1) A corporate cloud security policy states that communication between the company's VPC and KMS must travel entirely within the AWS network and not use the public service endpoints.

What combination of the following actions satisfies that requirement? (Select TWO.)

A) Add the aws:sourceVpce condition to the KMS key policy referencing the company's Amazon VPC endpoint ID.
B) Remove the VPC internet gateway from the VPC and add a virtual private gateway to the VPC to prevent direct, public internet connectivity.
C) Create a VPC endpoint for AWS KMS with private DNS enabled.
D) Use the KMS Import Key feature to securely transfer the AWS KMS key over a VPN.
E) Add the following condition to the AWS KMS key policy: "aws:SourceIp": "10.0.0.0/16".

2) An application team is designing a solution with two applications. The security team wants the applications' logs to be captured in two different places, because one of the applications produces logs with sensitive data.

What solution meets the requirement with the LEAST risk and effort?

A) Use Amazon CloudWatch logs to capture all logs, write an AWS Lambda function that parses the log file, and move sensitive data to a different log.
B) Use Amazon CloudWatch logs with two log groups, one for each application, and use an AWS IAM policy to control access to the log groups as required.
C) Aggregate logs into one file, then use Amazon CloudWatch Logs, and then design two CloudWatch metric filters to filter sensitive data from the logs.
D) Add logic to the application that saves sensitive data logs on the Amazon EC2 instances' local storage, and write a batch script that logs into the EC2 instances and moves sensitive logs to a secure location.

3) A Security Engineer must set up security group rules for a three-tier application:

Presentation Tier - Accessed by users over the web, protected by the security group, presentation-sg
Logic Tier - RESTful API accessed from the Presentation Tier via https, protected by the security group, logic-sg
Data Tier - SQL Server database accessed over port 1433 from the Logic Tier, protected by the security group, data-sg

What combination of the following security group rules will allow the application to be secure and functional? (Select THREE.)

A) presentation-sg: Allow ports 80 and 443 from 0.0.0.0/0
B) data-sg: Allow port 1433 from presentation-sg
C) data-sg: Allow port 1433 from logic-sg
D) presentation-sg: Allow port 1433 from data-sg
E) logic-sg: Allow port 443 from presentation-sg
F) logic-sg: Allow port 443 from 0.0.0.0/0

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved | aws.amazon.com

4) A Security Engineer is working with a product team building a web application on AWS. The application uses Amazon S3 to host HTML pages and other static content, API Gateway and AWS Lambda to provide RESTful services, and Amazon DynamoDB to store state and other data. The users already exist in a directory that is exposed via a SAML identity provider (IDP).

What should the Security Engineer do to enable users to be authenticated into the web application and call APIs? (Select THREE).

A) Create a custom authorization service using AWS Lambda.
B) Configure a SAML Identity Provider in Amazon Cognito to map attributes to the Cognito User Pool attributes.
C) Configure the identity provider to add the Amazon Cognito User Pool as a relying party.
D) Configure an Amazon Cognito Identity Pool to integrate with social login providers.
E) Update DynamoDB to store the user email addresses and passwords.
F) Update API Gateway to use an Amazon Cognito User Pools authorizer.

5) An organization is hosting a web application on AWS and is using an S3 bucket to store images. Users should have the ability to read objects in the bucket. A Security Engineer has written the following bucket policy to grant public read access:

```
{
   "ID":"Policy1502987489630",
   "Version":"2012-10-17",
   "Statement":[
      {
         "Sid":"Stmt1502987487640",
         "Action":[
            "s3:GetObject",
            "s3:GetObjectVersion"
         ],
         "Effect":"Allow",
         "Resource":"arn:aws:s3:::appbucket",
         "Principal":"*"
      }
   ]
}
```

Attempts to read an object, however, receive the error: "Action does not apply to any resource(s) in statement.”

What should the Engineer do to fix the error?

A) Change the IAM permissions by applying PutBucketPolicy permissions.
B) Verify that the policy has the same name as the bucket name. If not, make it the same.
C) Change the Resource section to "arn:aws:s3:::appbucket/*".
D) Add an action s3:ListBucket.

6) The decision was made to place database hosts in their own VPC, and to set up VPC peering to different VPCs containing the application and web tiers. The application servers are unable to connect to the database.

Which network troubleshooting steps should be taken to resolve the issue? (Select TWO.)

A) Check to see if the application servers are in a private subnet or public subnet.
B) Check the route tables for the application server subnets for routes to the VPC peering connection.
C) Check the network access control lists for the database subnets for rules that allow traffic from the Internet.
D) Check the database security groups for rules that allow traffic from the application servers.
E) Check to see if the database VPC has an Internet gateway

7) When testing a new AWS Lambda function that retrieves items from an Amazon DynamoDB table, the Security Engineer noticed that the function was not logging any data to Amazon CloudWatch logs.

Below is the policy that was assigned to the role assumed by the Lambda function:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Dynamo-1234567",
      "Action": [
        "dynamodb:GetItem"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
}
```

What would be the least-privileged policy addition that would allow this function to log properly?

A)
```
{
    "Sid": "Logging-12345",
    "Resource": "*",
    "Action": [
      "logs:*"
    ],
    "Effect": "Allow"
}
```

B)
```
{
    "Sid": "Logging-12345",
    "Resource": "*",
    "Action": [
      "logs:CreateLogStream"
    ],
```

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved | aws.amazon.com

```
      "Effect": "Allow"
    }
C)  {
      "Sid": "Logging-12345",
      "Resource": "*",
      "Action": [
        "logs:CreateLogGroup",
        "logs:CreateLogStream",
        "logs:PutLogEvents"
      ],
      "Effect": "Allow"
    }
D)  {
      "Sid": "Logging-12345",
      "Resource": "*",
      "Action": [
        "logs:CreateLogGroup",
        "logs:CreateLogStream",
        "logs:DeleteLogGroup",
        "logs:DeleteLogStream",
        "logs:getLogEvents",
        "logs:PutLogEvents"
      ],
      "Effect": "Allow"
    }
```

8) A company is building a data lake on Amazon S3. The data consists of millions of small files containing sensitive information. The security team has the following requirements for the architecture:

- **Data must be encrypted in transit.**
- **Data must be encrypted at rest.**
- **The bucket must be private, but if the bucket is accidentally made public, the data must remain confidential.**

What combination of steps would meet the requirements? (Select TWO.)

A) Enable AES-256 (SSE-S3) encryption on the S3 bucket.
B) Enable default encryption with AWS KMS-managed keys (SSE-KMS) on the S3 bucket.
C) Add a bucket policy that includes a deny if PutObject request does not include aws:SecureTransport.
D) Add a bucket policy with aws:SourceIp to only allow uploads and downloads from the corporate intranet.
E) Enable Amazon Macie to monitor and act on changes to the data lake's S3 bucket.

9) A Security Engineer must ensure that all API calls are collected across all company accounts, and that they are preserved online and are instantly available for analysis for 90 days. For compliance reasons, this data must be restorable for 7 years.

Which steps must be taken to meet the retention needs in a scalable, cost-effective way?

A) Enable AWS CloudTrail logging across all accounts to a centralized Amazon S3 bucket with versioning enabled. Set a lifecycle policy to move the data to Amazon Glacier daily, and expire the data after 90 days.
B) Enable AWS CloudTrail logging across all accounts to S3 buckets. Set a lifecycle policy to expire the data in each bucket after 7 years.
C) Enable AWS CloudTrail logging across all accounts to Glacier. Set a lifecycle policy to expire the data after 7 years.
D) Enable AWS CloudTrail logging across all accounts to a centralized Amazon S3 bucket. Set a lifecycle policy to move the data to Glacier after 90 days, and expire the data after 7 years.

10) A Security Engineer has been informed that a user's access key has been found on GitHub. The Engineer must ensure that this access key cannot continue to be used, and must assess whether the access key was used to perform any unauthorized activities.

What steps must be taken to perform these tasks?

A) Review the user's IAM permissions and delete any unrecognized or unauthorized resources.
B) Delete the user, review the Amazon CloudWatch logs in all regions, and report the abuse.
C) Delete or rotate the user's key, review the CloudTrail logs in all regions, and delete any unrecognized or unauthorized resources.
D) Instruct the user to remove the key from the GitHub submission, rotate keys, and re-deploy any instances that were launched.

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved | aws.amazon.com

Answers

1) AC – An IAM policy can deny access to KMS except through your VPC endpoint with the following condition statement:

```
"Condition": {
    "StringNotEquals": {
        "aws:sourceVpce": "vpce-0295a3caf8414c94a"
    }
}
```

If you select the Enable Private DNS Name option, the standard AWS KMS DNS hostname (https://kms.<region>.amazonaws.com) resolves to your VPC endpoint.

2) B – Each application's log can be configured to send the log to a specific CloudWatch log group

3) ACE – In an n-tier architecture, each tier's security group allows traffic only from the security group sending it traffic. The presentation tier opens traffic for HTTP and HTTPS from the internet. Since security groups are stateful, only inbound rules are required.

4) BCF – When Cognito receives a SAML assertion it needs to be able to map SAML attributes to user pool attributes:. When configuring Cognito to receive SAML assertions from an identity provider you need ensure that the IDP is configured to have Cognito as a relying party. API Gateway will need to be able to understand the authorization being passed from Cognito which is a configuration step:

5) C – The resource section should match with the type of operation. Change the ARN to include /* at the end as it is an object operation. https://aws.amazon.com/blogs/security/writing-iam-policies-how-to-grant-access-to-an-amazon-s3-bucket/.

6) BD –You must configure the route tables in each VPC to route to each other through the peering connection. You also must add rules to the security group for the databases to accept requests from the application server security group in the other VPC.

7) C – Basic Lambda permissions required to log to CloudWatch Logs include: CreateLogGroup, CreateLogStream, and PutLogEvents.

8) BC – Bucket encryption using KMS will protect both in case disks are stolen as well as if the bucket is public. This is because the KMS key would need to have privileges granted to it for users outside of AWS.

HTTPS will protect data in transit.

9) D – Meets all requirements and is cost effective by using lifecycle polices to transition to glacier.

10) C – Removes keys and audits the environment for malicious activities.

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved | aws.amazon.com