AWS Security - Specialty (Released April 2018) (SCS-C01)
Sample Exam Questions

1) A Network Security Engineer is asked to implement security groups to allow for a secure public website. The development team included rules that allow the application to work correctly and be administrable.

Which of the following security group configurations are the MOST secure but still functional to support these requirements?

A. Port 80 coming from 0.0.0.0/0
Port 443 coming from 0.0.0.0/0
Port 22 coming from 0.0.0.0/0
Port 1433 coming from 0.0.0.0/0

B. Port 80 coming from 0.0.0.0/0
Port 443 coming from 0.0.0.0/0
Port 22 coming from 10.0.0.0/16
Port 1433 coming from 0.0.0.0/0

C. Port 80 coming from 10.0.0.0/16
Port 443 coming from 10.0.0.0/16
Port 22 coming from 10.0.0.0/16
Port 1433 coming from 10.0.0.0/16

D. Port 80 coming from 0.0.0.0/0
Port 443 coming from 0.0.0.0/0
Port 22 coming from 10.0.0.0/16
Port 1433 coming from 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.**

Which 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**

Which 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

4) A company has set up a new AWS account and must provide its large IT staff with permissions to access various AWS resources. The company already maintains user identities outside of AWS in its corporate user directory that supports SAML.

How can access be provided to the IT team members MOST efficiently?

A. Create a custom sign-in code with the secret access key to allow the corporate identity management system access to the new AWS account.
B. Use the corporate user directory as an IAM identity provider that will manage the user identities outside of AWS. Create a federation and give the external user identities permissions to use AWS resources.
C. Create new IAM users that the IT team will use for AWS account access. Map the new accounts to the existing corporate user directory, establishing an Active Directory relationship, and give the external user identities permissions to use AWS resources.
D. Create new IAM administrative accounts that the IT department members will share for accessing AWS resources, minimizing the number of user accounts to manage.

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 an IAM policy to grant public read access using the following bucket policy:**

```
{   "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. Create the bucket "appbucket" and then apply the policy.

6) **The decision was made to place database hosts in their own VPC, and to set up VPC peering to application and web servers that are hosted in different VPCs. 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 point.
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": "*"
          }
      }
```

Which of the following 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"
],
"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 requires that data stored in AWS be encrypted at rest.

Which of the following approaches achieve this requirement? (Select TWO.)

A. When storing data in Amazon EBS, use only EBS–optimized Amazon EC2 instances.
B. When storing data in EBS, encrypt the volume by using AWS KMS.
C. When storing data in Amazon S3, use object versioning and MFA Delete.
D. When storing data in Amazon EC2 Instance Store, encrypt the volume by using KMS.
E. When storing data in S3, enable server-side encryption.

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 CloudTrail logging in all accounts into S3 buckets, and set a lifecycle policy to expire the data in each bucket after 7 years.
C. Enable CloudTrail logging to Glacier, and set a lifecycle policy to expire the data after 7 years.
D. Enable CloudTrail logging to a centralized 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.

Answers

1) D – Port 22 should not be open to the public. Port 1433, leaves an unknown security risk open as well. https://aws.amazon.com/articles/1233/
2) B – Each application's log can be configured to send the log to a specific CloudWatch log group http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/CloudWatchLogsConcepts.html
3) ACE – Each of these are required and do not allow outside connection when not needed. Limited access for needs. https://d0.awsstatic.com/whitepapers/Security/AWS_Security_Best_Practices.pdf
4) B – With an identity provider, user identities can be managed outside of AWS and gives external user identities permissions to use AWS resources in an account. This is useful if an organization already has its own identity system, such as a corporate user directory. https://aws.amazon.com/blogs/security/enabling-federation-to-aws-using-windows-active-directory-adfs-and-saml-2-0/, http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_saml.html, and http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers.html
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 – http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-partial-access.html
https://aws.amazon.com/about-aws/whats-new/2016/03/announcing-support-for-security-group-references-in-a-peered-vpc/
http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-peering-security-groups.html
7) C – http://docs.aws.amazon.com/lambda/latest/dg/policy-templates.html
8) BE – EBS volume encryption with KMS provides encryption at rest. Amazon EBS encryption uses AWS Key Management Service (AWS KMS) customer master keys (CMK) when creating encrypted volumes and any snapshots created from them. The first time and encrypted volume is created in a region, a default CMK is created automatically. Using S3 with client-side encryption provides encryption both in transit and at rest. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html and http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html Client-side encryption refers to encrypting data before sending it to Amazon S3. The following two options exist for using data encryption keys: Use an AWS KMS-managed customer master key or use a client-side master key. http://docs.aws.amazon.com/AmazonS3/latest/dev/bucket-encryption.html
9) D – Meets all requirements and is cost effective by using glacier. http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html
10) C – Removes keys and audits the environment for malicious activities. https://aws.amazon.com/premiumsupport/knowledge-center/potential-account-compromise/