

1) An enterprise has a large number of AWS accounts owned by separate business groups. One of the accounts was recently compromised. The attacker launched a large number of instances, resulting in a high bill for that account.

The security breach was addressed, but management has asked a solutions architect to develop a solution to prevent excessive spending in all accounts. Each business group wants to retain full control over its AWS account.

Which solution should the solutions architect recommend to meet these requirements?

- A) Use AWS Organizations to add each AWS account to the master account. Create a service control policy (SCP) that uses the `ec2:instanceType` condition key to prevent the launch of high-cost instance types in each account.
- B) Attach a new customer-managed IAM policy to an IAM group in each account that uses the `ec2:instanceType` condition key to prevent the launch of high-cost instance types. Place all of the existing IAM users in each group.
- C) Enable billing alerts on each AWS account. Create Amazon CloudWatch alarms that send an Amazon SNS notification to the account administrator whenever their account exceeds the spending budget.
- D) Enable Cost Explorer in each account. Regularly review the Cost Explorer reports for each account to ensure spending does not exceed the planned budget.

2) A company has multiple AWS accounts. The company has integrated its on-premises Active Directory with AWS SSO to grant Active Directory users least privilege abilities to manage infrastructure across all the accounts.

A solutions architect must integrate a third-party monitoring solution that requires read-only access across all AWS accounts. The monitoring solution will run in its own AWS account.

How can the monitoring solution be given the required permissions?

- A) Create a user in an AWS SSO directory and assign a read-only permissions set. Assign all AWS accounts to be monitored to the new user. Provide the third-party monitoring solution with the user name and password.
- B) Create an IAM role in the organization's master account. Allow the AWS account of the third-party monitoring solution to assume the role.
- C) Invite the AWS account of the third-party monitoring solution to join the organization. Enable all features.
- D) Create an AWS CloudFormation template that defines a new IAM role for the third-party monitoring solution with the account of the third party listed in the trust policy. Create the IAM role across all linked AWS accounts by using a stack set.

3) A team is building an HTML form hosted in a public Amazon S3 bucket. The form uses JavaScript to post data to an Amazon API Gateway endpoint. The endpoint is integrated with AWS Lambda functions. The team has tested each method in the API Gateway console and received valid responses.

Which combination of steps must be completed for the form to successfully post to the API Gateway and receive a valid response? (Select TWO.)

- A) Configure the S3 bucket to allow cross-origin resource sharing (CORS).
- B) Host the form on Amazon EC2 rather than Amazon S3.
- C) Request a limit increase for API Gateway.
- D) Enable cross-origin resource sharing (CORS) in API Gateway.
- E) Configure the S3 bucket for web hosting.

4) A retail company runs a serverless mobile app built on Amazon API Gateway, AWS Lambda, Amazon Cognito, and Amazon DynamoDB. During heavy holiday traffic spikes, the company receives complaints of intermittent system failures. Developers find that the API Gateway endpoint is returning 502 Bad Gateway errors to seemingly valid requests.

Which method should address this issue?

- A) Increase the concurrency limit for Lambda functions and configure notification alerts to be sent by Amazon CloudWatch when the `ConcurrentExecutions` metric approaches the limit.
- B) Configure notification alerts for the limit of transactions per second on the API Gateway endpoint and create a Lambda function that will increase this limit, as needed.
- C) Shard users to Amazon Cognito user pools in multiple AWS Regions to reduce user authentication latency.
- D) Use DynamoDB strongly consistent reads to ensure the latest data is always returned to the client application.

5) A web hosting company has enabled Amazon GuardDuty in every AWS Region for all of its accounts. A system administrator must create an automated response to high-severity events.

How should this be accomplished?

- A) Create rules through VPC Flow Logs that trigger an AWS Lambda function that programmatically addresses the issue.
- B) Create an Amazon CloudWatch Events rule that triggers an AWS Lambda function that programmatically addresses the issue.
- C) Configure AWS Trusted Advisor to trigger an AWS Lambda function that programmatically addresses the issue.
- D) Configure AWS CloudTrail to trigger an AWS Lambda function that programmatically addresses the issue.

6) A company is launching a new web service on an Amazon ECS cluster. Company policy requires that the security group on the cluster instances block all inbound traffic but HTTPS (port 443). The cluster consists of 100 Amazon EC2 instances. Security engineers are responsible for managing and updating the cluster instances. The security engineering team is small, so any management efforts must be minimized.

How can the service be designed to meet these operational requirements?

- A) Change the SSH port to 2222 on the cluster instances with a user data script. Log in to each instance using SSH over port 2222.
- B) Change the SSH port to 2222 on the cluster instances with a user data script. Use AWS Trusted Advisor to remotely manage the cluster instances over port 2222.
- C) Launch the cluster instances with no SSH key pairs. Use the Amazon Systems Manager Run Command to remotely manage the cluster instances.
- D) Launch the cluster instances with no SSH key pairs. Use AWS Trusted Advisor to remotely manage the cluster instances.

7) A company has two AWS accounts: one for production workloads and one for development workloads. Creating and managing these workloads are a development team and an operations team. The company needs a security strategy that meets the following requirements:

- **Developers need to create and delete development application infrastructure.**
- **Operators need to create and delete both development and production application infrastructure.**
- **Developers should have no access to production infrastructure.**
- **All users should have a single set of AWS credentials.**

What strategy meets these requirements?

- A) In the development account:
- Create a development IAM group with the ability to create and delete application infrastructure.
 - Create an IAM user for each operator and developer and assign them to the development group.
- In the production account:
- Create an operations IAM group with the ability to create and delete application infrastructure.
 - Create an IAM user for each operator and assign them to the operations group.
- B) In the development account:
- Create a development IAM group with the ability to create and delete application infrastructure.
 - Create an IAM user for each developer and assign them to the development group.
 - Create an IAM user for each operator and assign them to the development group and the operations group in the production account.
- In the production account:
- Create an operations IAM group with the ability to create and delete application infrastructure.
- C) In the development account:
- Create a shared IAM role with the ability to create and delete application infrastructure in the production account.
 - Create a development IAM group with the ability to create and delete application infrastructure.
 - Create an operations IAM group with the ability to assume the shared role.
 - Create an IAM user for each developer and assign them to the development group.
 - Create an IAM user for each operator and assign them to the development group and the operations group.
- D) In the development account:
- Create a development IAM group with the ability to create and delete application infrastructure.
 - Create an operations IAM group with the ability to assume the shared role in the production account.
 - Create an IAM user for each developer and assign them to the development group.
 - Create an IAM user for each operator and assign them to the development group and the operations group.
- In the production account:
- Create a shared IAM role with the ability to create and delete application infrastructure.
 - Add the development account to the trust policy for the shared role.

8) A company is migrating an Apache Hadoop cluster from its data center to AWS. The cluster consists of 60 VMware Linux virtual machines (VMs). During the migration cluster, downtime should be minimized.

Which process will minimize downtime?

- A) Use the AWS Management Portal for vCenter to migrate the VMs to AWS as Amazon EC2 instances.
- B) Use AWS Server Migration Service (AWS SMS) to migrate the VMs to AWS as AMIs. Launch the cluster on AWS as Amazon EC2 instances from the migrated AMIs.
- C) Create Open Virtualization Archive (OVA) files of the VMs. Upload the OVA files to Amazon S3. Use VM Import/Export to create AMIs from the OVA files. Launch the cluster on AWS as Amazon EC2 instances from the AMIs.
- D) Export the Hadoop Digital File System (HDFS) data from the VMs to a new Amazon Aurora DB cluster. Launch a new Hadoop cluster on Amazon EC2 instances. Import the data from the Aurora database to HDFS on the new cluster.

9) A solutions architect needs to reduce costs for a big data application. The application environment consists of hundreds of devices that send events to Amazon Kinesis Data Streams. The device ID is used as the partition key, so each device gets a separate shard. Each device sends between 50 KB and 450 KB of data per second. The shards are polled by an AWS Lambda function that processes the data and stores the result on Amazon S3.

Every hour, an AWS Lambda function runs an Amazon Athena query against the result data that identifies any outliers and places them in an Amazon SQS queue. An Amazon EC2 Auto Scaling group of two EC2 instances monitors the queue and runs a short (approximately 30-second) process to address the outliers. The devices submit an average of 10 outlying values every hour.

Which combination of changes to the application would MOST reduce costs? (Select TWO.)

- A) Change the Auto Scaling group launch configuration to use smaller instance types in the same instance family.
- B) Replace the Auto Scaling group with an AWS Lambda function triggered by messages arriving in the Amazon SQS queue.
- C) Reconfigure the devices and data stream to set a ratio of 10 devices to 1 data stream shard.
- D) Reconfigure the devices and data stream to set a ratio of 2 devices to 1 data stream shard.
- E) Change the desired capacity of the Auto Scaling group to a single EC2 instance.

10) A company operates an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. After an order is successfully processed, the application immediately posts order data to an external third-party affiliate tracking system that pays sales commissions for order referrals. During a highly successful marketing promotion, the number of EC2 instances increased from 2 to 20. The application continued to work correctly, but the increased request rate overwhelmed the third-party affiliate and resulted in failed requests.

Which combination of architectural changes could ensure that the entire process functions correctly under load? (Select TWO.)

- A) Move the code that calls the affiliate to a new AWS Lambda function. Modify the application to invoke the Lambda function asynchronously.
- B) Move the code that calls the affiliate to a new AWS Lambda function. Modify the application to place the order data in an Amazon SQS queue. Trigger the Lambda function from the queue.
- C) Increase the timeout of the new AWS Lambda function.
- D) Adjust the concurrency limit of the new AWS Lambda function.
- E) Increase the memory of the new AWS Lambda function.

Answers

- 1) C – [Billing alarms](#) will allow management to get alerted on excessive spend without taking control away from any of the business groups. A and B are incorrect because each business group wants to retain control of their account, and these solutions would not protect against launching a large number of instances. D is a manual process, and it could be a while before any unauthorized spend is discovered.
- 2) D – [AWS CloudFormation StackSets](#) can deploy the IAM role across multiple accounts with a single operation. A is incorrect because credentials supplied by AWS SSO are temporary, so the application would lose permissions and have to log in again. B would grant access to the master account only. C is incorrect because accounts belonging to an organization do not receive permissions in the other accounts.
- 3) D, E – [CORS must be enabled](#) to keep the browser from generating an error due to the same origin policy, which requires that the dynamic content should come from the same domain as the static content. Since API Gateway is using a domain of the form `[restapi-id].execute-api.amazonaws.com`, and the S3 bucket uses `[bucketname].s3.website-[region].amazonaws.com`, a CORS header must be sent with the API Gateway response for the browser to relax the restriction. E is required for the HTML form to be served using a [website endpoint](#). A is incorrect because the CORS header must be configured to be returned by the dynamic response from the API endpoint. Configuring CORS for the S3 bucket does not help. B is incorrect because there is no advantage to serving a static webpage from a web server running on EC2 versus an S3 bucket. C is incorrect because API Gateway has a [default per AWS Region limit](#) of 10,000 requests per second. If required for production, this limit can be increased.
- 4) A – [The 502 internal server errors](#) will be returned intermittently by API Gateway if the Lambda function exceeds concurrency limits. B is incorrect because, in this case, API Gateway would return a [429 error for too many requests](#). C is incorrect because the error occurs when calling the API Gateway endpoint, not during the authentication process. D is incorrect because stale data would not cause a bad gateway error.
- 5) B – GuardDuty findings can be sent to Amazon SNS topics and [CloudWatch Events](#). Neither VPC Flow Logs nor AWS CloudTrail can trigger a Lambda function. Trusted Advisor is a recommendation service, and is not suited for this scenario.
- 6) C – The [Systems Manager Run Command](#) requires no inbound ports to be open; it operates entirely over outbound HTTPS (which is open by default for security groups). A and B are ruled out because the requirements clearly state that the only inbound port to be open is 443. D is ruled out because Trusted Advisor does perform management functions.
- 7) D – This is the only response that will work and meets the requirements. It follows the [standard guidelines](#) for granting cross-account access between two accounts that you control. A requires two sets of credentials for operators, which breaks the requirements. B will not work, as an IAM user cannot be added to an IAM group in a different account. C will not work, as a role cannot grant access to resources in another account; the shared role must be in the account with resources it manages.
- 8) B – [AWS SMS](#) uploads each VM incrementally, so it can upload the servers while the data center cluster is still running. The data center cluster must be shut down prior to the final incremental sync of all the VMs only.
- 9) B, D – The average amount of compute used each hour is about 300 seconds (10 events x 30 seconds). While A and E would both reduce costs, they both involve paying for one or more EC2 instances sitting unused for 3,300 or more seconds per hour. B involves paying [for the small amount of compute time required](#) to process the outlying values only. Both C and D reduce the shard hour costs of the Kinesis data stream, but C will not work

because the amount of data would exceed the [1 MBps limit](#) of a single shard.

10) B, D – Putting the messages in a queue (B) will decouple the main application from calls to the affiliate. That will not only protect the main application from the reduced capacity of the affiliate, it will also allow failed requests to automatically go back to the queue. [Limiting number of concurrent executions](#) (D) will prevent overwhelming the affiliate application. A is incorrect because, while asynchronously invoking the Lambda function will reduce load on the EC2 instances, it will not lower the number of requests to the affiliate application. C is incorrect because, while it will allow the Lambda function to wait longer for the external call to return, it does not reduce the load on the affiliate application (which will still be overwhelmed). E is incorrect because adjusting the memory will have no effect on the interaction between the Lambda function and the affiliate application.