instance type.
- D.** Use EC2 Image Builder to prepare an Amazon Machine Image (AMI) that has pre-installed software.

Answer: C

NO.7 A company has a cluster of Linux Amazon EC2 Spot Instances that read many files from and write many files to attached Amazon Elastic Block Store (Amazon EBS) volumes. The EC2 instances are frequently started and stopped. As part of the process when an EC2 instance starts, an EBS volume is restored from a snapshot.

EBS volumes that are restored from snapshots are experiencing initial performance that is lower than expected. The company's workload needs almost all the provisioned IOPS on the attached EBS volumes. The EC2 instances are unable to support the workload when the performance of the EBS volumes is too low. A SysOps administrator must implement a solution to ensure that the EBS volumes provide the expected performance when they are restored from snapshots.

Which solution will meet these requirements?

- A.** Configure fast snapshot restore (FSR) on the snapshots that are used.
- B.** Restore each snapshot onto an unencrypted EBS volume. Encrypt the EBS volume when the performance stabilizes.
- C.** Format the EBS volumes as XFS file systems before restoring the snapshots.
- D.** Increase the Linux read-ahead buffer to 1 MiB.

Answer: A

NO.8 A company has implemented a data ingestion pipeline to process files in the form of messages. A frontend application accepts user input and stores the input in Amazon S3. A backend application uses Amazon EC2 instances to process the object that was uploaded to Amazon S3.

The company recently experienced a significant increase in customer traffic. The frontend application is now sending more messages at one time than the backend application can handle, resulting in some lost messages.

Which action will resolve this problem with the LEAST operational effort?

- A.** Redevelop the backend application as a series of AWS Lambda functions.
- B.** Implement an Amazon Kinesis data stream to replace the backend application.

C. Implement an Application Load Balancer to distribute message traffic across the backend application instances.

D. Implement an Amazon Simple Queue Service (Amazon SQS) queue between the frontend and backend components.

Answer: D

Explanation:

Implementing an Amazon SQS queue between the frontend and backend decouples the processing pipeline. The queue can buffer messages, allowing the backend to process them at its own pace, and preventing message loss during traffic spikes. This solution requires minimal operational effort and can be integrated with the existing architecture without a major redesign.

NO.9 A company needs to implement disaster recovery (DR) for its application. The application's production environment consists of Amazon EC2 instances in an Auto Scaling group and an Amazon RDS database that has three read replicas.

The DR environment uses a single EC2 instance in an Auto Scaling group and one cross-Region read replica. The company uses Amazon Route 53 for DNS resolution.

A SysOps administrator must automate the DR process. However, the SysOps administrator must maintain manual control over the final failover step of updating DNS records.

Which solution will meet these requirements with the LEAST downtime?

A. Create an AWS CloudFormation template that increases the Auto Scaling group's EC2 instance count. Create an AWS Systems Manager Run Command document that adds three RDS read replicas. Update the Route 53 DNS record.

B. Create an AWS CloudFormation template that deploys a new Auto Scaling group that contains the appropriate number of EC2 instances. Use the CloudFormation template to deploy a new RDS DB instance and the required read replicas. When the DR environment is ready to support traffic, send a command to update the Route 53 DNS record.

C. Create an AWS Systems Manager Automation runbook that increases the Auto Scaling group's EC2 instance count, adds three RDS read replicas, and fails over the database. When the DR environment is ready to support traffic, send a command to update the Route 53 DNS record.

D. Create an AWS Systems Manager Run Command document that increases the Auto Scaling group's EC2 instance count, adds three RDS read replicas, and fails over the database. Update the Route 53 DNS record.

Answer: C

Explanation:

An AWS Systems Manager Automation runbook can orchestrate the multi-step DR actions - scale the ASG, add replicas, and promote/fail over the DB - quickly and repeatably. You keep manual control by triggering the final Route 53 update only after the runbook finishes preparing the DR site. This yields minimal downtime with controlled DNS cutover.

NO.10 A SysOps administrator notices that the cache hit ratio for an Amazon CloudFront distribution is less than 10%. The SysOps administrator needs to increase the cache hit ratio for the distribution, improve network performance, and reduce the load on the origin.

Which combination of actions should the SysOps administrator take to meet these requirements? (Choose two.)

A. Enable CloudFront Origin Shield for the required AWS Regions.

- B. Change the viewer protocol policy to use HTTPS only.
- C. Add a second origin. Create an origin group that includes both origins. Activate CloudFront origin failover.
- D. Turn on automatic compression of objects in the cache behavior settings.
- E. Increase the CloudFront TTL values in the cache behavior settings.

Answer: AE

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/origin-shield.html>

NO.11 While setting up an AWS managed VPN connection, a SysOps administrator creates a customer gateway resource in AWS. The customer gateway device resides in a data center with a NAT gateway in front of it.

What address should be used to create the customer gateway resource?

- A. The private IP address of the customer gateway device
- B. The MAC address of the NAT device in front of the customer gateway device
- C. The public IP address of the customer gateway device
- D. The public IP address of the NAT device in front of the customer gateway device

Answer: D

Explanation:

If your customer gateway device is behind a network address translation (NAT) device, use the IP address of your NAT device.

<https://docs.aws.amazon.com/vpn/latest/s2svpn/cgw-options.html>

NO.12 A SysOps administrator wants to securely share an object from a private Amazon S3 bucket with a group of users who do not have an AWS account.

What is the MOST operationally efficient solution that will meet this requirement?

- A. Attach an S3 bucket policy that only allows object downloads from the users' IP addresses.
- B. Create an IAM role that has access to the object. Instruct the users to assume the role.
- C. Create an IAM user that has access to the object. Share the credentials with the users.
- D. Generate a presigned URL for the object. Share the URL with the users.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.html>

NO.13 A company manages a set of accounts on AWS by using AWS Organizations. The company's security team wants to use a native AWS service to regularly scan all AWS accounts against the Center for Internet Security (CIS) AWS Foundations Benchmark.

What is the MOST operationally efficient way to meet these requirements?

- A. Designate a central security account as the AWS Security Hub administrator account. Create a script that sends an invitation from the Security Hub administrator account and accepts the invitation from the member account. Run the script every time a new account is created. Configure Security Hub to run the CIS AWS Foundations Benchmark scans.
- B. Run the CIS AWS Foundations Benchmark across all accounts by using Amazon Inspector.
- C. Designate a central security account as the Amazon GuardDuty administrator account. Create a

script that sends an invitation from the GuardDuty administrator account and accepts the invitation from the member account. Run the script every time a new account is created.

Configure GuardDuty to run the CIS AWS Foundations Benchmark scans.

D. Designate an AWS Security Hub administrator account. Configure new accounts in the organization to automatically become member accounts. Enable CIS AWS Foundations Benchmark scans.

Answer: D

NO.14 A data analytics application is running on an Amazon EC2 instance. A SysOps administrator must add custom dimensions to the metrics collected by the Amazon CloudWatch agent.

How can the SysOps administrator meet this requirement?

A. Create a custom shell script to extract the dimensions and collect the metrics using the Amazon CloudWatch agent.

B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to evaluate the required custom dimensions and send the metrics to Amazon Simple Notification Service (Amazon SNS).

C. Create an AWS Lambda function to collect the metrics from AWS CloudTrail and send the metrics to an Amazon CloudWatch Logs group.

D. Create an `append_dimensions` field in the Amazon CloudWatch agent configuration file to collect the metrics.

Answer: D

Explanation:

In custom metrics, the `--dimensions` parameter is common. A dimension further clarifies what the metric is and what data it stores. You can have up to 30 dimensions assigned to one metric, and each dimension is defined by a name and value pair.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

NO.15 A company runs hundreds of Amazon EC2 instances in a single AWS Region. Each EC2 instance has two attached 1 GiB General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volumes. A critical workload is using all the available IOPS capacity on the EBS volumes.

According to company policy, the company cannot change instance types or EBS volume types without completing lengthy acceptance tests to validate that the company's applications will function properly. A SysOps administrator needs to increase the I/O performance of the EBS volumes as quickly as possible.

Which action should the SysOps administrator take to meet these requirements?

A. Increase the size of the 1 GiB EBS volumes.

B. Add two additional elastic network interfaces on each EC2 instance.

C. Turn on Transfer Acceleration on the EBS volumes in the Region.

D. Add all the EC2 instances to a cluster placement group.

Answer: A

Explanation:

With Amazon EBS Elastic Volumes, you can increase the volume size, change the volume type, or adjust the performance of your EBS volumes. If your instance supports Elastic Volumes, you can do so without detaching the volume or restarting the instance. This enables you to continue using your application while the changes take effect.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modify-volume.html> They have Elastic

Volumes in place (per the question) and that's exactly why it is specified in the question. As others have mentioned, increasing the volume size increases IOPS, up to the volume type max. For gp2, you can have a volume size of 1 GiB - 16 TiB with a max IOPS of 16,000 for the 16 TiB volume size.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html>

NO.16 A company analyzes sales data for its customers. Customers upload files to one of the company's Amazon S3 buckets, and a message is posted to an Amazon Simple Queue Service (Amazon SQS) queue that contains the object Amazon Resource Name (ARN). An application that runs on an Amazon EC2 instance polls the queue and processes the messages. The processing time depends on the size of the file.

Customers are reporting delays in the processing of their files. A SysOps administrator decides to configure Amazon EC2 Auto Scaling as the first step. The SysOps administrator creates an Amazon Machine Image (AMI) that is based on the existing EC2 instance. The SysOps administrator also creates a launch template that references the AMI.

How should the SysOps administrator configure the Auto Scaling policy to improve the response time?

A. Add several different instance sizes in the launch template.

Create an Auto Scaling policy based on the `ApproximateNumberOfMessagesVisible` metric to select the size of the instance based on the number of messages in the queue.

B. Create an Auto Scaling policy based on the `ApproximateNumberOfMessagesDelayed` metric to scale the number of instances based on the number of messages in the queue that have been delayed.

C. Create a custom metric based on the `ASGAverageCPUUtilization` metric and the `GroupPendingInstances` metric from the Auto Scaling group.

Modify the application to calculate the metric and post the metric to Amazon CloudWatch once each minute.

Create an Auto Scaling policy based on this metric to scale the number of instances.

D. Create a custom metric based on the `ApproximateNumberOfMessagesVisible` metric and the number of instances in the `InService` state in the Auto Scaling group.

Modify the application to calculate the metric and post the metric to Amazon CloudWatch once each minute.

Create an Auto Scaling policy based on this metric to scale the number of instances.

Answer: D

Explanation:

When there are delays in processing files due to a high volume of messages in the queue, adding more instances using Auto Scaling can help to reduce the processing time. The `ApproximateNumberOfMessagesVisible` metric is a good indicator of the workload on the EC2 instances. By creating an Auto Scaling policy based on this metric, the number of instances can be scaled up or down depending on the number of messages in the queue.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-target-tracking-metric-math.html#metric-math-sqs-queue-backlog>

NO.17 A company plans to deploy a database on an Amazon Aurora MySQL DB cluster. The database will store data for a demonstration environment. The data must be reset on a daily basis. What is the MOST operationally efficient solution that meets these requirements?

- A.** Create a manual snapshot of the DB cluster after the data has been populated. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the snapshot and then delete the previous DB cluster.
- B.** Enable the Backtrack feature during the creation of the DB cluster. Specify a target backtrack window of 48 hours. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to perform a backtrack operation.
- C.** Export a manual snapshot of the DB cluster to an Amazon S3 bucket after the data has been populated. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the snapshot from Amazon S3.
- D.** Set the DB cluster backup retention period to 2 days. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the DB cluster to a point in time and then delete the previous DB cluster.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>

NO.18 A company has a new requirement stating that all resources in AWS must be tagged according to a set policy.

Which AWS service should be used to enforce and continually identify all resources that are not in compliance with the policy?

- A.** AWS CloudTrail
- B.** Amazon Inspector
- C.** AWS Config
- D.** AWS Systems Manager

Answer: C

Explanation:

<https://aws.amazon.com/config>

NO.19 A company is using an Amazon Aurora MySQL DB cluster that has point-in-time recovery, backtracking, and automatic backup enabled.

A SysOps administrator needs to be able to roll back the DB cluster to a specific recovery point within the previous 72 hours. Restores must be completed in the same production DB cluster.

Which solution will meet these requirements?

- A.** Create an Aurora Replica. Promote the replica to replace the primary DB instance.
- B.** Create an AWS Lambda function to restore an automatic backup to the existing DB cluster.
- C.** Use backtracking to rewind the existing DB cluster to the desired recovery point.
- D.** Use point-in-time recovery to restore the existing DB cluster to the desired recovery point.

Answer: C

Explanation:

Resolution

Note: If you receive errors when running AWS Command Line Interface (AWS CLI) commands, [make sure that you're using the most recent version of the AWS CLI](#).

Amazon Aurora backs-up your cluster volume's changes automatically and continuously. The back-ups are retained for the length of your [backup retention period](#). This continuous backup also means that you are able to restore your data to a new cluster, to any point in time within the retention period specified. This avoids the need for a lengthy binlog roll-forward process. Because you create a new cluster, there is no impact to performance or interruption to your original database.

When you initiate a clone, snapshot, or point in time restore, Amazon RDS calls the following APIs on your behalf:

- Either [RestoreDBClusterFromSnapshot](#) or [RestoreDBClusterToPointInTime](#). This creates a new cluster and restores volume from Amazon Simple Storage Service (Amazon S3). This can take up to two hours to complete. This is because when you restore data to an Aurora cluster, all of the data must be brought in parallel from Amazon S3 to the six copies on your three AZs.
- [Cluster storage volume cloning](#) is a variation of [RestoreDBClusterToPointInTime](#). It uses the copy-on-write protocol, and usually completes in a few minutes.

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-mysql-slow-snapshot-restore/>
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>

NO.20 A company is supporting a business-critical application that runs on Amazon EC2 instances. The application receives data from a service that runs in an on-premises data center. End users are reporting intermittent issues that are related to data refreshes. The issues are occurring because of fluctuations in available network bandwidth between AWS and the on-premises data center. A SysOps administrator must improve the user experience and the application's performance while minimizing changes to the application stack.

Which solution will offer the MOST performance improvement while meeting these requirements?

- A.** Migrate the service to AWS Implement auto scaling.
- B.** Modify the service to use Amazon S3 Transfer Acceleration.
- C.** Set up an AWS Direct Connect connection with the on-premises data center.
- D.** Use AWS Storage Gateway to move the data into AWS.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html>

NO.21 A SysOps administrator is troubleshooting a VPC with public and private subnets that leverage custom network ACLs. Instances in the private subnet are unable to access the internet. There is an internet gateway attached to the public subnet. The private subnet has a route to a NAT gateway that is also attached to the public subnet. The Amazon EC2 instances are associated with the default security group for the VPC.

What is causing the issue in this scenario?

- A.** There is a network ACL on the private subnet set to deny all outbound traffic.

- B. There is no NAT gateway deployed in the private subnet of the VPC.
- C. The default security group for the VPC blocks all inbound traffic to the EC2 instances.
- D. The default security group for the VPC blocks all outbound traffic from the EC2 instances.

Answer: A

Explanation:

Network ACLs (Access Control Lists) are stateless and operate at the subnet level. If there is a network ACL on the private subnet that is configured to deny all outbound traffic, it would prevent instances in the private subnet from accessing the internet through the NAT gateway.

NO.22 A company has a web application with a database tier that consists of an Amazon EC2 instance that runs MySQL. A SysOps administrator needs to minimize potential data loss and the time that is required to recover in the event of a database failure.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon CloudWatch alarm for the StatusCheckFailed_System metric to invoke an AWS Lambda function that stops and starts the EC2 instance.
- B. Create an Amazon RDS for MySQL Multi-AZ DB instance. Use a MySQL native backup that is stored in Amazon S3 to restore the data to the new database. Update the connection string in the web application.
- C. Create an Amazon RDS for MySQL Single-AZ DB instance with a read replica. Use a MySQL native backup that is stored in Amazon S3 to restore the data to the new database. Update the connection string in the web application.
- D. Use Amazon Data Lifecycle Manager (Amazon DLM) to take a snapshot of the Amazon Elastic Block Store (Amazon EBS) volume every hour. In the event of an EC2 instance failure, restore the EBS volume from a snapshot.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/snapshot-lifecycle.html>

NO.23 An application accesses databases that run on an Amazon Aurora PostgreSQL Multi-AZ DB cluster. The application is gaining more users and is experiencing an increased load. A SysOps administrator needs to improve the application performance by pooling and sharing database user connections.

Which solution will meet this requirement?

- A. Increase the IOPS of the DB cluster.
- B. Use Amazon RDS Proxy to set up a proxy. Associate the proxy with the DB cluster.
- C. Enable Enhanced Monitoring on the DB cluster. Move the logs to Amazon CloudWatch.
- D. Enable Performance Insights for 35 days on the DB cluster.

Answer: B

Explanation:

Amazon RDS Proxy pools and reuses database connections for Aurora PostgreSQL, reducing overhead from frequent opens/closes and letting many application clients share fewer backend connections - directly improving performance under increased load.

NO.24 A company has a memory-intensive application that runs on a fleet of Amazon EC2 instances

behind an Elastic Load Balancer (ELB). The instances run in an Auto Scaling group. A SysOps administrator must ensure that the application can scale based on the number of users that connect to the application.

Which solution will meet these requirements?

- A.** Create a scaling policy that will scale the application based on the ActiveConnectionCount Amazon CloudWatch metric that is generated from the ELB.
- B.** Create a scaling policy that will scale the application based on the mem_used Amazon CloudWatch metric that is generated from the ELB.
- C.** Create a scheduled scaling policy to increase the number of EC2 instances in the Auto Scaling group to support additional connections.
- D.** Create and deploy a script on the ELB to expose the number of connected users as a custom Amazon CloudWatch metric. Create a scaling policy that uses the metric.

Answer: A

NO.25 A company is running a development application on an Amazon EC2 instance. The application uploads 500,000 files that are 1 GB in size into a target Amazon S3 bucket that has default encryption enabled. The EC2 instance is in the same AWS Region where the S3 bucket is deployed.

The company uses performance logging that is built into the application software. The logs show that the application is constantly waiting for the files to be written to the S3 bucket. A SysOps administrator needs to improve the application's throughput performance. The SysOps administrator validates that the networking on the EC2 instance is not constrained.

What should the SysOps administrator do to improve the S3 upload performance?

- A.** Enable S3 Transfer Acceleration on the S3 bucket.
- B.** Split the S3 write operations to use multiple bucket prefixes to write items in parallel.
- C.** Configure AWS PrivateLink for Amazon S3. Turn off encryption on the S3 bucket.
- D.** Configure AWS Global Accelerator in the Region. Turn off encryption on the S3 bucket.

Answer: B

NO.26 A company has implemented a Kubernetes cluster on Amazon Elastic Kubernetes Service (Amazon EKS) to host a microservices-based application. The company expects application traffic to increase significantly for the next month and wants to prevent the application from crashing because of the high number of requests.

Which solution will meet these requirements with the LEAST administrative overhead?

- A.** Create a second EKS cluster. Load balance the workload between the two clusters.
- B.** Implement the Kubernetes Horizontal Pod Autoscaler. Set a target CPU utilization percentage.
- C.** Migrate the application from Amazon EKS to Amazon EC2 for the next month. Migrate the application back to Amazon EKS when the month ends.
- D.** Implement the Kubernetes Vertical Pod Autoscaler. Set a target CPU utilization percentage.

Answer: B

Explanation:

The Kubernetes Horizontal Pod Autoscaler (HPA) is designed to automatically scale the number of pods in a deployment or replica set based on observed CPU or memory utilization. In this scenario, the company wants to prevent the application from crashing due to high request traffic.

The HPA can dynamically adjust the number of pods based on CPU utilization, ensuring that the

application can handle increased traffic while avoiding overloading the system.

NO.27 A company is using an AWS KMS customer master key (CMK) with imported key material. The company references the CMK by its alias in the Java application to encrypt data. The CMK must be rotated every 6 months.

What is the process to rotate the key?

- A.** Enable automatic key rotation for the CMK, and specify a period of 6 months.
- B.** Create a new CMK with new imported material, and update the key alias to point to the new CMK.
- C.** Delete the current key material, and import new material into the existing CMK.
- D.** Import a copy of the existing key material into a new CMK as a backup, and set the rotation schedule for 6 months.

Answer: B

Explanation:

If you choose to import keys to AWS KMS or asymmetric keys or use a custom key store, you can manually rotate them by creating a new KMS key and mapping an existing key alias from the old KMS key to the new KMS key.

<https://aws.amazon.com/kms/faqs/>

NO.28 A SysOps administrator needs to collect the content of log files from a custom application that is deployed across hundreds of Amazon EC2 instances running Ubuntu. The log files need to be stored in Amazon CloudWatch Logs.

How should the SysOps administrator collect the application log files with the LOWEST operational overhead?

- A.** Configure the syslogd service on each EC2 instance to collect and send the application log files to CloudWatch Logs.
- B.** Install the CloudWatch agent by using the Amazon Linux package manager on each EC2 instance. Configure each agent to collect the application log files.
- C.** Install the CloudWatch agent on each EC2 instance by using AWS Systems Manager. Create an agent configuration on each instance by using the CloudWatch configuration wizard. Configure each agent to collect the application log files.
- D.** Store a CloudWatch agent configuration in the AWS Systems Manager Parameter Store. Install the CloudWatch agent on each EC2 instance by using Systems Manager. Configure each agent to collect the application log files.

Answer: D

Explanation:

The most operationally efficient method is to centralize the CloudWatch agent configuration so that it can be deployed and managed uniformly across all instances. By storing the CloudWatch agent configuration in AWS Systems Manager Parameter Store, you can centrally manage and update the configuration for the agents. Then, using AWS Systems Manager, you can install the CloudWatch agent on each EC2 instance (even though they are running Ubuntu) without manually configuring each one.

NO.29 A company has a public website that recently experienced problems. Some links led to missing webpages, and other links rendered incorrect webpages. The application infrastructure was running properly, and all the provisioned resources were healthy. Application logs and dashboards

did not show any errors, and no monitoring alarms were raised. Systems administrators were not aware of any problems until end users reported the issues.

The company needs to proactively monitor the website for such issues in the future and must implement a solution as soon as possible.

Which solution will meet these requirements with the LEAST operational overhead?

A. Rewrite the application to surface a custom error to the application log when issues occur.

Automatically parse logs for errors.

Create an Amazon CloudWatch alarm to provide alerts when issues are detected.

B. Create an AWS Lambda function to test the website.

Configure the Lambda function to emit an Amazon CloudWatch custom metric when errors are detected.

Configure a CloudWatch alarm to provide alerts when issues are detected.

C. Create an Amazon CloudWatch Synthetics canary.

Use the CloudWatch Synthetics Recorder plugin to generate the script for the canary run.

Configure the canary in line with requirements.

Create an alarm to provide alerts when issues are detected.

D. In the Amazon CloudWatch console, turn on Application Insights.

Create a CloudWatch alarm to provide alerts when an issue is detected.

Answer: C

Explanation:

Canaries are scripts written in Node.js or Python.

They create Lambda functions in your account that use Node.js or Python as a framework.

Canaries work over both HTTP and HTTPS protocols, which makes it possible for you to continually verify your customer experience even when you don't have any customer traffic on your applications. By using canaries, you can discover issues before your customers do.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Synthetics_Canaries.html

NO.30 A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company has configured an Amazon CloudWatch alarm to monitor the HTTPCode_Target_5XX_Count metric. The application crashes every few days during business hours. The crashes trigger the CloudWatch alarm and result in service disruption.

The cause of the crashes is a memory leak in the application. While developers work to fix the problem, a SysOps administrator needs to implement a temporary solution. The solution must automatically reboot the EC2 instances every day and must minimize application disruption during business hours.

Which solution will meet these requirements?

A. Create an Amazon EventBridge rule that is scheduled to run outside of business hours. Configure the rule to invoke the StartInstances operation on the EC2 instances.

B. Use AWS Systems Manager to create a daily maintenance window that is outside of business hours. Register the EC2 instances as a target. Assign the AWS-RestartEC2Instance runbook to the maintenance window.

C. Configure an additional CloudWatch alarm to monitor the StatusCheckFailed_System metric for the EC2 instances. Configure an EC2 action on the additional alarm to reboot the instances.

D. Configure an additional CloudWatch alarm that is triggered every time the application

crashes. Configure an EC2 action on the additional alarm to restart the application on the EC2 instances.

Answer: B

Explanation:

Using AWS Systems Manager to create a daily maintenance window outside of business hours allows you to schedule a reboot of the EC2 instances with minimal disruption. By registering the instances as targets and assigning the AWS-RestartEC2Instance runbook, the instances will be automatically rebooted during off-peak hours, which helps mitigate the memory leak issue temporarily while preserving application availability during business hours.

NO.31 A company has developed a service that is deployed on a fleet of Linux-based Amazon EC2 instances that are in an Auto Scaling group. The service occasionally fails unexpectedly because of an error in the application code. The company's engineering team determines that resolving the underlying cause of the service failure could take several weeks.

A SysOps administrator needs to create a solution to automate recovery if the service crashes on any of the EC2 instances.

Which solutions will meet this requirement? (Choose two.)

- A.** Install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to monitor the service. Set the CloudWatch action to restart if the service health check fails.
- B.** Tag the EC2 instances. Create an AWS Lambda function that uses AWS Systems Manager Session Manager to log in to the tagged EC2 instances and restart the service. Schedule the Lambda function to run every 5 minutes.
- C.** Tag the EC2 instances. Use AWS Systems Manager State Manager to create an association that uses the AWS-RunShellScript document. Configure the association command with a script that checks if the service is running and that starts the service if the service is not running. For targets, specify the EC2 instance tag. Schedule the association to run every 5 minutes.
- D.** Update the EC2 user data that is specified in the Auto Scaling group's launch template to include a script that runs on a cron schedule every 5 minutes. Configure the script to check if the service is running and to start the service if the service is not running. Redeploy all the EC2 instances in the Auto Scaling group with the updated launch template.
- E.** Update the EC2 user data that is specified in the Auto Scaling group's launch template to ensure that the service runs during startup. Redeploy all the EC2 instances in the Auto Scaling group with the updated launch template.

Answer: AE

NO.32 A company's VPC has connectivity to an on-premises data center through an AWS Site-to-Site VPN. The company needs Amazon EC2 instances in the VPC to send DNS queries for example.com to the DNS servers in the data center.

Which solution will meet these requirements?

- A.** Create an Amazon Route 53 Resolver inbound endpoint. Create a conditional forwarding rule on the on-premises DNS servers to forward DNS requests for example.com to the inbound endpoints.
- B.** Create an Amazon Route 53 Resolver inbound endpoint. Create a forwarding rule on the resolver that sends all queries for example.com to the on-premises DNS servers. Associate this rule with the VPC.
- C.** Create an Amazon Route 53 Resolver outbound endpoint. Create a conditional forwarding rule on

the on-premises DNS servers to forward DNS requests for example.com to the outbound endpoints.

D. Create an Amazon Route 53 Resolver outbound endpoint. Create a forwarding rule on the resolver that sends all queries for example.com to the on-premises DNS servers. Associate this rule with the VPC.

Answer: D

Explanation:

To allow EC2 instances in the VPC to resolve DNS queries using on-premises DNS servers over an AWS Site-to-Site VPN, you need to configure an Amazon Route 53 Resolver outbound endpoint. This enables DNS queries to be forwarded from AWS to external DNS servers, such as those in an on-premises data center.

1. Create a Route 53 Resolver outbound endpoint ?This allows the VPC to send DNS queries to on-premises DNS servers.
2. Configure a forwarding rule ?The rule ensures that all queries for example.com are directed to the on-premises DNS servers.
3. Associate the rule with the VPC ?This ensures that the EC2 instances in the VPC use the resolver for DNS resolution.

NO.33 A company has an existing public web application for www.example.com. The Application Load Balancer (ALB) is configured with a single HTTP 80 listener. A SysOps administrator must ensure that all web requests to www.example.com are encrypted between the client and the ALB.

The SysOps administrator already has requested and validated a public certificate for www.example.com in AWS Certificate Manager (ACM). Existing users of the application must not be required to change the endpoint to which they are connecting.

Which additional set of steps should the SysOps administrator take to meet these requirements?

A. Create an additional ALB listener for HTTPS on port 443. Set the default action to forward all traffic to the target group. Specify the ACM certificate that was created for www.example.com as the default SSL certificate.

B. Create an additional ALB listener for HTTPS on port 443. Set the default action to forward all traffic to the target group. Specify the ACM certificate that was created for www.example.com as the default SSL certificate. Delete the original HTTP listener on port 80.

C. Modify the ALB default rule for the HTTP port 80 listener. Create a rule in the listener to forward all traffic for the host www.example.com to the target group. Specify the ACM certificate that was created for www.example.com as the default SSL certificate.

D. Modify the ALB default rule for the HTTP port 80 listener to redirect to HTTPS on port 443. Create an additional HTTPS listener on port 443. Set the default action to forward all traffic to the target group. Specify the ACM certificate that was created for www.example.com as the default SSL certificate.

Answer: D

NO.34 A company updates its security policy to prohibit the public exposure of any data in Amazon S3 buckets in the company's account.

What should a SysOps administrator do to meet this requirement?

A. Turn on S3 Block Public Access from the account level.

B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to enforce that all S3 objects are private.

C. Use Amazon Inspector to search for S3 buckets and to automatically reset S3 ACLs if any public S3 buckets are found.

D. Use S3 Object Lambda to examine S3 ACLs and to change any public S3 ACLs to private.

Answer: A

Explanation:

Using Amazon S3 Block Public Access as a centralized way to limit public access. Block Public Access settings override bucket policies and object permissions. Be sure to enable Block Public Access for all accounts and buckets that you don't want publicly accessible.

NO.35 A company is trying to connect two applications. One application runs in an on-premises data center that has a hostname of host1.onprem.private. The other application runs on an Amazon EC2 instance that has a hostname of host1.awscloud.private. An AWS Site-to-Site VPN connection is in place between the on-premises network and AWS.

The application that runs in the data center tries to connect to the application that runs on the EC2 instance, but DNS resolution fails. A SysOps administrator must implement DNS resolution between on-premises and AWS resources.

Which solution allows the on-premises application to resolve the EC2 instance hostname?

A. Set up an Amazon Route 53 inbound resolver endpoint with a forwarding rule for the onprem.private hosted zone. Associate the resolver with the VPC of the EC2 instance. Configure the on-premises DNS resolver to forward onprem.private DNS queries to the inbound resolver endpoint.

B. Set up an Amazon Route 53 inbound resolver endpoint. Associate the resolver with the VPC of the EC2 instance. Configure the on-premises DNS resolver to forward awscloud.private DNS queries to the inbound resolver endpoint.

C. Set up an Amazon Route 53 outbound resolver endpoint with a forwarding rule for the onprem.private hosted zone. Associate the resolver with the AWS Region of the EC2 instance. Configure the on-premises DNS resolver to forward onprem.private DNS queries to the outbound resolver endpoint.

D. Set up an Amazon Route 53 outbound resolver endpoint. Associate the resolver with the AWS Region of the EC2 instance. Configure the on-premises DNS resolver to forward awscloud.private DNS queries to the outbound resolver endpoint.

Answer: B

Explanation:

Route 53 resolver provides resolution for AWS resources and on-prem dns NS provides resolution for on-prem resources. When DNS NS gets a dns query for AWS resources, it forwards it to Route 53 resolver.