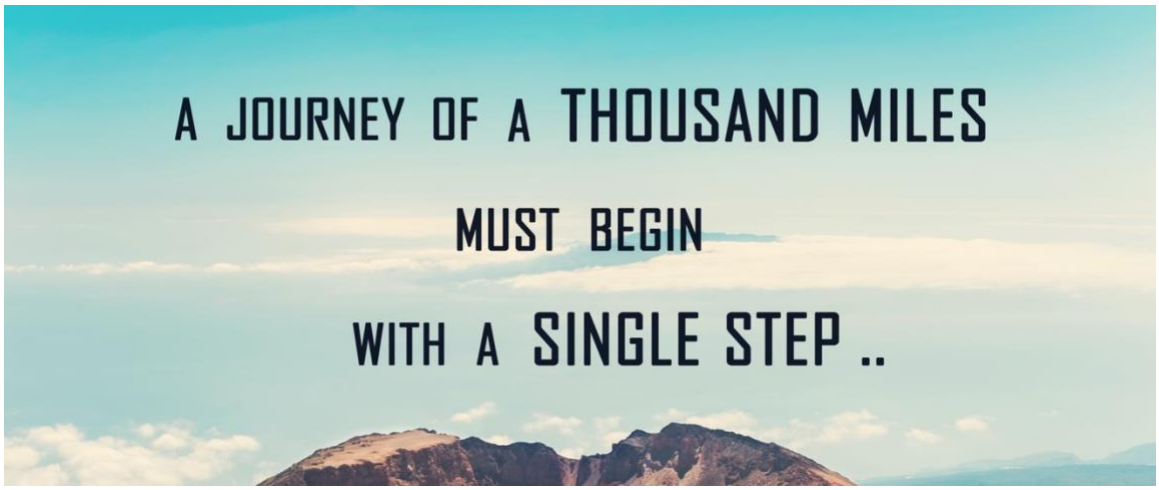

AWS Certified Cloud Practitioner Practice Test with Answers

Question 1:
Which of the following is not a pillar of AWS Well-Architected Framework?

A. Security
B. Reliability
C. Scalability
D. Sustainability

Answer: C

Explanation
Correct Option
Scalability: The AWS Well-Architected Framework is based on the following architectural pillars: Operational Excellence, Security, Reliability, Performance Efficiency, Cost Optimization, and Sustainability.

Reference:
https://aws.amazon.com/architecture/well-architected/

Incorrect Options

Security
Reliability
Sustainability

These are part of the AWS Well-Architected Framework.

Reference:
https://aws.amazon.com/architecture/well-architected/

Question 2:
Which of the following AWS services can perform multiple builds concurrently?

A. AWS CodeStar
B. AWS Code Build
C. Amazon CodeGuru
D. AWS CodeCommit

Answer: B

Explanation
Correct Option:
AWS Code Build : AWS CodeBuild, which is a fully managed continuous integration (CI) service, compiles source code, runs tests, and produces software packages ready to deploy. With AWS CodeBuild, you don't need to provision, manage, and scale your build servers. Instead, AWS CodeBuild scales and performs multiple builds concurrently. This helps in builds not left waiting in a queue.

Incorrect Options
AWS CodeStar: AWS CodeStar provides a unified interface to bring all AWS CI/CD-related services under one service.

It brings together the following services:

- AWS CodeCommit is essentially a managed git source code repository in AWS.
- AWS CodeBuild, which is like Jenkins. It builds the code and runs tests, and creates deployable artifacts.
- AWS CodeDeploy is an automated software deployment service to deploy your code on an EC2 instance or Elastic Beanstalk.
- AWS Code Pipeline checks out the code, builds it, tests it, and then deploys the code. It is a managed CI/CD pipeline.

Amazon CodeGuru: Amazon CodeGuru is a developer tool powered by machine learning that provides intelligent recommendations for improving code quality by identifying applications' most expensive lines of code. You can add Amazon CodeGuru to your workflow to identify expensive lines of code to reduce cost. When you add a pull request, Amazon CodeGuru will analyze code from the repository for the critical issues and provides a recommendation report. The report shows why the issue is flagged, the cost of incurring, and suggestive resolution steps.

AWS CodeCommit: AWS CodeCommit is a fully managed, secure source control service that hosts git-based repositories. You can create a git repository, add files, clone a repository, create a pull request, and merge pull requests to the branch. In other words, using AWS CodeCommit, you can do all sorts of git operations that you usually do as a developer.

Reference:
https://aws.amazon.com/codestar/
https://aws.amazon.com/codebuild/
https://aws.amazon.com/codeguru/
https://aws.amazon.com/codecommit/

Question 3:
Which of the following AWS services can you use to find trends related to your AWS cost and usage?

A. AWS CloudWatch Dashboard
B. AWS Cost Explorer
C. AWS Organizations
D. AWS Budgets

Answer: B

Explanation
Correct Option
AWS Cost Explorer: AWS Cost Explorer service lets you visualize, understand, and manage your AWS costs and usage over time. It has an easy-to-use interface. You can explore your usage and

costs using graphs or various Cost Explorer canned reports. For example, you can find your total AWS cost for the last 12 months and forecast for the current month. You can also get AWS costs for individual services, for example, how much you spend on EC2, S3, or any other AWS services.

The Cost Explorer also shows different trends, for example, if your current month's cost usage is down or up for a particular service. If, for the current month, you have uploaded more content on S3, then there is a potential for your S3 usage cost to go up. And it may show up in the trend. These trends can help you understand your charges and usage pattern to analyze them. The analysis may help you reduce AWS costs -- proactively.

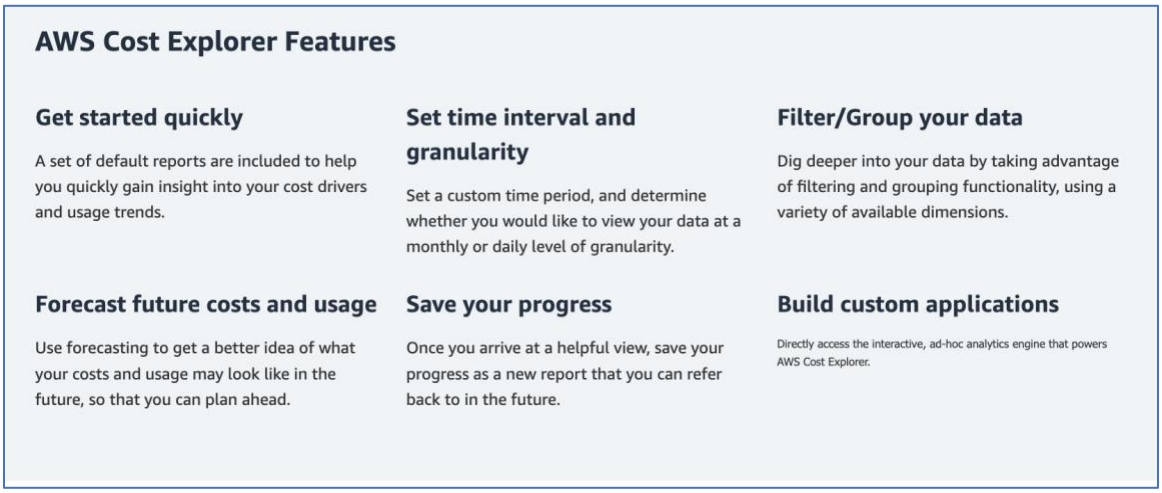

Screenshot from:
https://aws.amazon.com/aws-cost-management/aws-cost-explorer/

Incorrect Options
AWS CloudWatch Dashboard: This is one of the features of AWS CloudWatch. With this feature, you can create a custom dashboard based on your requirements for the kind of metrics you are interested in, particularly resources. For example, if you are interested in EC2 instances CPU utilization metrics or ConsumeWriteCapacityUnits metrics of DynamoDB. You can set up various graphs of resources you are interested in to look at a glance to idea about the assessment of the health of the resources. There is a lot to it – you can add live data to the dashboard. In summary, with the dashboard, you can quickly get the health of your system.

AWS Organizations: AWS Organizations is a management service using which you can consolidate multiple AWS accounts. The consolidation helps in the central management of accounts. As a result, AWS Organizations can help simplify account management – particularly for organizations with multiple AWS accounts. For example, AWS Organizations can help create automated account creation, apply policies to the group of accounts, and consolidate billing. Thus, AWS Organizations provides centralized account and billing management control for organizations and companies with multiple AWS accounts.

AWS Budgets: AWS Budgets give the ability to set custom budgets that alert you when your costs or usage exceed your budgeted or forecasted amount. With AWS Budgets, you can be

alerted by email or SNS notification when actual or forecasted cost and usage exceed your budgeted threshold or when your actual RI and Savings Plans' utilization or coverage drops below your desired threshold. AWS Budgets can be created at different levels, for example, the monthly, quarterly, or yearly levels, and can customize the start and end dates as well. Additionally, you can further refine your budget to track costs associated with multiple dimensions, such as AWS service, linked account, tag, and others.

Reference:
https://aws.amazon.com/aws-cost-management/aws-cost-explorer/
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Dashboards.html
https://aws.amazon.com/organizations/
https://aws.amazon.com/aws-cost-management/aws-budgets/

Question 4:
You have deployed a web application on an EC2 instance which allows users to upload images into S3 buckets. You have users all over the world for this application. Which of the following services would you use to make the experience as good as possible for worldwide users?

A. Amazon CloudFront
B. AWS S3 Accelerator
C. AWS Global Accelerator
D. Edge Locations

Answer: B

Explanation
Correct Option
AWS S3 Accelerator: Amazon S3 Transfer Acceleration (S3TA) is a file transfer service over long distances between the source location and your Amazon S3 bucket. It is fast, easy, and secure. S3 Transfer Acceleration leverages Amazon CloudFront's globally distributed AWS Edge Locations. As data arrives at an AWS Edge Location, data is routed to your Amazon S3 bucket over an optimized network path.

Amazon S3 Transfer Acceleration is designed to optimize transfer speeds from across the world into S3 buckets. So if you are uploading to a centralized bucket from geographically dispersed locations, or if you regularly transfer GBs or TBs of data across continents, you may save hours or days of data transfer time with S3 Transfer Acceleration.

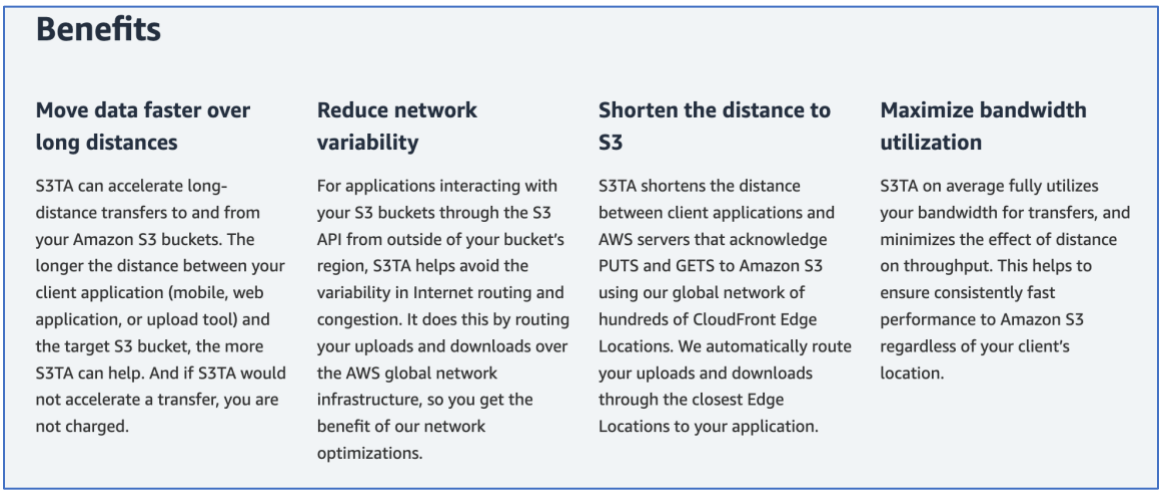

Screenshot from: https://aws.amazon.com/s3/transfer-acceleration/

Incorrect Options

Amazon CloudFront: Amazon CloudFront is a fast CDN (Content Delivery Network) service. It securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds, all within a developer-friendly environment. CloudFront caches data, and a subsequent request for a webpage will not go to the origin server but will be served from the cache. CloudFront uses Edge Locations to cache content.

AWS Global Accelerator: It is not used for S3. When you deploy AWS Global Accelerator, you bring your remote user as close as possible to the AWS backbone. It is a network layer service that you can deploy in front of your Internet-facing applications to improve availability and performance for your globally distributed users. You deploy Global Accelerator between users on the Internet and the public-facing applications deployed and hosted on AWS. Then, the Global Accelerator allows an optimized experience for those users. In addition, it can optimize user experience both for TCP and UDP types of applications.

Like CloudFront, it uses AWS Global network and edge locations for enhanced performance. However, it's an overall performance enhancer than an upload speed accelerator. You cannot use Global Accelerator to speed up media file uploads into S3.

Edge Locations: AWS Services use Edge Locations for different purposes. For example, CloudFront uses Edge Locations to cache content while Global Accelerator uses Edge Locations to find an optimal pathway to the nearest regional endpoint. Amazon S3 Transfer Acceleration uses AWS Edge locations to accelerate upload to S3 buckets. This is an incorrect option for this question.

Reference:

https://aws.amazon.com/global-accelerator
https://aws.amazon.com/cloudfront/
https://aws.amazon.com/s3/transfer-acceleration/

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html

Question 5:

Which of the following design principles is related to the Operational Excellence Pillar of the AWS Well-Architected Framework?

A. Perform operations as code
B. Protect data in transit and at rest
C. Experiment more often
D. Maximize utilization

Answer: A

Explanation

Correct Option

Perform operations as code: The Operational Excellence pillar deals with the ability to run and monitor systems to provide business value. The pillar also talks about continually improving supporting processes and procedures. You can apply this engineering discipline to develop the entire application code in your cloud environment.

You can define your entire workload, for example, applications, and infrastructure, as code and also update them with code. You can also implement your operations procedures as code and automate their execution by triggering them in response to events.

The Operational Excellence pillar has the following design principles:

- Perform operations as code
- Make a frequent small, reversible change
- Refine operations procedures frequently
- Anticipate failure
- Learn from all operational failures

Incorrect Options

Protect data in transit and at rest: This is one of the design principles of the Security Pillar.
Experiment more often: This is one of the design principles of the Performance Efficiency Pillar.
Maximize utilization: This is one of the design principles of the Sustainability Pillar.

Reference:
https://aws.amazon.com/architecture/well-architected/

Question 6:

Your application is writing logs to CloudWatch. However, there is an issue with the application. To troubleshoot the issue, you need to search through around 1000 log files on the CloudWatch. Which options on the AWS CloudWatch can you use to run regular expressions like query to search through 1000 logs?

A. Log Groups

B. Insights
C. Rules
D. Event Buses

Answer: B

Explanation
Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights. There are many components or features of the CloudWatch service. It has Dashboard, Alarms, Logs, Metrics, X-Ray traces, Events, Application monitoring, and Insights sections – these are the main sections.

The CloudWatch service collects data in logs, metrics, and events. After collecting the data, the service provides a unified view of AWS resources, applications, and services that run on AWS and on-premises servers. You can use CloudWatch to detect anomalies or anomalous behavior in your environment. You can set alarms. You can view logs. You can view metrics, take automated action, troubleshot issues, and discover insights to keep your applications running as best as possible with respect to performance, resource utilization, and cost.

The question is how it works. There are four sections to how CloudWatch works: Collect, Monitor, Act, and Analyze.

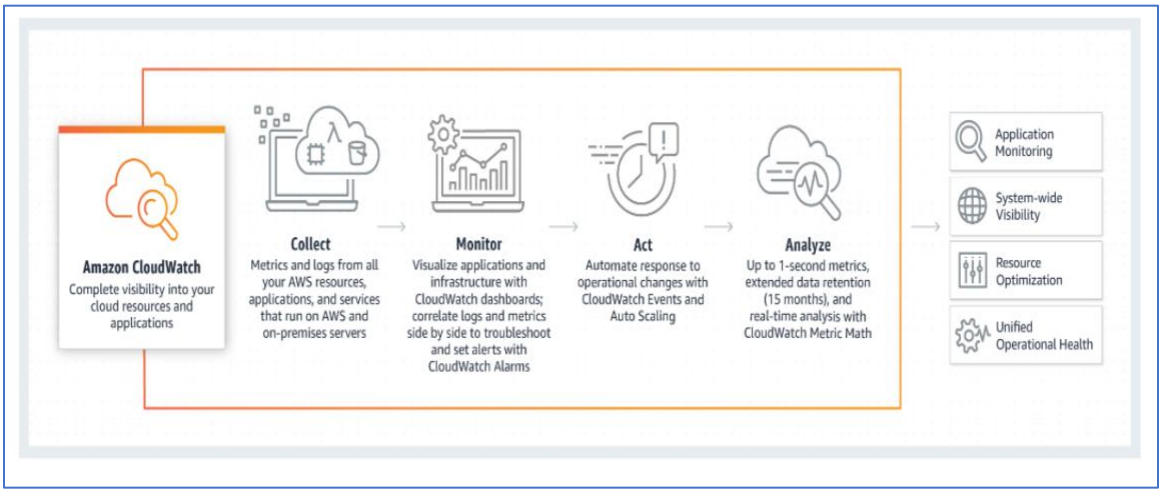

In the Collect section, the services collect all kinds of data (CloudWatch uses more specific data analytic term "metrics") logs from various services and resources in the system. You can also send custom log data from your applications to CloudWatch, for example, log data from AWS Lambda or EMR. For example, EC2 collects data for 38 different types of metrics. If you are interested to know CPU utilization metrics, you can find them in the CloudWatch.

Then there is a Monitor section. Once you have collected data, you can monitor the data with the help of a dashboard in monitoring resources and applications. You can view raw data, query them, and view graphs. In addition, you can view them based on periods and other ways.

Then there is the "Act" part. You can also automate them in response as data is ingested into the CloudWatch system. For example, you can set up an Auto Scaling service based on the CloudWatch event. If the CloudWatch event finds that your CPU utilization exceeds over 75%, new a new EC2 instance can be launched by the Auto Scaling service. You can also act on those data by setting up an alarm in another example. For example, if your bill exceeds 70% of the maximum target budget for the current month, you can utilize CloudWatch alarm to send you a notification in the form of an email.

You can act on data collected, but you can also analyze collected data for up to 1-second metrics and data retentions for up to 15 months for various metrics. These metrics can help you in real-time analytics and host other analyses you can do with the collected data. So, it's a constant feedback loop of getting the data, monitoring, acting, analyzing, and then tuning the system.

This is what CloudWatch is what it does in a nutshell.

So, the main idea about the workflow of the CloudWatch system is that you collect data, monitor them, act on them, analyze them, and tune the system to work the way you want the resource utilization to optimize the performance and reduce your overall cost.

Now let's look into the CloudWatch console to get an idea about various features of the CloudWatch service. First, the CloudWatch console or home page shows the summary of what's going on with different resources if any recent alarms are fired. This page is more geared to if you have configured alarms in your account, for example, a Billing alarm to get a notification you bill is crossing 75% of your total budget for the current month.

CloudWatch Home page

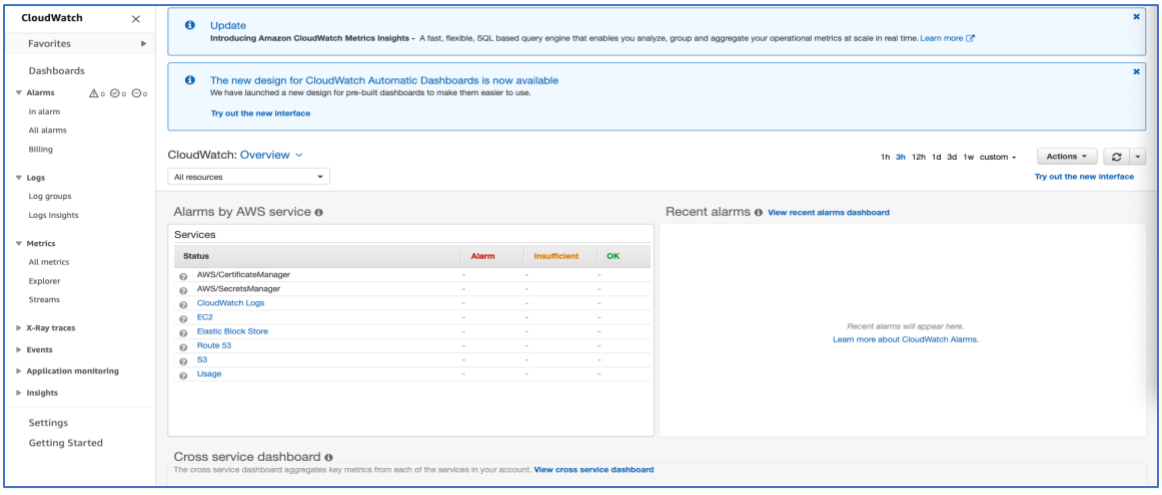

CloudWatch can be considered an umbrella service consisting of many small components or features. As you can see in the screenshot above on the left side, they are many. It has Dashboard, Alarms, Logs, Metrics, X-Ray traces, Events, Application monitoring, and Insights sections – these are the main sections.

CloudWatch can be considered an umbrella service consisting of many small components or features. As you can see in the screenshot above on the left side, they are many. It has Dashboard, Alarms, Logs, Metrics, X-Ray traces, Events, Application monitoring, and Insights sections – these are the main sections.

Correct Option
Insights: The logs section is handy for assessing what's going on in your application. The Logs section is separated into two groups: Log groups and Insights. Log groups are just data dumps for your log files. Insights is a relatively new feature, where you can run SQL looking of kind of query to search through log files. It is very powerful -- you can run regular expression types of query, group, etc. This is very handy as searching through logs for a piece of particular information is a bit involved and not easy as there may be log files in 100 or 1000. And searching for information through many files is a bit involved process.

Incorrect Options
Log Groups: The logs section is handy for assessing what's going on in your application. The Logs section is separated into two groups: Log groups and Insights. Log groups are just data dumps for your log files.
Rules
Event Buses

Both of these are incorrect options. The CloudWatch Events section is divided into two sections Rules and Event Buses. You can create a Rule and add an event pattern to hook that to a Lambda function, or you can add a scheduled event such as running a Lambda function every 30 min. You can configure it to pass input to the Lambda function as well. The Event Buses section is related to event delegation.

Reference:
https://aws.amazon.com/cloudwatch/

Question 7:
Which of the following is the most efficient way to access DynamoDB from an application running on an EC2 instance?

A. Internet Gateway Endpoint
B. VPC Gateway Endpoint
C. VPC Interface Endpoint
D. Virtual Private Endpoint

Answer: B

Explanation
Correct Option
VPC Gateway Endpoint: Many AWS services live outside VPC in the public address space, such as DynamoDB, S3, CloudWatch, AWS Lambda, etc. We know how we can access them using AWS Direct Connect from the on-premises data center – but what about accessing them within your VPC.

For this use case, there are different types of endpoints. Internet Gateway Endpoint, VPC Gateway Endpoint, VPC Interface Endpoints. Internet Gateway Endpoint uses the Internet Gateway of the VPC and uses the Internet to connect services in the public address space. VPC Gateway Endpoint is used only with S3 and DynamoDB. It is much more efficient as it privately communicates without going through the Internet and using public IP addresses. VPC Interface Endpoints are used by all the other available services (except S3 and DynamoDB) such as AWS Web Service API inside your VPC.

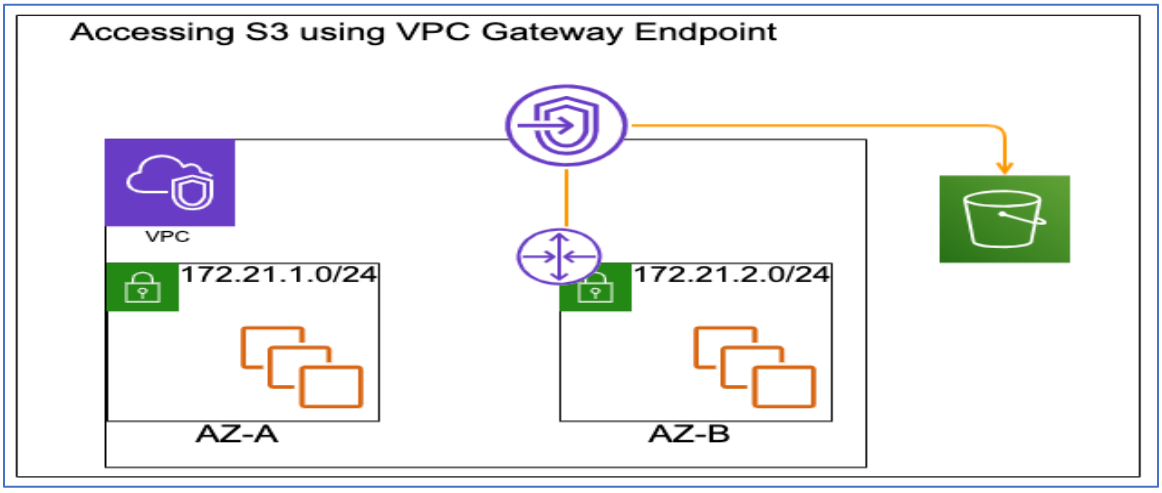

Incorrect Options
Internet Gateway Endpoint: This is incorrect. This is not efficient as it uses the Internet connection.
VPC Interface Endpoint: This is incorrect. VPC Interface Endpoints are used by all the other available services (except S3 and DynamoDB) such as AWS Web Service API inside your VPC.
Virtual Private Endpoint: This is a made-up option.

Reference:
https://docs.aws.amazon.com/vpc/latest/privatelink/gateway-endpoints.html

Question 8:
Which AWS services should you use to detect customer sentiment and analyze customer interactions to categorize inbound support requests automatically?

A. Amazon Kendra
B. Amazon Textract
C. Amazon Transcribe
D. Amazon Comprehend

Answer: D

Explanation

Correct Option

Amazon Comprehend: Amazon Comprehend applies natural-language processing (NLP) to uncover valuable insights and relationships in unstructured text. Regarding the use cases of this service, it can be used to mine business and call center analytics, for example, to detect customer sentiment and analyze customer interactions to categorize inbound support requests automatically. In another use case, the service can index and search product reviews by focusing on context and sentiment, not just keywords. You can also use the Amazon Comprehend service to secure your documents by identifying and redacting Personally Identifiable Information (PII).

Incorrect Options

Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

Amazon Textract: Amazon Textract service enables you to add document text detection and analysis to your applications easily. Using Amazon Textract, customers can automatically extract text and data from millions of scanned documents in just hours. Amazon Textract has many use cases. For example, you can use Amazon Textract to detect typed and handwritten text in various documents. In another use case, using the Amazon Textract Document Analysis API, you can extract text, forms, and tables from structured data documents. You can process invoices and receipts with the AnalyzeExpense API in another use case. Finally, by using the AnalyzeID API, you can process ID documents such as driver's licenses and passports issued by the U.S. government.

Amazon Transcribe: Amazon Transcribe service helps you quickly add high-quality speech-to-text capabilities to your applications. For example, you can quickly extract actionable insights from customer conversations. In another use case, content producers can use this service to convert audio and video assets into fully searchable content automatically. For example, you can create subtitles for your broadcast content to increase accessibility and improve customer experience. Amazon Transcribe service can be used in the medical field as well. For example, medical doctors and practitioners can use Amazon Transcribe Medical to quickly document clinical conversations into electronic health record (EHR) systems for analysis.

Reference:

https://aws.amazon.com/comprehend/
https://aws.amazon.com/kendra/
https://aws.amazon.com/textract/
https://aws.amazon.com/transcribe/

Question 9:

Which of the following is a persistent block storage service?

A. Amazon EFS
B. Amazon EBS

C. Amazon S3
D. Amazon EC2 Instance Store

Answer: B

Explanation
Correct Option
Amazon EBS: Amazon EBS is an easy-to-use, high-performance, block-storage service designed to store persistent data with Amazon EC2 instances. It is a block-storage service and not a file storage service. It is designed to work with EC2 for both throughput and transaction-intensive workloads at any scale. Many workloads, such as enterprise applications, containerized applications, and many other types of applications, are widely deployed on Amazon EBS. An EBS can only be mounted to one EC2 instance at a time.

Incorrect Options
Amazon EFS: Amazon EFS is a cloud-native, serverless, and fully managed file system that is accessible from Linux instances via the NFS protocol. You only pay for what you use for the storage, for reading and writing access to data stored in Infrequent Access storage classes, and for any provisioned throughput.

It is built to scale on-demand to petabytes without disrupting applications. The EFS scales out and scales in automatically as you add and remove files. As a result, you don't need to provision and manage capacity to accommodate growth. In addition, Amazon EFS is designed to provide massively parallel shared access to thousands of EC2 instances. Thus, enabling your applications to achieve high aggregate throughput and IOPS with consistent low latency. Furthermore, there is no minimum fee or setup charge.

Amazon S3: Amazon S3 is an object storage service offering scalability, data availability, security, and performance. Customers having various use cases can use S3 to store, protect, and retrieve any amount of data for different use cases at any time, from anywhere for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.

Amazon EC2 Instance Store: An instance store is temporary block-level storage for an EC2 instance. An instance store provides temporary block-level storage for your EC2 instances. The instance store storage is located on disks that are physically attached to the host computer. An instance store is ideal for the temporary storage of information that frequently changes, such as buffers, caches, and other temporary content. The key point to note about Instance Store is that it is temporary storage. What it means is that the data is lost if the instance experiences failure or when the instance is terminated.

An *instance store* provides temporary block-level storage for your instance. This storage is located on disks that are physically attached to the host computer. Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content, or for data that is replicated across a fleet of instances, such as a load-balanced pool of web servers.

An instance store consists of one or more instance store volumes exposed as block devices. The size of an instance store as well as the number of devices available varies by instance type.

The virtual devices for instance store volumes are `ephemeral[0-23]`. Instance types that support one instance store volume have `ephemeral0`. Instance types that support two instance store volumes have `ephemeral0` and `ephemeral1`, and so on.

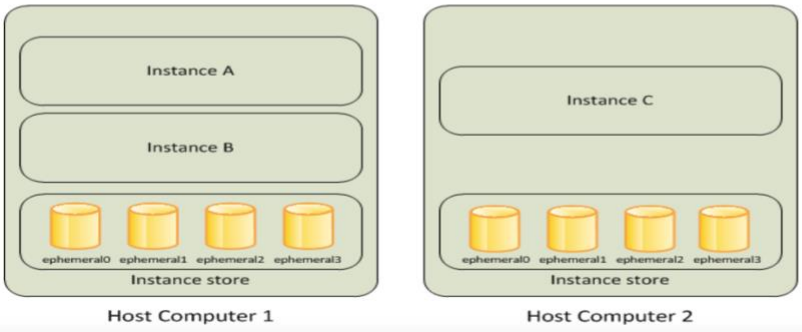

Screenshot from:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html

Reference:
https://aws.amazon.com/efs/
https://aws.amazon.com/ebs/
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html
https://aws.amazon.com/s3/

Question 10:
You have been planning to deploy an event-driven microservices applications to the AWS cloud. You have around 200 microservices in the application. Which of the following services of AWS you can use to help troubleshoot performance issues of a particular microservice?

A. Amazon Macie
B. AWS Lambda
C. AWS CodeStar
D. AWS X-Ray

Answer: D

Explanation
Correct Option
AWS X-Ray: AWS X-Ray helps developers analyze and debug distributed applications, such as applications developed using microservices. AWS X-Ray can help you understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. In addition, x-Ray provides an end-to-end view of requests traveling through your application and shows a map of your application's underlying components.

Incorrect Options
Amazon Macie: Amazon Macie, which is a fully-managed data security and data privacy service, uses machine learning and pattern matching to discover, classify, monitor, protect and report your sensitive data in AWS S3 buckets.

Managing volumes of data as it grows can be challenging, complex, and expensive, particularly at scale, particularly in industries that have strong regulations such as HIPPA and GDPR. Amazon Macie helps automate the discovery of sensitive data at scale and thus lowers the cost of protecting your data.

Amazon Macie automatically offers a list of unencrypted buckets, publicly accessible buckets, and buckets shared with AWS accounts outside of AWS Organizations. Macie then alerts you about personally identifiable information (PII), such as personal data, financial data, and health-related data, by applying machine learning techniques, including pattern matching to the buckets you select.

AWS Lambda: AWS Lambda enables you to execute code without provisioning or managing servers. Instead, you are charged based on the number of requests for your functions and the duration it takes for your code to execute. AWS Lambda executes code in response to events such as object uploads to S3, updates to DynamoDB tables, or other events such as website clicks. Once you upload code to Lambda, Lambda handles all the capacity related to resource provisioning, scaling, patching, and administration of infrastructure to run your code. It also provides visibility into performance by publishing real-time metrics and logs to Amazon CloudWatch. Your job is only to write code.

AWS CodeStar: AWS CodeStar provides a unified interface to bring all AWS CI/CD-related services under one service.

It brings together the following services:

- AWS CodeCommit is essentially a managed git source code repository in AWS.
- AWS CodeBuild, which is like Jenkins. It builds the code and runs tests, and creates deployable artifacts.
- AWS CodeDeploy is an automated software deployment service to deploy your code on an EC2 instance or Elastic Beanstalk.
- AWS Code Pipeline checks out the code, builds it, tests it, and then deploys the code. It is a managed CI/CD pipeline.

Reference:
https://aws.amazon.com/codestar/
https://aws.amazon.com/lambda/
https://aws.amazon.com/macie/
https://aws.amazon.com/xray/

Question 11:
Which of the following services can you use to get recommendations about what Reserved Instances to purchase based on your historical AWS usage?

A. AWS CloudWatch Dashboard
B. AWS Cost Explorer
C. AWS Organizations
D. AWS Budgets

Answer: B

Explanation
Correct Option
AWS Cost Explorer: AWS Cost Explorer service lets you visualize, understand, and manage your AWS costs and usage over time. It has an easy-to-use interface. You can explore your usage and costs using graphs or various Cost Explorer canned reports. For example, you can find your total AWS cost for the last 12 months and forecast for the current month. You can also get AWS costs for individual services, for example, how much you spend on EC2, S3, or any other AWS services.

The Cost Explorer also shows different trends, for example, if your current month's cost usage is down or up for a particular service. If, for the current month, you have uploaded more content on S3, then there is a potential for your S3 usage cost to go up. And it may show up in the trend. These trends can help you understand your charges and usage pattern to analyze them. The analysis may help you reduce AWS costs -- proactively.

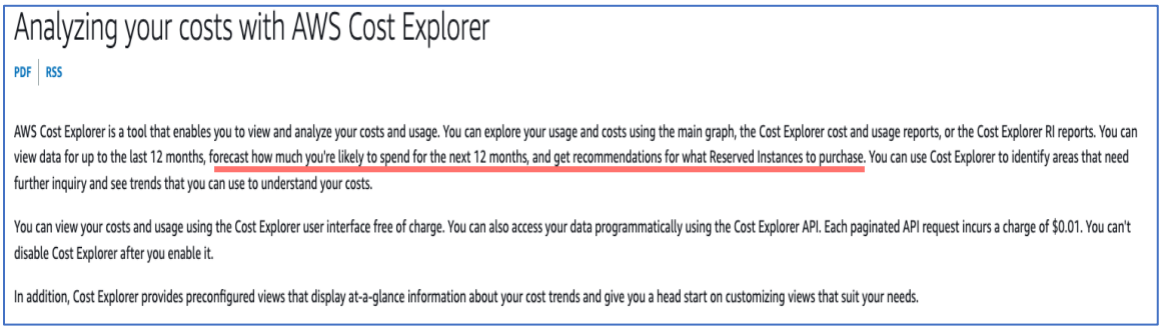

Analyzing your costs with AWS Cost Explorer

PDF | RSS

AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your costs.

You can view your costs and usage using the Cost Explorer user interface free of charge. You can also access your data programmatically using the Cost Explorer API. Each paginated API request incurs a charge of $0.01. You can't disable Cost Explorer after you enable it.

In addition, Cost Explorer provides preconfigured views that display at-a-glance information about your cost trends and give you a head start on customizing views that suit your needs.

Screenshot from:
https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html

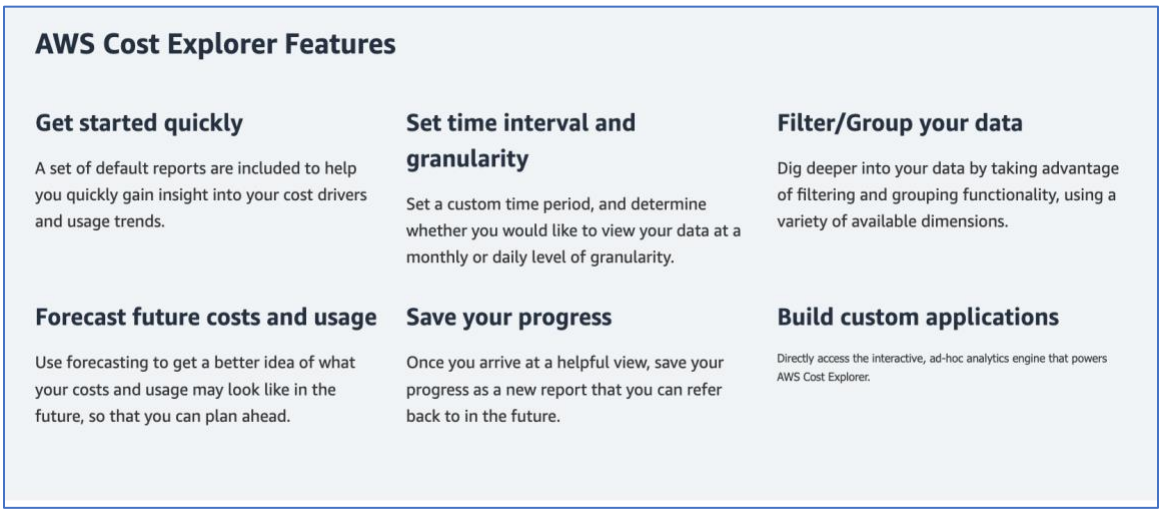

Screenshot from:
https://aws.amazon.com/aws-cost-management/aws-cost-explorer/

Incorrect Options
AWS CloudWatch Dashboard: This is one of the features of AWS CloudWatch. With this feature, you can create a custom dashboard based on your requirements for the kind of metrics you are interested in, particularly resources. For example, if you are interested in EC2 instances CPU utilization metrics or ConsumeWriteCapacityUnits metrics of DynamoDB. You can set up various graphs of resources you are interested in to look at a glance to idea about the assessment of the health of the resources. There is a lot to it – you can add live data to the dashboard. In summary, with the dashboard, you can quickly get the health of your system.

AWS Organizations: AWS Organizations is a management service using which you can consolidate multiple AWS accounts. The consolidation helps in the central management of accounts. As a result, AWS Organizations can help simplify account management – particularly for organizations with multiple AWS accounts. For example, AWS Organizations can help create automated account creation, apply policies to the group of accounts, and consolidate billing. Thus, AWS Organizations provides centralized account and billing management control for organizations and companies with multiple AWS accounts.

AWS Budgets: AWS Budgets give the ability to set custom budgets that alert you when your costs or usage exceed your budgeted or forecasted amount. With AWS Budgets, you can be alerted by email or SNS notification when actual or forecasted cost and usage exceed your budgeted threshold or when your actual RI and Savings Plans' utilization or coverage drops below your desired threshold. AWS Budgets can be created at different levels, for example, the monthly, quarterly, or yearly levels, and can customize the start and end dates as well. Additionally, you can further refine your budget to track costs associated with multiple dimensions, such as AWS service, linked account, tag, and others.

Reference:
https://aws.amazon.com/aws-cost-management/aws-cost-explorer/
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Dashboards.html
https://aws.amazon.com/organizations/
https://aws.amazon.com/aws-cost-management/aws-budgets/

Question 12:
Which of the following statements is not true about Security Group?

A. EC2 Security Groups are stateless.
B. You can assign multiple Security Groups to an EC2 instance.
C. There are quotas about how many rules per Security Groups allowed.
D. When you create VPC, it comes with a default Security Group.

Answer: A

Explanation
Correct Option
Security Groups are stateless: Security Groups are stateful. Security groups are AWS distributed firewalls. They protect your EC2 instances from controlling incoming and outgoing traffic. Inbound rules control the incoming traffic to your instance, and outbound rules control the outgoing traffic from your instance. Security groups act at the instance level, not the subnet level, controlling inbound and outbound traffic. The important point to know about the AWS Security Group, or in general, about any firewall, is that they are stateful. So, if a request is allowed, a response is automatically allowed.

Incorrect Options
You can assign multiple Security Groups to an EC2 instance: You can assign one or more than one security group to an EC2 instance.
There are quotas about how many rules per Security Groups allowed: You can have 60 inbounds and 60 outbound rules per security group.
When you create VPC, it comes with a default Security Group: Each VPC comes with a default Security Group.

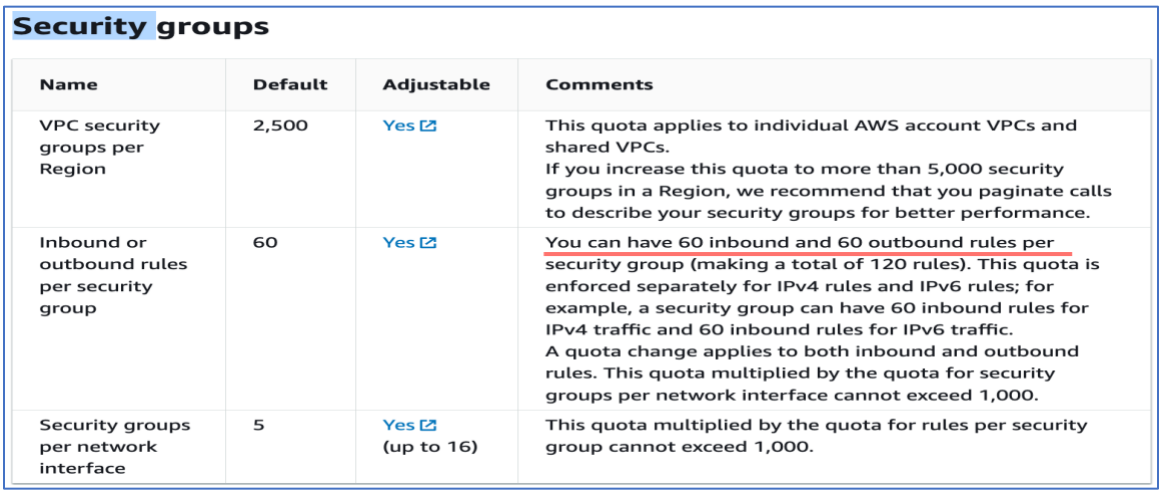

Security groups

Name	Default	Adjustable	Comments
VPC security groups per Region	2,500	Yes	This quota applies to individual AWS account VPCs and shared VPCs. If you increase this quota to more than 5,000 security groups in a Region, we recommend that you paginate calls to describe your security groups for better performance.
Inbound or outbound rules per security group	60	Yes	You can have 60 inbound and 60 outbound rules per security group (making a total of 120 rules). This quota is enforced separately for IPv4 rules and IPv6 rules; for example, a security group can have 60 inbound rules for IPv4 traffic and 60 inbound rules for IPv6 traffic. A quota change applies to both inbound and outbound rules. This quota multiplied by the quota for security groups per network interface cannot exceed 1,000.
Security groups per network interface	5	Yes (up to 16)	This quota multiplied by the quota for rules per security group cannot exceed 1,000.

Reference:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/working-with-security-groups.html
https://docs.aws.amazon.com/vpc/latest/userguide/amazon-vpc-limits.html

Question 13:
You have a use case where you need to use on-premises data to build machine learning models using AWS SageMaker. Which of the following AWS Storage Gateway types can you use to copy on-premises files to S3 cost-effectively?

A. Tape Gateway
B. File Gateway
C. Volume Gateway
D. AWS Direct Connect

Answer: B

Explanation
Correct Option
File Gateway: The File Gateway is one of the Storage Gateway types. The File Gateway is used to integrate on-premises IT infrastructure to Amazon S3. For example, you can use File Gateway -- using NFS and SMB protocol -- to copy on-premises data such as files to Amazon S3. You can take advantage of different S3 storage classes to cost-effectively manage accessibility, durability, and retention.

Regarding the use case of File Gateway, you can use File Gateway to manage a hybrid workload, which requires access to both on-premises and AWS Cloud environments. The File Gateway can be used for big data analytic use cases. For example, you can use File Gateway to copy on-premises data to S3. And there, you leverage AWS services such as EMR, Athena, or Glue to build and run ETL jobs or can perform ad hoc analytics.

You can also leverage AWS machine learning services such as SageMaker, Forecast, and Rekognition to build and run machine models.

You can also use data copied on S3 as backup copy by applying various retention policies.

Incorrect Options
Tape Gateway: The Tape Gateway is one of the Storage Gateway types. The Tape Gateway helps offload your tape back up to the AWS Cloud without disturbing your existing on-premises backup workflow. The Tape Gateway allows you to continue to rely on the current backup workflow, yet you can copy your data to Amazon S3 and then archive it to Glacier as a backup. In addition, his Tape Gateway is compatible with typical backup applications such as Dell EMC NetWorker and Microsoft System Center Data Protection Manager.

Volume Gateway: The Volume Gateway is one of the Storage Gateway types The File Gateway presents your on-premises application to access S3 storage. On the other hand, the Volume Gateway offers your on-premises application to access iSCSI block storage such as EBS volume. In other words, File Gateway is used to access object storage such as S3, while Volume Gateway is used to access block storage such as EBS.

AWS Direct Connect: The AWS Direct Connect is not one of the AWS Storage Gateway types. AWS Direct Connect helps make it easy to set up a dedicated network connection from your on-premises data center to the AWS cloud, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. The AWS Direct Connect is a physical connection. However, AWS Direct Connect does not encrypt your traffic in transit.

Reference:
https://aws.amazon.com/storagegateway/
https://aws.amazon.com/directconnect/

Question 14:
You have many developers in your organization who are busy architecting, designing, and developing the code using Java as the main programming language. Your senior developers who are involved in the code review process complain about not having enough bandwidth for quality code review. In addition, there are many critical pull requests which are pending to be reviewed. Which of the following AWS tools can you use to help speed the code review process?

A. AWS CodeStar
B. AWS CodeBuild
C. Amazon CodeGuru
D. AWS CodeCommit

Answer: C

Explanation
Correct Option
Amazon CodeGuru: Every day, millions of developers analyze codes and find issues to help solve performance issues, yet code issues still occur. This results in additional infrastructure costs and poor customer satisfaction.

Amazon CodeGuru is a developer tool powered by machine learning that provides intelligent recommendations for improving code quality by identifying applications' most expensive lines of code. You can add Amazon CodeGuru to your workflow to identify expensive lines of code to reduce cost. When you add a pull request, Amazon CodeGuru will analyze code from the repository for the critical issues and provides a recommendation report. The report shows why the issue is flagged, the cost of incurring, and suggestive resolution steps.

CodeGuru Reviewer uses machine learning to help identify critical issues such as security vulnerabilities and hard-to-find bugs during application development and provides recommendations to improve code quality.

CodeGuru Profiler pinpoints an application's most expensive lines of code by helping developers understand the runtime behavior of their applications, identify and remove code inefficiencies, and improve performance. In addition, it significantly decreases compute costs.

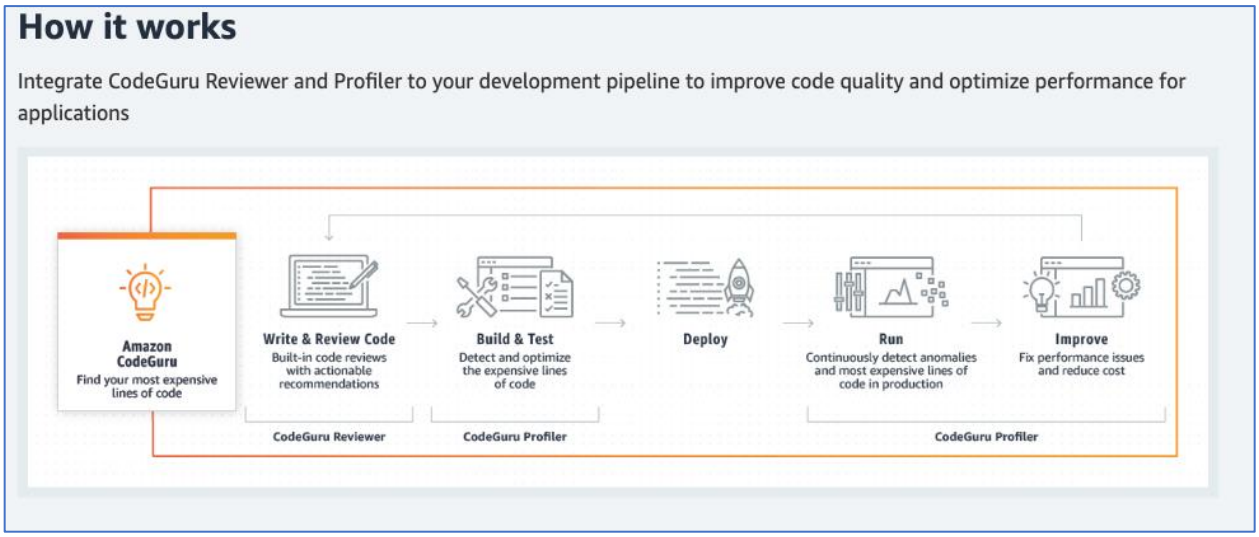

Screenshot from: https://aws.amazon.com/codeguru/

Incorrect Options

AWS CodeStar: AWS CodeStar provides a unified interface to bring all AWS CI/CD-related services under one service.

AWS CodeBuild: AWS CodeBuild, which is a fully managed continuous integration (CI) service, compiles source code, runs tests, and produces software packages ready to deploy. With AWS CodeBuild, you don't need to provision, manage, and scale your build servers. Instead, AWS CodeBuild scales and performs multiple builds concurrently. This helps in builds not left waiting in a queue.

AWS CodeCommit: AWS CodeCommit is a fully managed, secure source control service that hosts git-based repositories. You can create a git repository, add files, clone a repository, create a pull request, and merge pull requests to the branch. In other words, using AWS CodeCommit, you can do all sorts of git operations that you usually do as a developer.

Reference:
https://aws.amazon.com/codestar/
https://aws.amazon.com/codebuild/
https://aws.amazon.com/codeguru/
https://aws.amazon.com/codecommit/

Question 15:

Which of the following services is FREE?

A. AWS WAF
B. AWS Shield Advanced
C. AWS Shield Standard
D. Amazon S3

Answer: C

Explanation
Correct Option
AWS Shield Standard: All AWS customers get the default AWS Shield protection for most common DDoS attacks, by default with no additional charge. It protects from Layer 3 and Layer 4 DDoS attacks for any AWS resource in any AWS Region. It protects against common network attacks such as SYN floods, UDP floods, Reflection attacks, etc., at layer three and layer 4 for any Resources in any AWS Region.

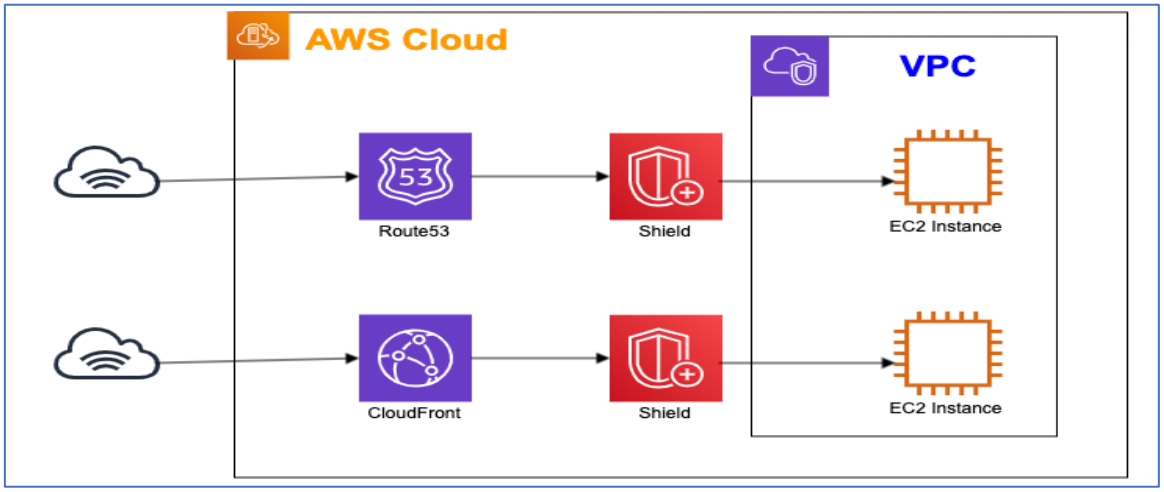

It provides comprehensive defense on layer 3 (Network) and layer 4 (Transport) for the most common DDoS network attacks for CloudFront and Route 53. In other words, if you have your EC2 instance fronted with CloudFront or if you have a DNS request hitting Route 53, you are protected at Layer 3 and Layer 4 of DDoS attacks.

Incorrect Options

AWS WAF: AWS WAF is a paid service**.** AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies.

AWS Shield Advanced: AWS Shield Advanced is a paid service. Being a paid service, it provides additional protection and a number of benefits. It also provides WAF. With AWS Shield Advanced, AWS does additional detection and monitoring. In particular, AWS WAF looks into network flows and data streams more at a customer level, which helps in anomaly detection at a little more granular level. This helps in providing visibility to AWS Shield Advanced customers, which is not available for AWS Shield Standard customers.

Amazon S3: Amazon S3 is a paid service. It is an object storage service offering scalability, data availability, security, and performance. Customers having various use cases can use S3 to store, protect, and retrieve any amount of data for different use cases at any time, from anywhere for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.

Reference:

https://aws.amazon.com/waf/
https://aws.amazon.com/s3/
https://aws.amazon.com/shield/
https://docs.aws.amazon.com/waf/latest/developerguide/getting-started-ddos.html

Question 16:

You are working as a DevOps lead in your company. The software engineering team is deploying a new application in the test environment. However, the developer who is deploying the application doesn't have access to the EC2 instance of the test environment. Which features or services can you use to provide temporary credentials to the developer so that the developer can deploy the application in the test environment?

A. AWS Security Token Service
B. AWS Secrets Manager
C. AWS Web Application Firewall
D. Amazon Cognito

Answer: A

Explanation

Correct Option

AWS Security Token Service: AWS Security Token Service (STS) generates temporary security credentials. It is a web service that allows you to request temporary, limited-privilege credentials

for AWS Identity and Access Management (IAM) users or for federated users (users that are outside of AWS). AWS STS is a global service, and all AWS STS requests go to a single endpoint at https://sts.amazonaws.com. You can only access it programmatically. An STS will return: AccessKeyID, SecretAccessKey, SessionToken, Expiration, and AssumeRoleUser.

You can use this service to create and provide temporary security credentials to trusted users so that they can get access to AWS resources. Suppose you have a user that doesn't have access to upload a document in a particular S3 bucket. Then by using the concept of AssumeRole, you can generate temporary AccessKeyID and SecretAccessKey by the STS service. Using this temporary SecretAccessID and SecretAccessKey can upload the document in the S3 bucket. Since STS credentials are temporary, it expires based on the expiration time.

Temporary security credentials are short-term credentials, as the name implies. They can be configured to last for a few minutes to several hours. After that, AWS no longer recognizes the credentials or allows access from API requests after the credentials expire.

Incorrect Options
AWS Secrets Manager: AWS Secrets Manager, a secrets management service, protects your AWS account's access to applications, services, and IT resources. The service helps you easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs. This call eliminates the need to hardcode sensitive information in plain text. The AWS Secrets Manager also offers secret rotation with built-in integration for Amazon RDS, Redshift, and Amazon DocumentDB.

AWS Web Application Firewall: AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies.

Amazon Cognito: Amazon Cognito is a simple user identity service. It lets you add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. With Amazon Cognito, you also can authenticate users through social identity providers such as Facebook, Twitter, or Amazon, with SAML identity solutions, or by using your identity system.

Reference:
https://stackoverflow.com/questions/50082732/what-is-exactly-assume-a-role-in-aws
https://serverfault.com/questions/1066180/aws-iam-roles-what-is-a-trusted-entity-exactly
https://aws.amazon.com/cognito/
https://aws.amazon.com/waf/

Question 17:
You are looking for the most cost-effective option to run around 1000 ETL jobs for about 1 month. Each of them runs hourly every day. You are ok if these jobs are interrupted as you have added

a hook to handle the interruption to save the state of a job so that next time the job gets the instance, the job resumes from the state it was interrupted. Which of the following EC2 Instance types would be the most cost-effective option?

A. On-Demand Instance
B. Reserved Instance
C. Spot Instance
D. Dedicated Host

Answer: B

Explanation
Correct Option
Spot Instance: Spot Instance lets you use spare EC2 capacity at a meager price – up to 90% off the price of On-Demand instances. Because Spot Instances enable you to request unused EC2 instances at steep discounts, it helps lower your Amazon EC2 costs. You can get Spot Instances at up to a 90% discount compared to On-Demand prices.

You could use Spot Instances for the various stateless, fault-tolerant applications that don't get impacted if they terminate suddenly. For example, if you have an ETL batch job that can be interrupted and resumed to start from the point where it stopped, you can be a good candidate for using Spot Instances if you are looking to save costs for running these types of jobs. On the hand, since these instances can be terminated at short notice, Spot Instances are not suitable for workloads that need to run at a specific point in time or if the workload result is impacted because of sudden termination of its execution.

Spot Instances can be used in the following scenario:

- Suppose your application has flexibility for its start and stop time. In other words, it's ok if your application is stopped, interrupted, or terminated at any time. The reason is spot instances are acquired using the bidding process. Therefore, there is a probability that you may not get the spot instance at the price you bid. Additionally, spot instances can be terminated because of the bidding nature to acquire instances.
- If you are looking for a large number of computing resources, immediately, you can use Spot Instances.

Incorrect Options
On-Demand Instance: On-Demand Instances are more expensive than Spot Instances. As the name says, An On-Demand Instance is an instance that you use on-demand. You have complete control over its lifecycle — you decide when to launch, stop, hibernate, start, reboot, or terminate it. There is no long-term commitment required when you purchase On-Demand Instances. There is no upfront payment when using On-Demand instances. You can pay by the hour or the second (minimum 60 seconds) depending on which instances you run, with no long-term commitment. On-demand instances are not interrupted.

Reserved Instance: Reserved Instances offer you significant savings on your Amazon EC2 costs compared to On-Demand Instance pricing. For example, Reserved Instances save you up to 75% compared to On-Demand Instances. Reserved Instances is a billing discount applied to the use of

On-Demand Instances in your account -- they are not physical instances. Reserved Instanced can be purchased for a one-year or three-year commitment. You get a more significant discount when you choose a three-year commitment offering. When using Reserved Instances, you will be charged for the entire duration, irrespective of your usage.

Dedicated Host: EC2 Dedicated Hosts is a physical server with EC2 instance capacity fully dedicated for your use. In addition, Amazon EC2 Dedicated Hosts allow you to use your eligible software licenses from vendors such as Microsoft and Oracle on Amazon EC2. This concept, also known as BYOL (bring your licenses), helps you get the flexibility and cost-effectiveness of using your licenses, along with the simplicity, resiliency, and elasticity of AWS. An Amazon EC2 Dedicated Host is a physical server fully dedicated for your use. It can help in addressing corporate compliance requirements as well. With respect to pricing, Dedicated Hosts can be purchased On-Demand at an hourly rate and can be purchased as Reservations with 70% off from the On-Demand price.

Reference:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html
https://aws.amazon.com/ec2/spot/
https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/standard-vs.-convertible-offering-classes.html
https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/standard-vs.-convertible-offering-classes.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html
https://aws.amazon.com/ec2/pricing/on-demand/

Question 18:
The AWS team in your organization is involved in automating many processes such as account creation and applying policies to the group of accounts. Which AWS services can you use to create AWS accounts programmatically?

A. AWS IAM
B. AWS Roles
C. AWS Management Console
D. AWS Organizations

Answer: D

Explanation
Correct Option
AWS Organizations: You can also use AWS Organizations API to help automate the creation of AWS accounts. With simple API calls, you can create new accounts programmatically and then apply policies to these new accounts automatically.

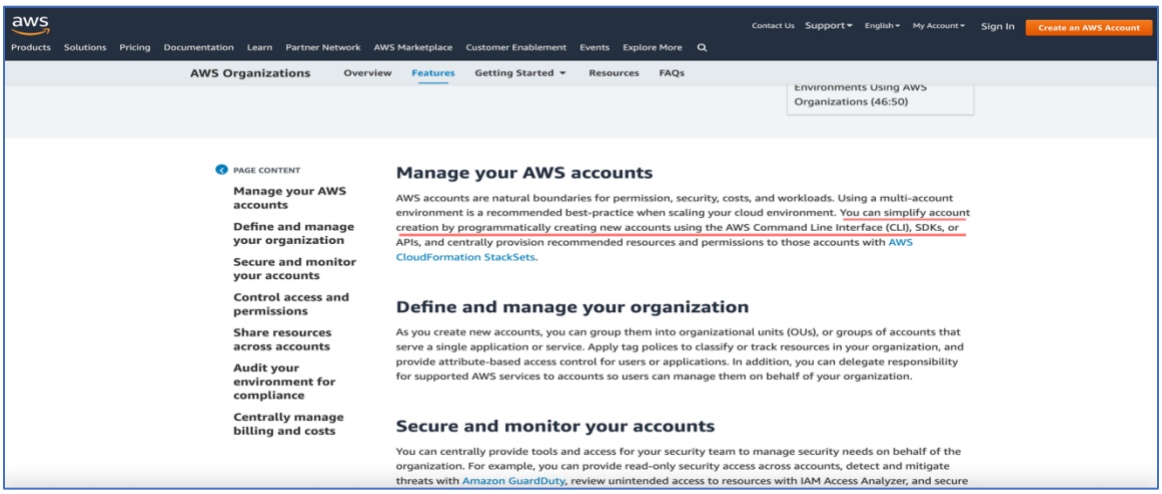

Screenshot from:
https://aws.amazon.com/organizations/features/

Incorrect Options

AWS IAM
AWS Roles
AWS Management Console

These are incorrect options.

Reference:
https://aws.amazon.com/organizations

Question 19:
What is the time limit of an AWS Lambda function per execution?

A. 5 min
B. 10 min
C. 15 min
D. 20 min

Answer: C

Explanation
Correct Option
15 min: In the AWS Lambda service, you are charged based on the number of requests and time taken to execute the function (measured in increments of 100ms). Maximum execution time should not be more than 15 min per execution.

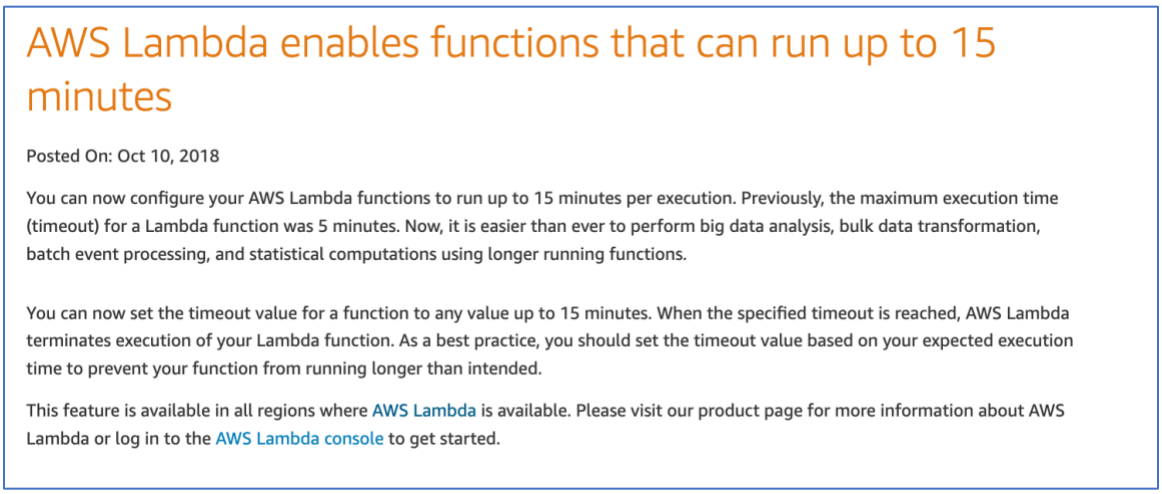

AWS Lambda enables functions that can run up to 15 minutes

Posted On: Oct 10, 2018

You can now configure your AWS Lambda functions to run up to 15 minutes per execution. Previously, the maximum execution time (timeout) for a Lambda function was 5 minutes. Now, it is easier than ever to perform big data analysis, bulk data transformation, batch event processing, and statistical computations using longer running functions.

You can now set the timeout value for a function to any value up to 15 minutes. When the specified timeout is reached, AWS Lambda terminates execution of your Lambda function. As a best practice, you should set the timeout value based on your expected execution time to prevent your function from running longer than intended.

This feature is available in all regions where AWS Lambda is available. Please visit our product page for more information about AWS Lambda or log in to the AWS Lambda console to get started.

Screenshot from:
https://aws.amazon.com/about-aws/whats-new/2018/10/aws-lambda-supports-functions-that-can-run-up-to-15-minutes/

Incorrect Options
All the other options are incorrect.

Reference:
https://aws.amazon.com/about-aws/whats-new/2018/10/aws-lambda-supports-functions-that-can-run-up-to-15-minutes/

Question 20:

You are working as VP of software engineering. You have been asked to find out the organization-wide security posture of the AWS environment in your organization by automated checks based on a security best practice. Which of the following AWS services can you use for this use case?

A. AWS Security Hub
B. AWS Encryption SDK
C. AWS Secrets Manager
D. AWS Artifact

Answer: A

Explanation
Correct Option
AWS Security Hub: AWS Security Hub provides a centralized and organizational-wide cloud security posture management of all your workloads. It performs security best practice checks, aggregates alerts, and enables automated remediation. You can run AWS Security Hub in standalone mode, which means for a single account. Or you can run it in organization mode, where it aggregates data from all your AWS account.

AWS Security Hub provides you with a view of the comprehensive state of security in your AWS account. Additionally, it helps you check your AWS environment against security industry standards and best practices. The AWS Security Hub collects security data from AWS accounts, services, and supported third-party partner products. The data helps analyze security trends and identify the highest priority security issues. AWS Security Hub, once enabled, begins to consume, aggregate, organize and prioritize findings from AWS services that you have enabled, such as Amazon GuardDuty, and Amazon Inspector. AWS Security Hub also generates its own findings by running continuous, automated security checks based on AWS best practices and supported industry standards. Security Hub then correlates and consolidates findings across providers to help you to prioritize the most significant results.

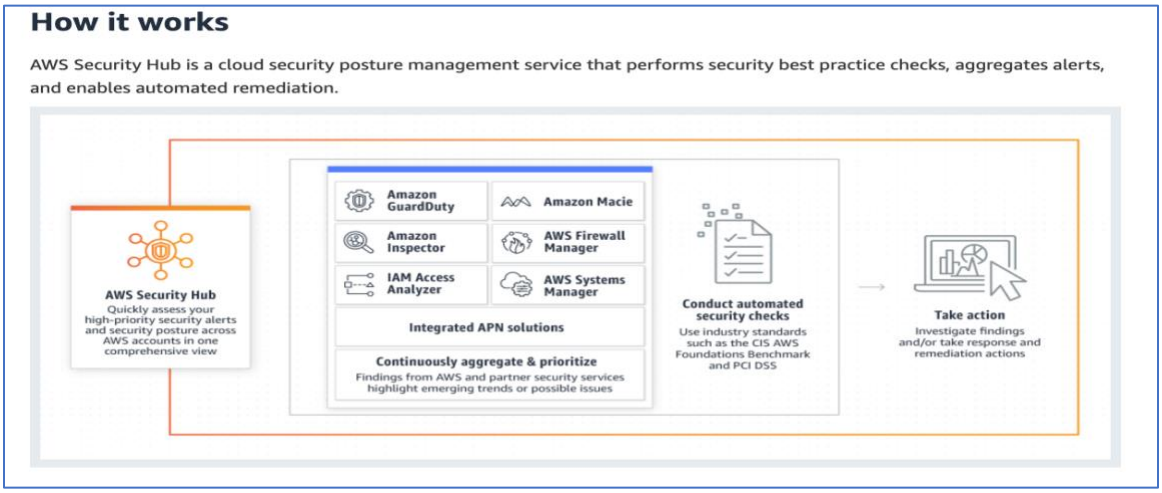

Screenshot from:
https://aws.amazon.com/security-hub/

Incorrect Options

AWS Encryption SDK: The AWS Encryption SDK is a client-side encryption library that you can use to encrypt and decrypt data using industry standards and best practices. The use of AWS Encryption SDK is provided of charge by AWS under the Apache 2.0 license.

AWS Secrets Manager: AWS Secrets Manager, a secrets management service, protects your AWS account's access to applications, services, and IT resources. The service helps you easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs. This call eliminates the need to hardcode sensitive information in plain text. The AWS Secrets Manager also offers secret rotation with built-in integration for Amazon RDS, Redshift, and Amazon DocumentDB.

AWS Artifact: AWS Artifact is a central place for compliance-related information that matters to you. It provides on-demand access to AWS' security and compliance reports and select online agreements. AWS Artifact is a portal using which an enterprise can access security and compliance reports related to the AWS public cloud.

Reference:

https://aws.amazon.com/security-hub/

https://docs.aws.amazon.com/encryption-sdk/latest/developer-guide/introduction.html
https://aws.amazon.com/artifact/
https://aws.amazon.com/secrets-manager/

Question 21:
Which of the following statements is NOT correct?

A. AWS WAF can be deployed on Amazon CloudFront.
B. AWS WAF can be deployed on Amazon API Gateway
C. AWS WAF can be deployed on Amazon S3
D. AWS WAF can be deployed on Application Load Balancer

Answer: C

Explanation
Correct Option
AWS WAF can be deployed on Amazon S3: AWS WAF cannot be deployed on Amazon EC2. AWS WAF can be deployed on Amazon CloudFront, Amazon API Gateway, Application Load Balancer, and AWS AppSync GraphQL to protect your applications against common web exploits and bots, whether applications run in the cloud or on-premises.

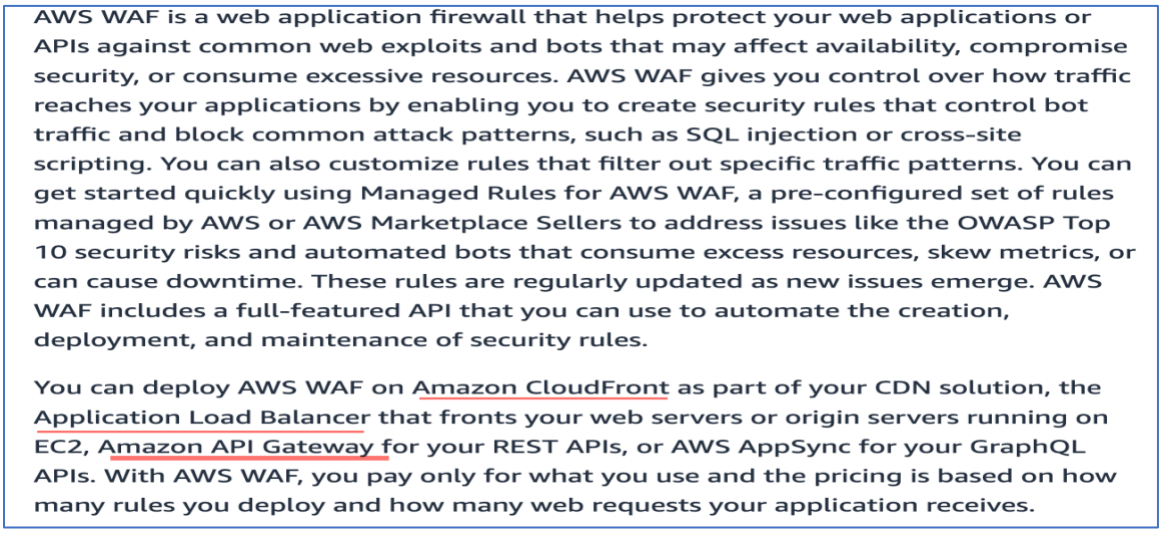
AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that control bot traffic and block common attack patterns, such as SQL injection or cross-site scripting. You can also customize rules that filter out specific traffic patterns. You can get started quickly using Managed Rules for AWS WAF, a pre-configured set of rules managed by AWS or AWS Marketplace Sellers to address issues like the OWASP Top 10 security risks and automated bots that consume excess resources, skew metrics, or can cause downtime. These rules are regularly updated as new issues emerge. AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of security rules.

You can deploy AWS WAF on Amazon CloudFront as part of your CDN solution, the Application Load Balancer that fronts your web servers or origin servers running on EC2, Amazon API Gateway for your REST APIs, or AWS AppSync for your GraphQL APIs. With AWS WAF, you pay only for what you use and the pricing is based on how many rules you deploy and how many web requests your application receives.

Screenshot from:
https://aws.amazon.com/waf/

Incorrect Options
AWS WAF can be deployed on Amazon CloudFront.
AWS WAF can be deployed on Amazon API Gateway
AWS WAF can be deployed on Application Load Balancer

AWS WAF can be deployed on Amazon CloudFront, Amazon API Gateway, and Application Load Balancer.

Reference:
https://aws.amazon.com/waf/

Question 22:
You would like to bring your Windows license, which is based on a number of cores, to AWS Cloud. Which of the following instance types can you use for your Windows license?

A. On-Demand
B. Spot Instance
C. Reserved Instance
D. Dedicated Host

Answer: D

Explanation
Correct Option
Dedicated Host: EC2 Dedicated Hosts is a physical server with EC2 instance capacity fully dedicated for your use. In addition, Amazon EC2 Dedicated Hosts allow you to use your eligible software licenses from vendors such as Microsoft and Oracle on Amazon EC2. This concept, also known as BYOL (bring your licenses), helps you get the flexibility and cost-effectiveness of using your licenses, along with the simplicity, resiliency, and elasticity of AWS.

Incorrect Options
On-Demand
Spot Instance
Reserved Instance

You cannot use these instance types to install a Windows license.

Reference:
https://aws.amazon.com/ec2/spot/
https://aws.amazon.com/ec2/pricing/on-demand/
https://aws.amazon.com/aws-cost-management/aws-cost-optimization/reserved-instances/
https://aws.amazon.com/ec2/dedicated-hosts/

Question 23:
You are executing an AWS Lambda function to process a file when it is uploaded to an S3 bucket. Which of the following options is correct about how an AWS Lambda function's execution is charged?

A. The number of times the function is executed and the time taken to execute the function.
B. The time is taken to execute the function.
C. The number of times the function is executed.
D. The number of times the function is executed, the time taken to execute the function, and the memory consumed by the function during the execution.

Answer: A

Explanation
Correct Option

The number of times the function is executed and the time taken to execute the function: The AWS Lambda doesn't require any upfront cost – it is a low-cost solution to run your code. In the AWS Lambda service, you are charged based on the number of requests and time taken to execute the function (measured in increments of 100ms). Maximum execution time should not be more than 15 min per execution.

Incorrect Options
All the other options are incorrect.

Reference:
https://aws.amazon.com/lambda/

Question 24:
Which of the following AWS service uses machine learning, anomaly detection techniques, and threat intelligence techniques to identify traffic having potential threats?

A. AWS Shield Advanced
B. Amazon GaurdDuty
C. AWS WAF
D. AWS Shield Standard

Answer: B

Explanation
Correct Option
Amazon GaurdDuty: Amazon GuardDuty service, to identify and prioritize potential threats, uses the services related to machine learning, anomaly detection techniques, and other various threat intelligence techniques.

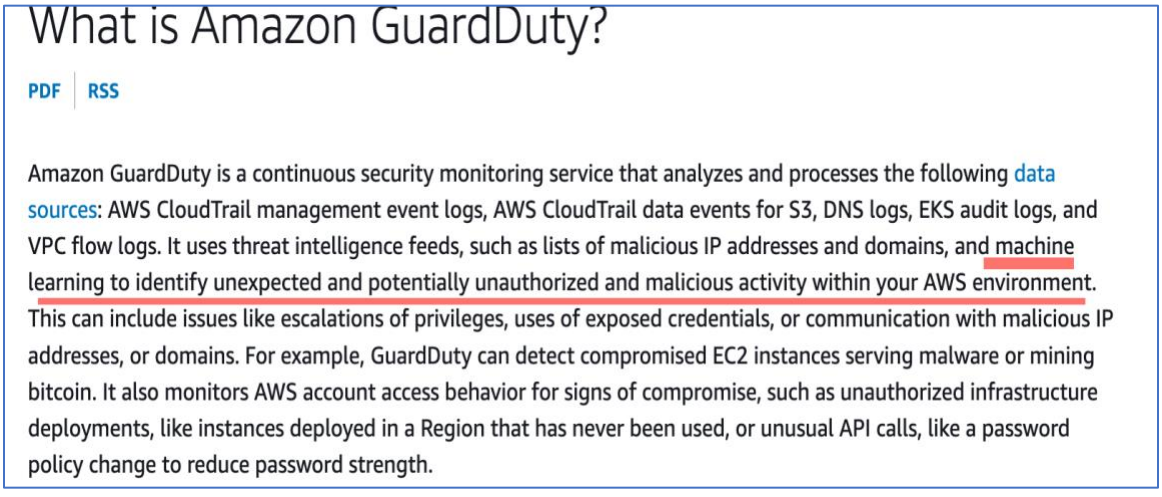

Incorrect Options
AWS Shield Advanced: AWS Shield Advanced is a paid service. Being a paid service, it provides additional protection and a number of benefits. It also provides WAF. With AWS Shield Advanced, AWS does additional detection and monitoring. In particular, AWS looks into

network flows and data streams more at a customer level, which helps in anomaly detection at a little more granular level. This helps in providing visibility to AWS Shield Advanced customers, which is not available for AWS Shield Standard customers.

AWS WAF: AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies.

AWS Shield Standard: All AWS customers get the default AWS Shield protection for most common DDoS attacks, by default with no additional charge. It protects from Layer 3 and Layer 4 DDoS attacks for any AWS resource in any AWS Region. It protects against common network attacks such as SYN floods, UDP floods, Reflection attacks, etc., at layer three and layer 4 for any Resources in any AWS Region.

Reference:
https://aws.amazon.com/waf/
https://aws.amazon.com/shield/
https://aws.amazon.com/guardduty/

Question 25:
Which of the following AWS services can you use to secure documents by identifying and redacting Personally Identifiable Information (PII)?

A. Amazon Kendra
B. Amazon Textract
C. Amazon Transcribe
D. Amazon Comprehend

Answer: D

Explanation
Correct Option
Amazon Comprehend: Amazon Comprehend applies natural-language processing (NLP) to uncover valuable insights and relationships in unstructured text. Regarding the use cases of this service, it can be used to mine business and call center analytics, for example, to detect customer sentiment and analyze customer interactions to categorize inbound support requests automatically. In another use case, the service can index and search product reviews by focusing on context and sentiment, not just keywords. You can also use the Amazon Comprehend service to secure your documents by identifying and redacting Personally Identifiable Information (PII).

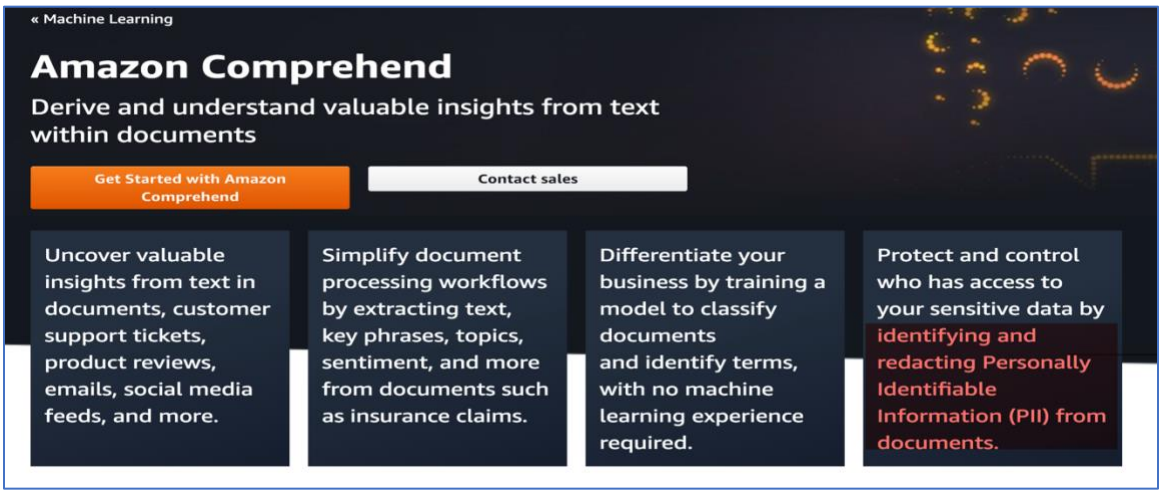

Screenshot from:
https://aws.amazon.com/comprehend/

Incorrect Options
Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

Amazon Textract: Amazon Textract service enables you to add document text detection and analysis to your applications easily. Using Amazon Textract, customers can automatically extract text and data from millions of scanned documents in just hours. Amazon Textract has many use cases. For example, you can use Amazon Textract to detect typed and handwritten text in various documents. In another use case, using the Amazon Textract Document Analysis API, you can extract text, forms, and tables from structured data documents. You can process invoices and receipts with the AnalyzeExpense API in another use case. Finally, by using the AnalyzeID API, you can process ID documents such as driver's licenses and passports issued by the U.S. government.

Amazon Transcribe: Amazon Transcribe service helps you quickly add high-quality speech-to-text capabilities to your applications. For example, you can quickly extract actionable insights from customer conversations. In another use case, content producers can use this service to convert audio and video assets into fully searchable content automatically. For example, you can create subtitles to your broadcast content to increase accessibility and improve customer experience. Amazon Transcribe service can be used in the medical field as well. For example, medical doctors and practitioners can use Amazon Transcribe Medical to quickly document clinical conversations into electronic health record (EHR) systems for analysis.

Reference:
https://aws.amazon.com/comprehend/
https://aws.amazon.com/kendra/
https://aws.amazon.com/textract/

https://aws.amazon.com/transcribe/

Question 26:
Which programming language cannot be used to write an AWS Lambda function?

A. Java
B. Python
C. C++
D. Ruby

Answer: C

Explanation
Correct Option
C++ : The languages used to write AWS function are: Java, Go, PowerShell, Node. js, C#, Python, and Ruby

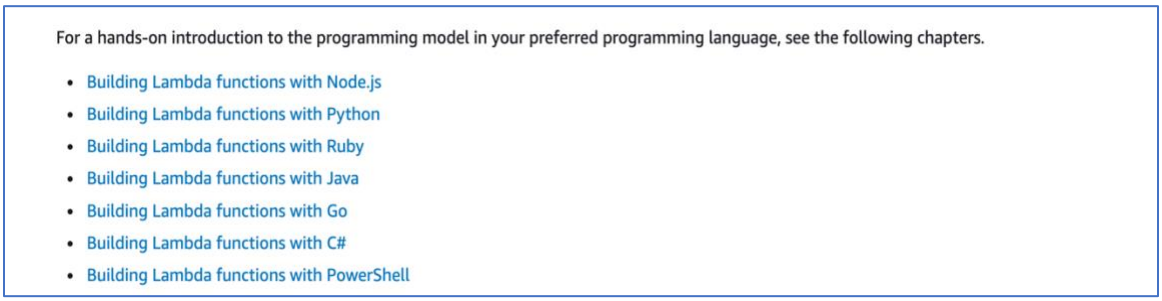

Screenshot from:
https://docs.aws.amazon.com/lambda/latest/dg/foundation-progmodel.html

Incorrect Options
Java – You can use Java programming language for an AWS Lambda function.
Python – You can use Python programming language for an AWS Lambda function.
Ruby – You can use Ruby programming language for an AWS Lambda function.

Reference:
https://docs.aws.amazon.com/lambda/latest/dg/foundation-progmodel.html

Question 27:
You are looking for an EC2 instance for 1 month for doing integration testing of the application that your team has recently worked on. You don't want the testing EC2 instances to be interrupted. Which of the following instance types will be the best fit for this use case?

A. On-Demand Instance
B. Reserved Instance
C. Spot Instance
D. Dedicated Host

Answer: A

Explanation
Correct Option
On-Demand Instance: Since the integration testing of the application is only for 1 month and you don't want integration testing to be interrupted. This is the best option for this use case. In the other option, Reserved Instance you cannot get for one month. The Spot Instance option cannot be used because Spot Instances can be interrupted. A Dedicated Host is expensive and usually used for BYOL use cases.

Incorrect Options

Reserved Instance
Spot Instance
Dedicated Host

These options are not the best fit for this use case.

Reference:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html
https://aws.amazon.com/ec2/pricing/on-demand/
https://aws.amazon.com/ec2/spot/
https://aws.amazon.com/aws-cost-management/aws-cost-optimization/reserved-instances/

Question 28:
Which of the following statements is true related to AWS Shield Standard?

A. AWS Shield Standard cannot protect the Network layer from DDoS attacks.
B. AWS Shield Standard cannot protect the Transport layer from DDoS attacks.
C. AWS Shield Standard cannot protect the Application layer from DDoS attacks.
D. AWS Shield Standard cannot protect CloudFront from layer 3 and layer 4 DDoS attacks.

Answer: C

Explanation
Correct Option
AWS Shield Standard cannot protect Application layer from DDoS attack: All AWS customers get the default AWS Shield protection for most common DDoS attacks, by default with no additional charge. It protects from Layer 3 and Layer 4 DDoS attacks for any AWS resource in any AWS Region.

It protects against common network attacks such as SYN floods, UDP floods, Reflection attacks, etc., at layer three and layer 4 for any Resources in any AWS Region.

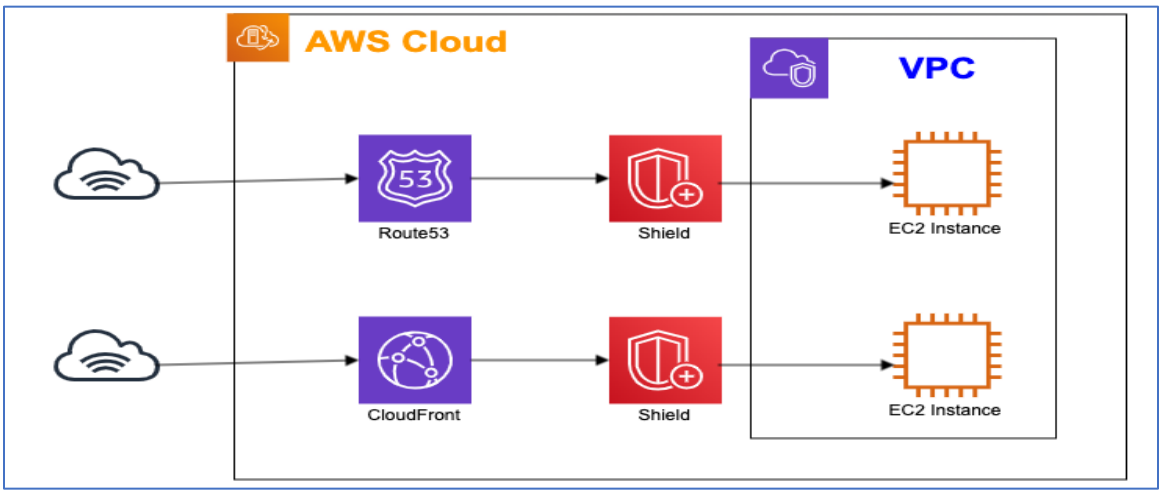

It provides comprehensive defense on layer 3 (Network) and layer 4 (Transport) for the most common DDoS network attacks for CloudFront and Route 53. In other words, if you have your EC2 instance fronted with CloudFront or if you have a DNS request hitting Route 53, you are protected at Layer 3 and Layer 4 of DDoS attacks.

If you have a concern about layer seven attacks, you can subscribe to AWS WAF. AWS WAF is a paid self-service with a pay-as-you model. AWS WAF protects layer 7 – the application layer – from DDoS attacks.

Incorrect Options

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield - Standard and Advanced.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your web site or applications. When you use AWS Shield Standard with Amazon CloudFront and Amazon Route 53, you receive comprehensive availability protection against all known infrastructure (Layer 3 and 4) attacks.

For higher levels of protection against attacks targeting your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator and Amazon Route 53 resources, you can subscribe to AWS Shield Advanced. In addition to the network and transport layer protections that come with Standard, AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, near real-time visibility into attacks, and integration with AWS WAF, a web application firewall. AWS Shield Advanced also gives you 24x7 access to the AWS Shield Response Team (SRT) and protection against DDoS related spikes in your Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator and Amazon Route 53 charges.

Screenshot from: https://aws.amazon.com/shield/

AWS Shield Standard cannot protect the Network layer from DDoS attacks: AWS Shield Standard can protect the Network layer from DDoS attacks.

AWS Shield Standard cannot protect the Transport layer from DDoS attacks: AWS Shield Standard can protect the Transport layer from DDoS attack

AWS Shield Standard cannot protect CloudFront from layer 3 and layer 4 DDoS attacks: AWS Shield Standard can protect CloudFront from layer 3 and layer 4 DDoS attacks.

Reference:
https://aws.amazon.com/shield/

Question 29:
You have a use case where you need to use cloud storage for your current tape backup without making any change in the existing backup and archive workflow. Which of the following AWS Storage features can you use to replace the current on-premises tape backup solution with the AWS cloud backup solution cost-effectively?

A. Tape Gateway
B. File Gateway
C. Volume Gateway
D. AWS Direct Connect

Answer: A

Explanation
Correct Option
Tape Gateway: AWS Storage Gateway, which is a hybrid cloud storage service, gives you on-premises access to AWS cloud storage. AWS Storage Gateway service connects your existing on-premises IT infrastructure with the AWS Cloud. The AWS Storage Gateway provides a standard set of storage protocols such as iSCSI, SMB, and NFS. AWS Storage Gateway allows you to use AWS storage without rewriting your existing applications.

The Tape Gateway helps offload your tape back up to the AWS Cloud without disturbing your existing on-premises backup workflow. The Tape Gateway allows you to continue to rely on the current backup workflow, yet you can copy your data to Amazon S3 and then archive it to Glacier as a backup. In addition, his Tape Gateway is compatible with typical backup applications such as Dell EMC NetWorker and Microsoft System Center Data Protection Manager.

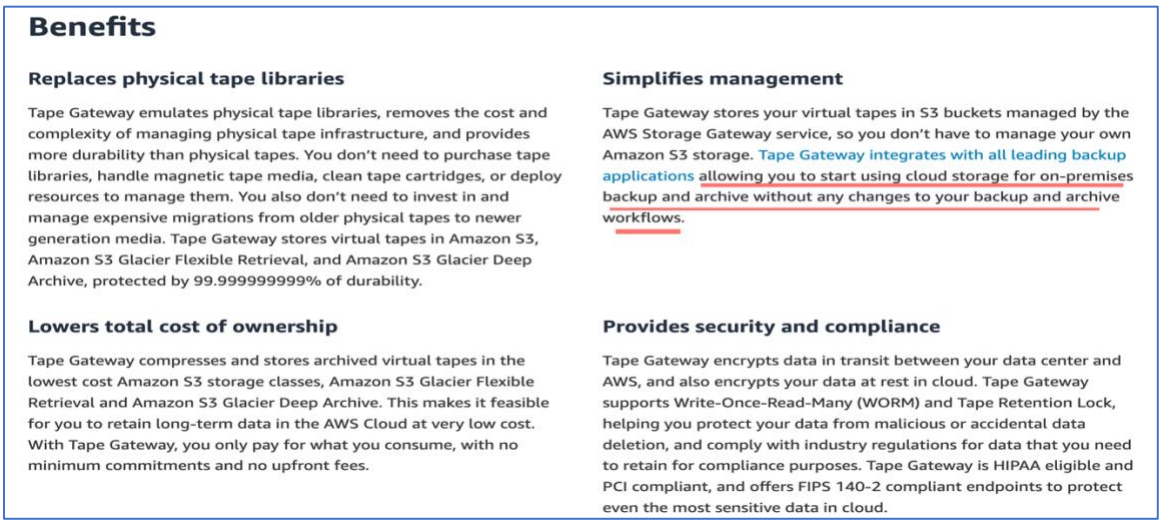

Benefits

Replaces physical tape libraries

Tape Gateway emulates physical tape libraries, removes the cost and complexity of managing physical tape infrastructure, and provides more durability than physical tapes. You don't need to purchase tape libraries, handle magnetic tape media, clean tape cartridges, or deploy resources to manage them. You also don't need to invest in and manage expensive migrations from older physical tapes to newer generation media. Tape Gateway stores virtual tapes in Amazon S3, Amazon S3 Glacier Flexible Retrieval, and Amazon S3 Glacier Deep Archive, protected by 99.999999999% of durability.

Simplifies management

Tape Gateway stores your virtual tapes in S3 buckets managed by the AWS Storage Gateway service, so you don't have to manage your own Amazon S3 storage. Tape Gateway integrates with all leading backup applications allowing you to start using cloud storage for on-premises backup and archive without any changes to your backup and archive workflows.

Lowers total cost of ownership

Tape Gateway compresses and stores archived virtual tapes in the lowest cost Amazon S3 storage classes, Amazon S3 Glacier Flexible Retrieval and Amazon S3 Glacier Deep Archive. This makes it feasible for you to retain long-term data in the AWS Cloud at very low cost. With Tape Gateway, you only pay for what you consume, with no minimum commitments and no upfront fees.

Provides security and compliance

Tape Gateway encrypts data in transit between your data center and AWS, and also encrypts your data at rest in cloud. Tape Gateway supports Write-Once-Read-Many (WORM) and Tape Retention Lock, helping you protect your data from malicious or accidental data deletion, and comply with industry regulations for data that you need to retain for compliance purposes. Tape Gateway is HIPAA eligible and PCI compliant, and offers FIPS 140-2 compliant endpoints to protect even the most sensitive data in cloud.

Screenshot from:
https://aws.amazon.com/storagegateway/vtl/

Incorrect Options
File Gateway: The File Gateway is one of the Storage Gateway types. The File Gateway is used to integrate on-premises IT infrastructure to Amazon S3. For example, you can use File Gateway -- using NFS and SMB protocol -- to copy on-premises data such as files to Amazon S3. You can take advantage of different S3 storage classes to cost-effectively manage accessibility, durability, and retention.

Volume Gateway: The File Gateway presents your on-premises application to access S3 storage. On the other hand, the Volume Gateway offers your on-premises application to access iSCSI block storage such as EBS volume. In other words, File Gateway is used to access object storage such as S3, while Volume Gateway is used to access block storage such as EBS.

AWS Direct Connect: AWS Direct Connect helps make it easy to set up a dedicated network connection from your on-premises data center to the AWS cloud, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. The AWS Direct Connect is a physical connection. However, AWS Direct Connect does not encrypt your traffic in transit.

Reference:
https://aws.amazon.com/storagegateway/vtl/
https://aws.amazon.com/storagegateway/
https://aws.amazon.com/directconnect/

Question 30:
You are having an availability issue with one of AWS services. Which of the following can help you find out if a particular service is available or not?

A. AWS CloudWatch

B. AWS CloudTrail
C. AWS Service Health Dashboard
D. AWS Systems Manager

Answer: C

Explanation
Correct Option
AWS Service Health Dashboard: The AWS Service Health Dashboard is an AWS general service event dashboard where you can view the overall health of AWS services. In other words, the AWS Service Health Dashboard is the single place to find out about the availability and health of AWS services. In addition, you can view the overall status of AWS services. Amazon Web Services publishes up-to-the-minute information on service availability using its Health Dashboard page. You can check the page to get current status information about AWS services or subscribe to an RSS feed to be notified about any interruption of AWS services.

The page can be accessed via the URL - https://status.aws.amazon.com/.

Incorrect Options
AWS CloudWatch: This service is not used to find availability issues with one of AWS services. It is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights.

AWS CloudTrail: This service is not used to find availability issues with one of AWS services. AWS CloudTrail provides auditing, security monitoring, and operational troubleshooting by tracking user activity and API usage.

AWS Systems Manager: This service is not used to find availability issues with one of AWS services. AWS Systems Manager is a set of fully managed services and capabilities that simplify the management of your Windows and Linux instances regardless of whether they are running on EC2 or on-premises.

Reference:
https://health.aws.amazon.com/health/status
https://aws.amazon.com/cloudwatch/
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html
https://aws.amazon.com/systems-manager/

Question 31:
Which of the following statements is not true with regards to EC2 instance data transfer?
A. There is no charge for inbound data transfer across all services in all Regions.
B. Data transfer from AWS to the internet is charged per service.
C. If the internet gateway is used to access the public endpoint of the AWS services in the same Region, there are no data transfer charges.

D. If a NAT gateway is used to access the same services, there is a data no processing charge.

Answer: D

Explanation
Correct Option

If a NAT gateway is used to access the same services, there is a data no processing charge.

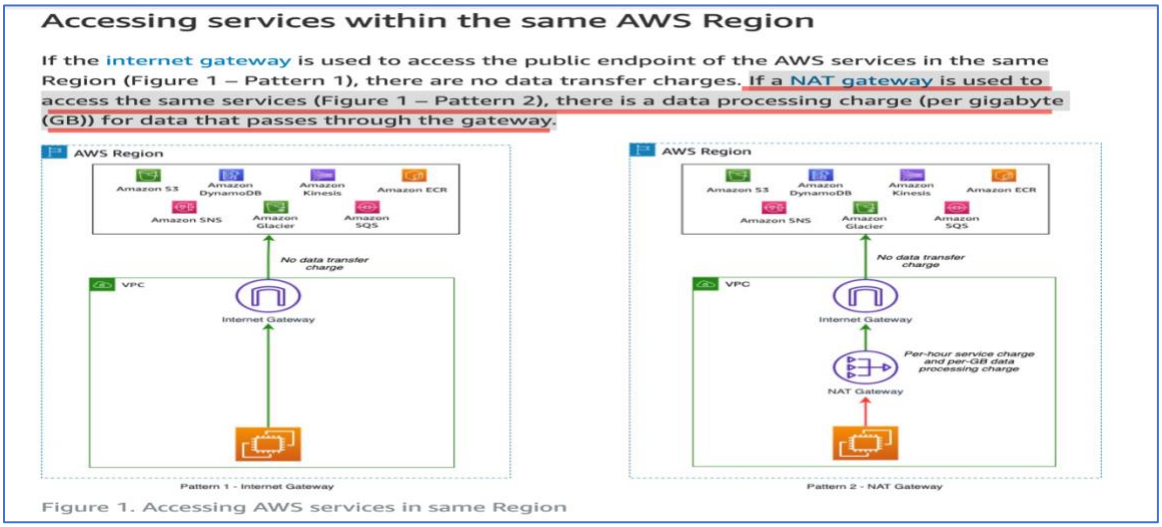

Figure 1. Accessing AWS services in same Region

Screenshot from:
https://aws.amazon.com/blogs/architecture/overview-of-data-transfer-costs-for-common-architectures/

Incorrect Options

There is no charge for inbound data transfer across all services in all Regions.
Data transfer from AWS to the internet is charged per service.

Data transfer between AWS and internet

There is no charge for inbound data transfer across all services in all Regions. Data transfer from AWS to the internet is charged per service, with rates specific to the originating Region. Refer to the pricing pages for each service—for example, the pricing page for Amazon Elastic Compute Cloud (Amazon EC2)—for more details.

If the internet gateway is used to access the public endpoint of the AWS services in the same Region, there are no data transfer charges.

All these statements are correct.

Reference:
https://aws.amazon.com/blogs/architecture/overview-of-data-transfer-costs-for-common-architectures/

Question 32:
You have deployed a microservices application on three EC2 instances. You have fronted this with Application Load Balancer. You would like to protect the login URL from brute force attacks. Which of the following services can you provide protection and monitoring?

A. AWS Shield Standard
B. AWS Web Application Firewall (WAF)
C. AWS Firewall Manager
D. AWS CloudWatch

Answer: B

Explanation
Correct Option
AWS Web Application Firewall (WAF): The brute force attack on the application's login URL is a layer 7, application-layer attack. The AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies.

Incorrect Options
AWS Shield Standard: AWS Shield Standard doesn't protect from application-layer attacks. It protects from Layer 3 and Layer 4 DDoS attacks for any AWS resource in any AWS Region. It protects against common network attacks such as SYN floods, UDP floods, Reflection attacks, etc., at layer three and layer 4 for any Resources in any AWS Region. All AWS customers get the default AWS Shield protection for most common DDoS attacks, by default with no additional charge.

AWS Firewall Manager: It doesn't protect from application-layer attacks. AWS Firewall Manager, a security management service, allows you to configure and manage firewall rules -- centrally -- across your accounts and applications in AWS Organizations. Using AWS Firewall Manager, you can build firewall rules, create security policies, and enforce them consistently across your entire AWS infrastructure from a central administrator account. When new applications are created, AWS Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules. You can easily implement AWS WAF rules for your application load balancers, API Gateways, and Amazon CloudFront distributions using AWS Firewall Manager.

AWS CloudWatch: It doesn't protect from application-layer attacks. Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights.

Reference:
https://aws.amazon.com/cloudwatch/
https://aws.amazon.com/firewall-manager/
https://aws.amazon.com/waf/
https://aws.amazon.com/shield

Question 33:
You are working as an AWS consultant for a company that is involved in a cloud migration project. The company would like to extend its on-premises IT infrastructure to connect to the AWS VPC to speed up some of its projects. The company would like to have a consistent high-bandwidth connection set up between on-premises and the AWS VPC. Which of the following options would you recommend for this use case?

A. AWS Direct Connect
B. AWS Site-to-Site VPN
C. Virtual Private Gateway
D. Customer Gateway

Answer: D

Explanation
Correct Option
AWS Direct Connect: The AWS Direct Connect is a physical connection. The AWS Direct Connect provides a more predictable and consistent experience for on-premises connectivity to the AWS. Using AWS Direct Connect, you can connect on-premises and AWS VPC. AWS Direct Connect uses a physical connection. To use AWS Direct Connect, you can use the AWS Direct Connect location near your data center location. AWS has many Direct Connection locations; each of these locations has Routers managed by AWS. Then you request port on one of these Routers – maybe 1G/sec or 10 G/sec. And then, you can set up AWS connectivity either by yourself or by taking the help of an AWS Connectivity partner.

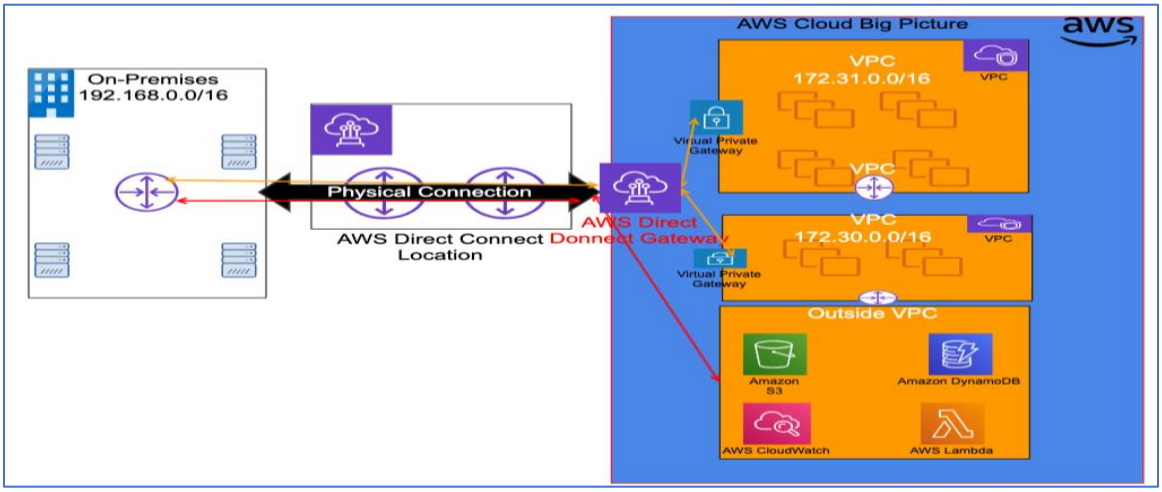

Incorrect Options

AWS Site-to-Site VPN: When extending an on-premises environment to the AWS Cloud, VPN is generally fast and easy to set up for many AWS customers.

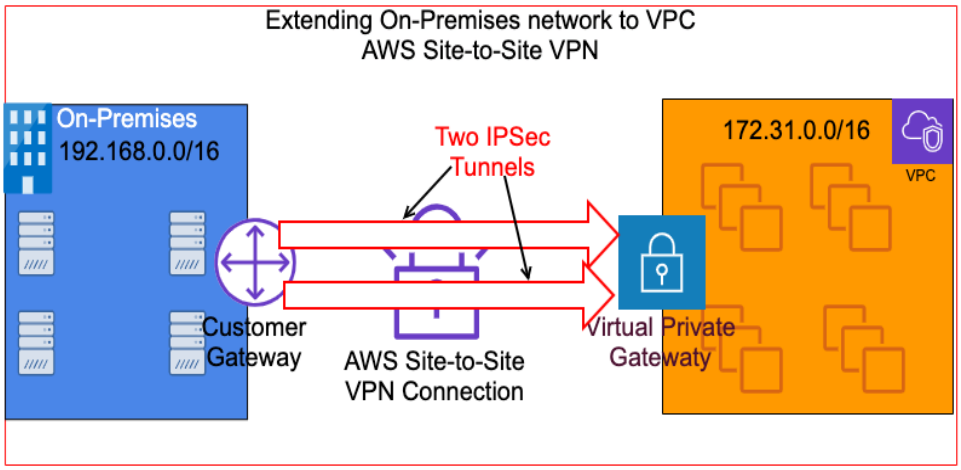

You need two things: Customer Gateway and Virtual Private Gateway. The Customer Gateway could be your Router, firewall, or other things that support IPSec in your on-premises environment. The Customer Gateway resides at the on-premises end, and it is used in the AWS Site-to-Site VPN connection. The next is Virtual Private Gateway which you create on AWS and associate that Virtual Private Gateway to a VPC.

Virtual Private Gateway: The Virtual Private Gateway is not full connectivity option between the on-premises and AWS VPC. A virtual gateway allows resources that are outside of your VPC to communicate to resources that are inside of your VPC.

Customer Gateway: The customer gateway is not a full connectivity option between the on-premises and AWS VPC. It is at the customer end of the connectivity. A customer gateway allows resources that are outside of the on-premises environment to communicate to resources that are inside of the on-premises environment.

Reference:

https://docs.aws.amazon.com/vpn/latest/s2svpn/VPC_VPN.html
https://aws.amazon.com/directconnect/
https://docs.aws.amazon.com/directconnect/latest/UserGuide/virtualgateways.html
https://docs.aws.amazon.com/vpn/latest/s2svpn/your-cgw.html

Question 34:
You are looking for a large number of computing resources immediately. Which of the following instance types will be the best fit for this use case?

A. On-Demand Instance
B. Reserved Instance
C. Spot Instance
D. Dedicated Host

Answer: C

Explanation
Correct Option
Spot Instance: Spot Instances are a good option if you are looking for a large number of computing resources immediately. Spot Instance lets you use spare EC2 capacity at a meager price – up to 90% off the price of On-Demand instances. Because Spot Instances enable you to request unused EC2 instances at steep discounts, it helps lower your Amazon EC2 costs. You can get Spot Instances at up to a 90% discount compared to On-Demand prices.

Spot Instances can be used in the following scenario:

- Suppose your application has flexibility for its start and stop time. In other words, it's ok if your application is stopped, interrupted, or terminated at any time. The reason is spot instances are acquired using the bidding process. Therefore, there is a probability that you may not get the spot instance at the price you bid. Additionally, spot instances can be terminated because of the bidding nature to acquire instances.
- If you are looking for a large number of computing resources, immediately, you can use Spot Instances.

Incorrect Options
On-Demand Instance
Reserved Instance
Dedicated Host

These cannot provide a large amount of computing resources, immediately.

Reference:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html
https://aws.amazon.com/ec2/spot/
https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/standard-vs.-convertible-offering-classes.html
https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/standard-vs.-convertible-offering-classes.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html
https://aws.amazon.com/ec2/pricing/on-demand/

Question 35:
Your company has many departments and each of these departments has many AWS accounts. There are budget issues, and your finance controller needs to see consolidated AWS billing to centralize the cost. You being the DevOps lead, which of the following AWS services /features will you use to consolidate AWS bills of multiple AWS accounts?

A. AWS Budgets
B. AWS Organizations
C. Amazon CloudWatch
D. AWS Cost and Usage

Answer: B

Explanation
Correct Option
AWS Organizations: AWS Organizations is a management service using which you can consolidate multiple AWS accounts. The consolidation helps in the central management of accounts. As a result, AWS Organizations can help simplify account management – particularly for organizations with multiple AWS accounts. For example, AWS Organizations can help create automated account creation, apply policies to the group of accounts, and consolidate billing. Thus, AWS Organizations provides centralized account and billing management control for organizations and companies with multiple AWS accounts.

Many organizations have found themselves using multiple AWS accounts as they have scaled up their AWS usage for various reasons. For example, some customers have added AWS accounts incrementally as more users or departments started using AWS. Other customers have created separate AWS accounts for Dev, Test, and Prod environments to meet strict guidelines such as HIPPA, PCI, or other compliance.

As these AWS accounts grow, these customers would like to add policies and manage billing across their accounts in a simple and more scalable way – without requiring manual processes or custom scripts. And they also would like to add or create new accounts with the policies applied. AWS Organizations can help with account management. Organizations want policy-based management for multiple AWS accounts. You can create a group of accounts and then add policies to those accounts that centrally control the use of AWS services down to the API level across multiple accounts. For example, you can create a collection of production accounts and then apply policies about which AWS services, resources, and API calls those accounts can use.

You can also use AWS Organizations API to help automate the creation of AWS accounts. With simple API calls, you can create new accounts programmatically and then apply policies to these new accounts automatically. With AWS Organizations, you can set up a single payment to all AWS accounts to get consolidated billing.

AWS Organizations is free and available to all AWS customers at no additional charge.

Incorrect Options
AWS Budgets: AWS Budgets service is not used to consolidate AWS billing to centralize the cost. AWS Budgets give the ability to set custom budgets that alert you when your costs or usage

exceed your budgeted or forecasted amount. With AWS Budgets, you can be alerted by email or SNS notification when actual or forecasted cost and usage exceed your budgeted threshold or when your actual RI and Savings Plans' utilization or coverage drops below your desired threshold.

Amazon CloudWatch: AWS CloudWatch service is not used to consolidate AWS billing to centralize the cost. Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights.

AWS Cost and Usage: This is used for reports. AWS Cost and Usage Report contain the most comprehensive cost and usage data available. You can use Cost and Usage Report to publish your AWS billing reports to an Amazon Simple Storage Service (Amazon S3) bucket that you own.

Reference:
https://aws.amazon.com/organizations
https://aws.amazon.com/cloudwatch/
https://aws.amazon.com/aws-cost-management/aws-budgets/
https://aws.amazon.com/aws-cost-management/aws-cost-and-usage-reporting/

Question 36:
You deployed a static web application on S3 and used CloudFront to handle global traffic efficiently. You want to protect an application from common web exploits against OWASP's top 10 security risks. Which of the following services can you use to protect an application from common web exploits?

A. AWS Shield Standard
B. AWS WAF
C. AWS Firewall Manager
D. AWS CloudWatch

Answer: B

Explanation
Correct Option
AWS WAF: The brute force attack on the application's login URL is a layer 7, application-layer attack. The AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies.

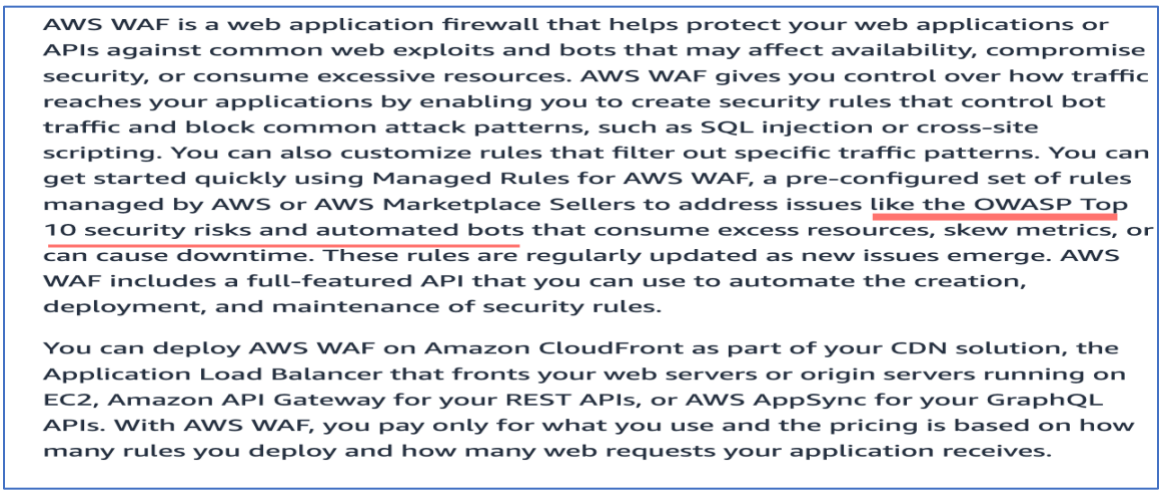

Screenshot from:
https://aws.amazon.com/waf/

Incorrect Options
AWS Shield Standard: AWS Shield Standard doesn't protect from application-layer attacks. It protects from Layer 3 and Layer 4 DDoS attacks for any AWS resource in any AWS Region. It protects against common network attacks such as SYN floods, UDP floods, Reflection attacks, etc., at layer three and layer 4 for any Resources in any AWS Region. All AWS customers get the default AWS Shield protection for most common DDoS attacks, by default with no additional charge.

AWS Firewall Manager: It doesn't protect from application-layer attacks. AWS Firewall Manager, a security management service, allows you to configure and manage firewall rules -- centrally -- across your accounts and applications in AWS Organizations. Using AWS Firewall Manager, you can build firewall rules, create security policies, and enforce them consistently across your entire AWS infrastructure from a central administrator account. When new applications are created, AWS Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules. You can easily implement AWS WAF rules for your application load balancers, API Gateways, and Amazon CloudFront distributions using AWS Firewall Manager.

AWS CloudWatch: It doesn't protect from application-layer attacks. Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights.

Reference:
https://aws.amazon.com/cloudwatch/
https://aws.amazon.com/firewall-manager/
https://aws.amazon.com/waf/
https://aws.amazon.com/shield

Question 37:
Which of the following AWS services don't require a VPC to run? (Select Two)

A. Amazon EC2
B. Amazon RDS
C. Amazon S3
D. Amazon DynamoDB
E. Elastic Load Balancer

Answer: C, D

Explanation
Correct Options
Amazon S3
Amazon DynamoDB

If we look at AWS, by default, VPCs are provided by AWS. And we can also create our VPCs. In VPC, we launch EC2 and services which needs EC2, such as Elastic Load Balancer, NAT Gateway, Amazon RDS, Amazon EMR, Amazon Redshift, Amazon Elasticsearch, and AWS Elastic Beanstalk.

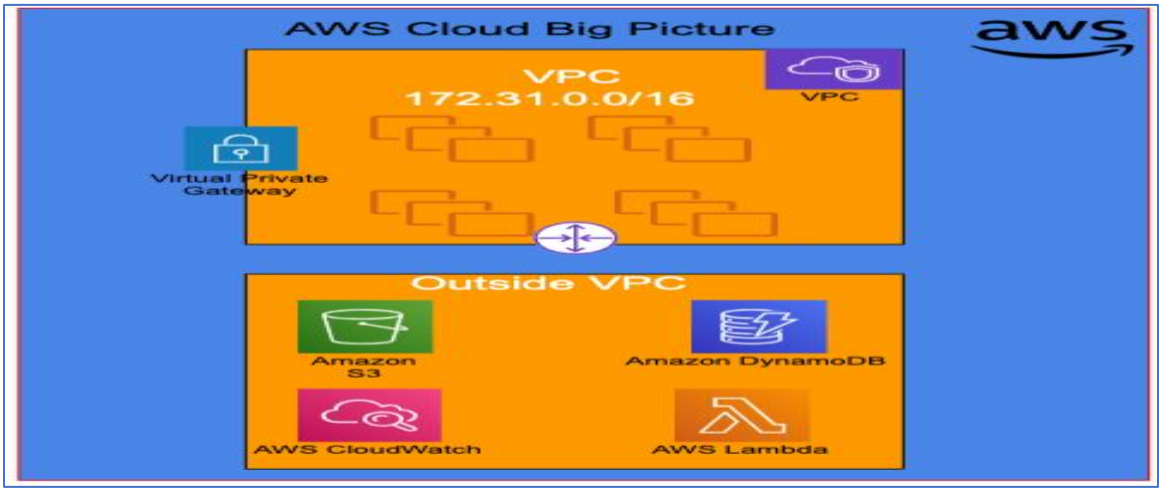

But outside of VPC, the private environment, AWS also provides managed services such as Amazon S3, Amazon DynamoDB, AWS Lambda, CloudWatch, and many other managed services. These services live outside of VPC – in the AWS public address space.

Incorrect Options
Amazon EC2
Amazon RDS
Elastic Load Balancer

These services live in VPC.

Reference:
https://stackoverflow.com/questions/56461741/list-of-aws-services-that-don-t-require-a-vpc-to-run

Question 38:
You have a fleet of 10 EC2 instances running in one AWS Region. You need to apply a patch script, which is stored on GitHub, to all of the EC2 instances, but you are only allowed to apply patch remotely. Which of the following services can you use to apply the patch script on the fleet of all EC2 instances remotely?

A. AWS Config
B. Use the Run command of AWS Systems Manager
C. AWS Web Application Firewall (WAF)
D. Amazon CloudWatch

Answer: B

Explanation
Correct Option
Use Run command of AWS Systems Manager: AWS Systems Manager is a set of fully managed services and capabilities that simplify the management of your Windows and Linux instances regardless of whether they are running on EC2 or on-premises.

AWS Systems Manager allows performing tasks such as collecting system inventory, applying operating system patches, automation of creating Amazon Machine Images (AMIs), and configuring operating systems and applications at scale. This helps accelerate the cloud journey by addressing the shortcomings of the traditional system management approach. It provides a flexible and easy-to-use automation-focused approach for both traditional and cloud-based workloads.

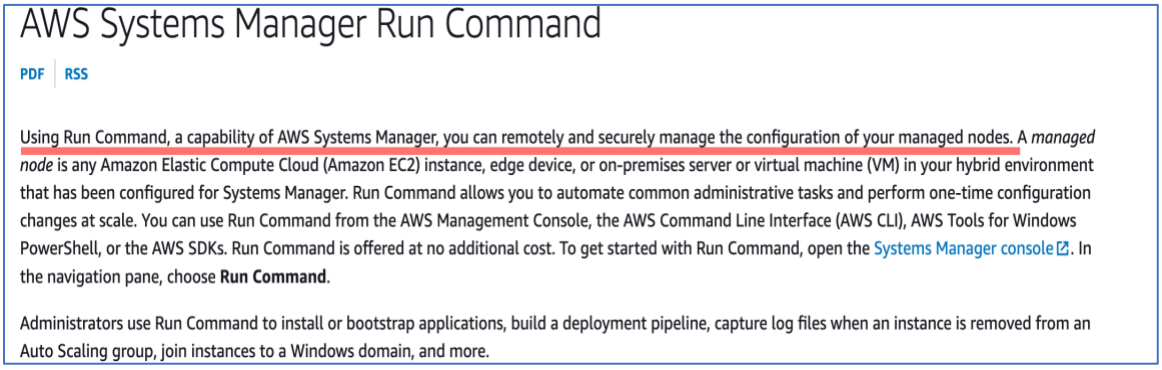
AWS Systems Manager Run Command

PDF | RSS

Using Run Command, a capability of AWS Systems Manager, you can remotely and securely manage the configuration of your managed nodes. A *managed node* is any Amazon Elastic Compute Cloud (Amazon EC2) instance, edge device, or on-premises server or virtual machine (VM) in your hybrid environment that has been configured for Systems Manager. Run Command allows you to automate common administrative tasks and perform one-time configuration changes at scale. You can use Run Command from the AWS Management Console, the AWS Command Line Interface (AWS CLI), AWS Tools for Windows PowerShell, or the AWS SDKs. Run Command is offered at no additional cost. To get started with Run Command, open the Systems Manager console. In the navigation pane, choose **Run Command**.

Administrators use Run Command to install or bootstrap applications, build a deployment pipeline, capture log files when an instance is removed from an Auto Scaling group, join instances to a Windows domain, and more.

Screenshot from:
https://docs.aws.amazon.com/systems-manager/latest/userguide/execute-remote-commands.html

Incorrect Options
AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource

configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations. This cannot be used to apply patches remotely.

AWS Web Application Firewall (WAF): AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies. This cannot be used to apply patches remotely.

Amazon CloudWatch: Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights. This cannot be used to apply patches remotely.

Reference:
https://docs.aws.amazon.com/systems-manager/latest/userguide/execute-remote-commands.html
https://aws.amazon.com/cloudwatch/
https://aws.amazon.com/waf
https://aws.amazon.com/config/

Question 39:
You have an application that stores information in DynamoDB. The application also uses S3 to store files and images. You are designing a feature for this application where if a user uploads an image, the image thumbnail should be displayed quickly. Which of the following AWS services will you help you implement this feature cost-effectively?

A. Amazon EFS
B. Amazon SageMaker
C. AWS Elastic Beanstalk
D. AWS Lambda

Answer: D

Explanation
Correct Option
AWS Lambda: You can use AWS Lambda for an event-driven programming model to execute an AWS Lambda function to process the image and create its thumbnail when an image is uploaded to S3.

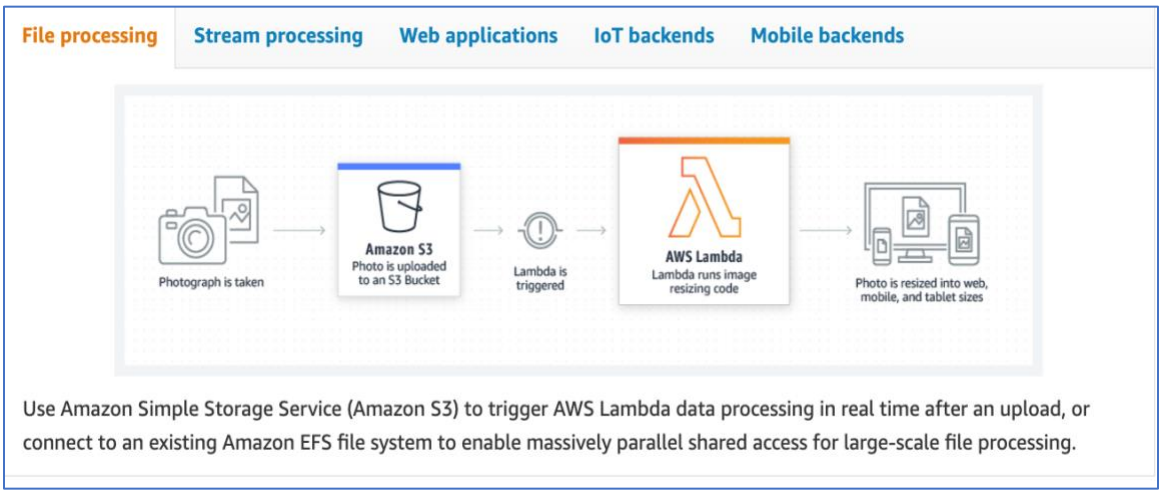

Screenshot from:
https://aws.amazon.com/lambda/

AWS Lambda enables you to execute code without provisioning or managing servers. Instead, you are charged based on the number of requests for your functions and the duration it takes for your code to execute. AWS Lambda executes code in response to events such as object uploads to S3, updates to DynamoDB tables, or other events such as website clicks.

Incorrect Options

Amazon EFS: Amazon EFS is a cloud-native, serverless, and fully managed file system that is accessible from Linux instances via the NFS protocol. It is built to scale on-demand to petabytes without disrupting applications. The EFS scales out and scales in automatically as you add and remove files. As a result, you don't need to provision and manage capacity to accommodate growth. In addition, Amazon EFS is designed to provide massively parallel shared access to thousands of EC2 instances. Thus, enabling your applications to achieve high aggregate throughput and IOPS with consistent low latency.

Amazon SageMaker: Amazon SageMaker enables business analysts, data scientists, and ML engineers to build, train, and deploy machine learning (ML) models for different use cases with fully managed infrastructure, tools, and workflows. Business analysts can use SageMaker to make ML predictions using a visual interface with SageMaker Canvas. Data scientists can use SageMaker to prepare data and build, train, and deploy models with SageMaker Studio. ML engineers can use SageMaker to deploy and manage models at scale with SageMaker with MLOps.

AWS Elastic Beanstalk: AWS Elastic Beanstalk is a PaaS (Platform-as-a-Service) type of AWS service that is used for deploying and scaling web applications. This AWS service takes your application code and deploys it while provisioning the compute resources required for the application to run.

AWS Elastic Beanstalk helps deploy and scale web applications that are developed with Java, NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar web servers such as Apache, Nginx, Passenger, and IIS. You simply upload the code, and Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto-scaling to

application health monitoring. Elastic Beanstalk qualifies as a Platform-as-a-Service cloud computing type.

Reference:
https://aws.amazon.com/lambda/
https://aws.amazon.com/efs/
https://aws.amazon.com/sagemaker/
https://aws.amazon.com/elasticbeanstalk/

Question 40:
You have two database servers on EC2 instances in a private subnet. You would like these instances to connect to the Internet so that they can download the latest patches. Which of the following can you use to allow EC2 instances in the private subnet to connect to the Internet?

A. AWS Direct Connect
B. NAT Gateway
C. Customer Gateway
D. Transit Gateway

Answer: B

Explanation
Correct Option
NAT Gateway: A NAT gateway is a Network Address Translation (NAT) service. You can use a NAT gateway so that instances in a private subnet can connect to services outside of your VPC. However, the external services cannot initiate a connection with instances inside the private subnet. The reason is NAT Gateway is a one-way street to connect. What it means inbound traffic from the Internet to the NAT Gateway is not allowed.

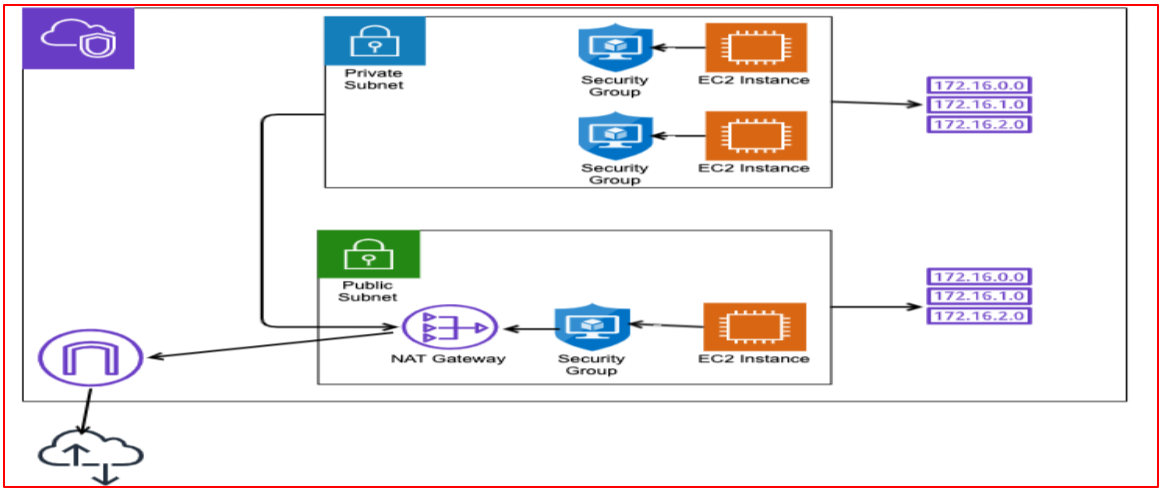

The above diagram shows setting up NAT Gateway in the public subnet to access the Internet from the hosts in the private subnet. If you notice, traffic from the private subnet goes to the NAT Gateway. And from there, then, it is sent to the Internet Gateway connected to the Internet.

Incorrect Options

AWS Direct Connect: AWS Direct Connect is a physical connection between on-premises to the AWS cloud. It uses AWS infrastructure. It is used to extend the on-premises connection to the AWS Cloud. AWS Direct Connect helps make it easy to set up a dedicated network connection from your on-premises data center to the AWS cloud, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. The AWS Direct Connect is a physical connection. However, AWS Direct Connect does not encrypt your traffic in transit.

Customer Gateway: Customer Gateway such as Router is used at the customer side in setting up connectivity between on-premises and AWS VPC. The Customer Gateway could be your Router, firewall, or other things that support IPSec in your on-premises environment. The Customer Gateway resides at the on-premises end, and it is used in an AWS Site-to-Site VPN connection.

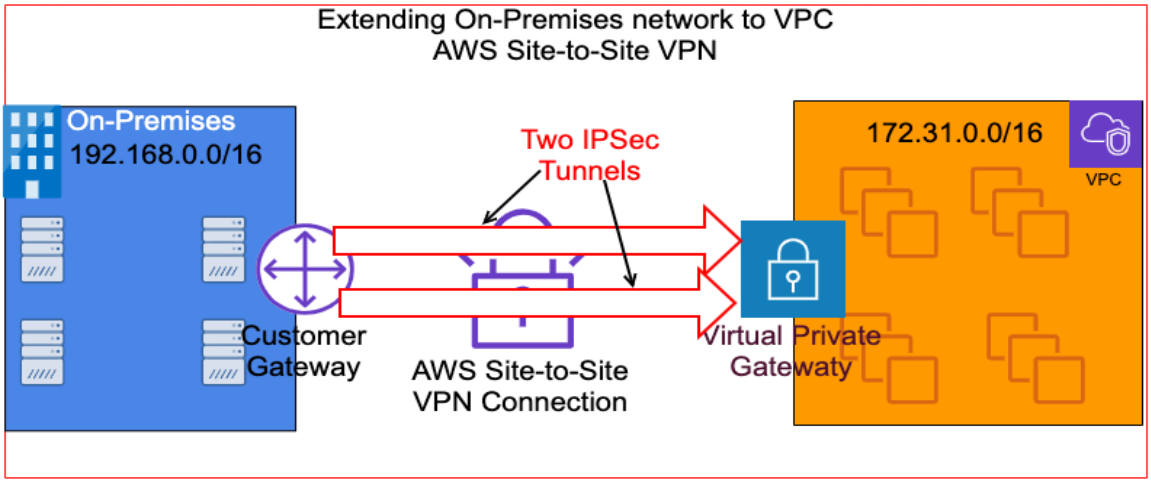

Transit Gateway: Transit Gateway is used to connect VPCs. Transit Gateway is distributed managed routing service that you deploy into a Region. So, you can connect and attach VPCs in the same Region to your Transit Gateway. And then, you can allow any-to-any connection between VPCs from a routing perspective.

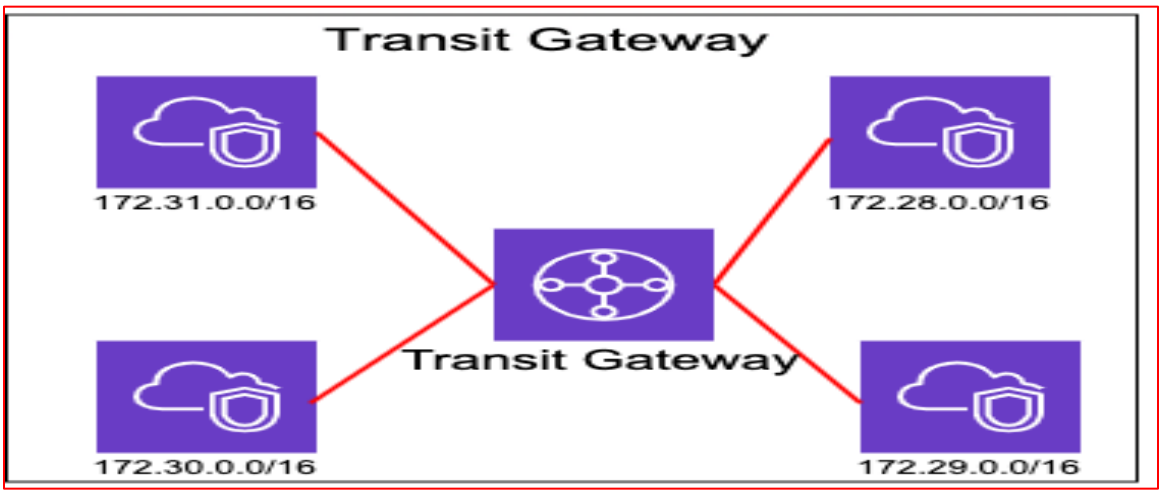

Reference:
https://docs.aws.amazon.com/vpc/latest/userguide/working-with-vpcs.html
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html
https://aws.amazon.com/directconnect/
https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html

Question 41:
You have a use case where you need to connect your on-premises IT infrastructure to the multiple VPC using a consistent high bandwidth connection. Which of the following options would you recommend?

A. AWS Site-to-Site VPN
B. AWS Direct Connect with Direct Connect Gateway
C. Virtual Private Gateway
D. Customer Gateway

Answer: B

Explanation
Correct Option
AWS Direct Connect with Direct Connect Gateway: The AWS Direct Connect is a physical connection. The AWS Direct Connect provides a more predictable and consistent experience for on-premises connectivity to the AWS VPC.

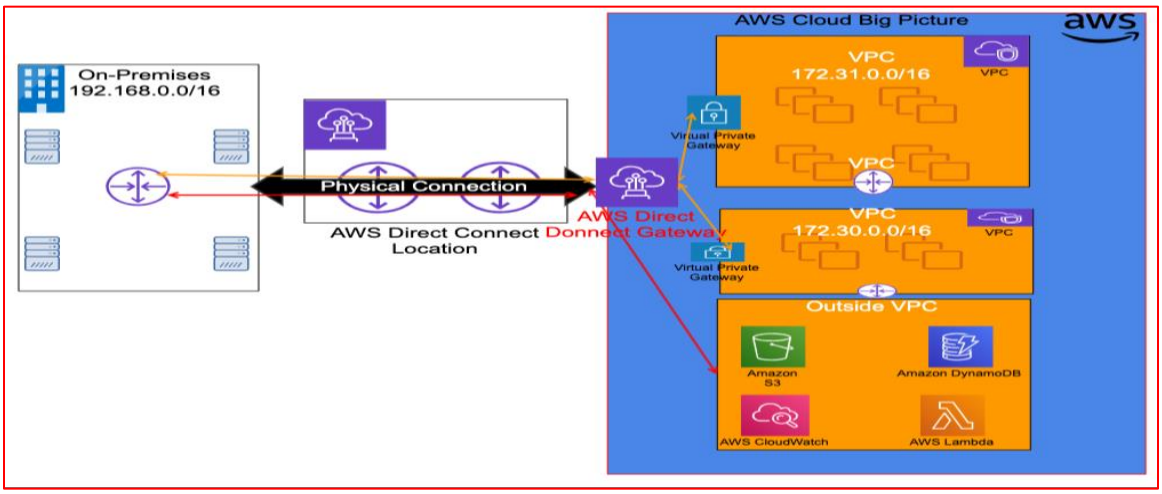

With AWS Direct Connect you can connect your on-premises IT environment to the AWS VPC. However, if you connect more than one VPC from your on-premises environment, you can use AWS Direct Connect Gateway. That way, you can connect up to 10 VPCs that reside in the same Region or in the different AWS Regions across the world (except China).

Incorrect Options

AWS Site-to-Site VPN: AWS Site-to-Site VPN can only connect your IT on-premises to one VPC. Also, AWS Site-to-Site VPN uses the Internet whether connection bandwidth might not be consistent because of the inherent nature of how traffic on the Internet flows.

Virtual Private Gateway: The Virtual Private Gateway is not a full connectivity option between the on-premises and AWS VPC. A virtual gateway allows resources that are outside of your VPC to communicate to resources that are inside of your VPC.

Customer Gateway: The customer gateway is not a full connectivity option between the on-premises and AWS VPC. It is at the customer end of the connectivity. A customer gateway allows resources that are outside of the on-premises environment to communicate to resources that are inside of the on-premises environment.

Reference:

https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

https://blog.opstree.com/2020/09/01/why-we-should-use-transit-direct-connect-gateways/

Question 42:

You are planning to run a build job. The job is of predictable nature in terms of compute resource requirements. You have a Reserved EC2 Instance available. Which of the following is a cost-effective to run the job using the Docker container on AWS?

A. ECS on EC2
B. AWS Fargate
C. AWS Lambda
D. AWS Lambda or AWS Fargate

Answer: A

Explanation
Correct Option
ECS on EC2: Since you already have a Reserved EC2 Instance available, running an ECS container on EC2 is the most cost-effective option. AWS ECS is a container management service that facilitates the management of Docker containers. In order to launch compute infrastructure, you have two options: ECS on EC2 or Fargate.

When using EC2 to launch ECS, you will have to set up EC2 instances separately, such as installing the ECS agent and setting up Auto Scaling. However, when using AWS Fargate to run containers, you don't need to deal with Amazon EC2 instances to manage servers or clusters. For example, you no longer have to provision, configure, or scale clusters of VMs to run containers.

When you run your ECS tasks and services with the Fargate launch type, you package your application in containers, specify the OS, CPU, and memory requirements, configure networking, define IAM policies, and launch the application. Each Fargate task has its own isolation boundary and does not share the underlying kernel, CPU resources, memory resources, or elastic network interface with another task.

When to use ECS on EC2 and when to use Fargate
If you would like to have explicit control of the running container, ECS on EC2 is a better option. If you would like to have flexibility, Fargate is the way to go. If you have spare EC2 instances, for example, Reserved EC2 instances, you can save on cost when using ECS on EC2. If you don't have time to set up an EC2 instance, for example, installing an ECS agent and setting up Auto Scaling, Fargate is the quickest option. In Fargate, you pay based on the provisioned compute resources.

When using EC2, you specify EC2 infrastructure. Typically, you launch one ECS task on an EC2 machine. However, depending on your task requirements and available capacity on the EC2 instance, you can run more than one task on the EC2 instance.

EC2 on ECS, you have the self-managed infrastructure; in the case of Fargate, you have AWS-managed infrastructure.

Incorrect Options
AWS Fargate
AWS Lambda
AWS Lambda or AWS Fargate

These are relatively expensive options for this use case.

Reference:
https://docs.aws.amazon.com/eks/latest/userguide/fargate.html
https://aws.amazon.com/blogs/aws/cloud-container-management/

Question 43:
Your company is involved in modernization projects for many applications to migrate them to the AWS cloud. You have been asked to build a design solution for one of the applications in such a way as to use AWS services wherever possible so that the application can be more cloud-native. The application that you have been assigned to modernize uses LDAP for authentication. Your use case is to replace the LDAP with an AWS service in your new design. Which of the following services can you use for this use case?

A. Amazon DynamoDB
B. Amazon Cognito
C. Amazon Cloud Directory
D. AWS Secrets Manager

Answer: C

Explanation
Correct Option
Amazon Cloud Directory: Amazon Cloud Directory is like Microsoft AD (Active Directory) or LDAP in terms of its core use case. However, Amazon Cloud Directory is a more advanced solution on the cloud for hierarchical data. For example, traditional directory solutions, such as Microsoft AD or LDAP, are limited to a single hierarchy, which is single. Cloud Directory, however, offers the flexibility to create directories with hierarchies that span multiple dimensions. For example, you can create an organizational chart that can be navigated through separate hierarchies for reporting structure, location, and cost center.

Amazon Cloud Directory has extensive scalability -- it can automatically scale to hundreds of millions of objects. It provides an extensible schema that can be shared with multiple applications. It is a fully-managed service – for example, it eliminates time-consuming and expensive administrative tasks, such as scaling infrastructure and managing servers.

Working with Amazon Cloud Directory is very easy -- first, you define the schema, then create a directory. After that, you can populate the directory by making calls to the Cloud Directory API. It is integrated with AWS CloudTrail. AWS CloudTrail can help you log the date, time, and identity of users who accesses your directory data. With resource tagging, you can tag your directories and schemas to better track and manage resources.

With regards to its use cases, you can use it to efficiently organize hierarchies of data across multiple dimensions and search your directory for objects and relationships.

Incorrect Options
Amazon DynamoDB: DynamoDB is an AWS NoSQL database. It delivers single-digit millisecond performance at any scale. It's a fully managed, multi-region, multi-active, durable database with built-in security, backup and restores, and in-memory caching for internet-scale applications. It can handle more than 10 trillion requests per day and support peaks of more than 20 million requests per second.

It is optimized for performance over consistency. Using this model, DynamoDB will look for the nearest and most available location to fulfill the read request during a read. DynamoDB keeps

three replicas of the table in a geographically distributed fashion. That being the case, there is a possibility that during the read, the data in the nearest and available replica may not have been updated with the latest. The DynamoDB database has an eventual consistency model.

Amazon DynamoDB cannot be used for this use case to replace LDAP functionality.

Amazon Cognito: Amazon Cognito is a simple user identity service. It lets you add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. With Amazon Cognito, you also can authenticate users through social identity providers such as Facebook, Twitter, or Amazon, with SAML identity solutions, or by using your identity system.

Amazon Cognito cannot be used for this use case to replace LDAP functionality.

AWS Secrets Manager: AWS Secrets Manager, a secrets management service, protects your AWS account's access to applications, services, and IT resources. The service helps you easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs. This call eliminates the need to hardcode sensitive information in plain text. The AWS Secrets Manager also offers secret rotation with built-in integration for Amazon RDS, Redshift, and Amazon DocumentDB.

AWS Secrets Manager cannot be used for this use case to replace LDAP functionality.

Reference:
https://aws.amazon.com/cloud-directory/
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html
https://aws.amazon.com/cognito/
https://aws.amazon.com/secrets-manager/

Question 44:
You have concerns about a possible DDoS attack on your application and are interested to find out bad actors if that happens. Which of the following AWS services can you use to address the issue?

A. AWS Shield Advanced
B. AWS Shield Standard
C. Amazon CloudWatch
D. Amazon Cognito

Answer: A

Explanation
Correct Option:
AWS Shield Advanced: AWS Shield prevents DDoS (Distributed Denial of Service) attacks on AWS resources for both global and regional resources. AWS Shield is available in two levels: AWS Shield Standard and AWS Shield Advanced. AWS Shield Standard is free and open to all customers. AWS Shield Standard protects all AWS customers against common and most frequently occurring infrastructure (layer 3 and 4) attacks like SYN/UDP floods, reflection attacks, and others to support the high availability of your applications on AWS. AWS Shield

Advanced is a paid service. Being a paid service, it provides every feature that AWS Shield Standard. Additionally, it offers additional protection such as visibility and monitoring of DDoS attacks, 24x7 support from the AWS DDoS support team, and AWS WAF subscription at no cost.

Shield Advanced gives you access to advanced, real-time metrics and reports for extensive visibility into events and attacks on your protected AWS resources. You can access this information through the Shield Advanced API and console, and through Amazon CloudWatch metrics.

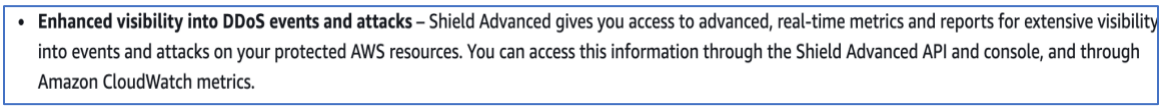

Screenshot from:
https://docs.aws.amazon.com/waf/latest/developerguide/ddos-advanced-summary-capabilities.html

Incorrect Options
AWS Shield Standard: All AWS customers get the default AWS Shield protection for most common DDoS attacks, by default with no additional charge. It protects from Layer 3 and Layer 4 DDoS attacks for any AWS resource in any AWS Region. It protects against common network attacks such as SYN floods, UDP floods, Reflection attacks, etc., at layer three and layer 4 for any Resources in any AWS Region. You cannot get detail about the DDoS attack on your application, such as about bad actors if that happens. This provides basic protection at layer 3 and layer 4.

Amazon CloudWatch: Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights. There are many components or features of the CloudWatch service. Amazon CloudWatch is an excellent service for building Resilient systems. If you are looking for resource performance monitoring, events, and alerts – CloudWatch is a go-to service. For example, you can configure a CloudWatch alarm that sends an email message using Amazon SNS when CPU utilization crosses the threshold of 80%. It cannot be used for the use case in the question.

Amazon Cognito: Amazon Cognito is a simple user identity service. It lets you add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. With Amazon Cognito, you also can authenticate users through social identity providers such as Facebook, Twitter, or Amazon, with SAML identity solutions, or by using your identity system. It cannot be used for the use case in the question.

Reference:
https://aws.amazon.com/cognito/
https://aws.amazon.com/shield/
https://aws.amazon.com/cloudwatch/

Question 45:

You have a fleet of 10 EC2 instances running in one AWS Region. You need to turn off SSH on all the instances, but you are only allowed to do it remotely. Which of the following services can you use to turn off SSH on all the EC2 instances remotely?

A. AWS Config
B. Use State Manager of AWS Systems Manager
C. AWS Web Application Firewall (WAF)
D. Amazon CloudWatch

Answer: B

Explanation
Correct Option
Use State Manager of AWS Systems Manager: AWS Systems Manager is a set of fully managed services and capabilities that simplify the management of your Windows and Linux instances regardless of whether they are running on EC2 or on-premises. AWS Systems Manager allows performing tasks such as collecting system inventory, applying operating system patches, automation of creating Amazon Machine Images (AMIs), and configuring operating systems and applications at scale.

How can State Manager benefit my organization?

By using pre-configured Systems Manager documents (SSM documents), State Manager offers the following benefits for managing your nodes:

- Bootstrap nodes with specific software at start-up.
- Download and update agents on a defined schedule, including the SSM Agent.
- Configure network settings.
- Join nodes to a Microsoft Active Directory domain.
- Patch nodes with software updates throughout their lifecycle.
- Run scripts on Linux, macOS, and Windows managed nodes throughout their lifecycle.

To manage configuration drift across other AWS resources, you can use Automation, a capability of Systems Manager, with State Manager to perform the following types of tasks:

- Attach a Systems Manager role to Amazon Elastic Compute Cloud (Amazon EC2) instances to make them *managed nodes*.
- Enforce desired ingress and egress rules for a security group.
- Create or delete Amazon DynamoDB backups.
- Create or delete Amazon Elastic Block Store (Amazon EBS) snapshots.
- Turn off read and write permissions on Amazon Simple Storage Service (Amazon S3) buckets.
- Start, restart, or stop managed nodes and Amazon Relational Database Service (Amazon RDS) instances.
- Apply patches to Linux, macOS, and Window AMIs.

Screenshot from:
https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-state.html

Incorrect Options
AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations. AWS Config service cannot be used for the use case in the question.

AWS Web Application Firewall (WAF): AWS Web Application Firewall (WAF), which is a layer seven firewall, helps protect web applications and APIs against common web exploits and bots that may affect applications' availability and compromise their security. Sometimes they may consume excessive resources to impact the performance of the overall system. AWS WAF helps

protect web applications from attacks by allowing you to configure rules that will enable, block, or monitor web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection, and cross-site scripting. In addition, you can use the IP address-based match rule to block specific geographies. AWS WAF service cannot be used for the use case in the question.

Amazon CloudWatch: Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights. There are many components or features of the CloudWatch service. Amazon CloudWatch is an excellent service for building Resilient systems. If you are looking for resource performance monitoring, events, and alerts – CloudWatch is a go-to service. For example, you can configure a CloudWatch alarm that sends an email message using Amazon SNS when CPU utilization crosses the threshold of 80%. Amazon CloudWatch service cannot be used for the use case in the question.

Reference:
https://aws.amazon.com/config/
https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-state.html
https://aws.amazon.com/cloudwatch/
https://aws.amazon.com/waf/

Question 46:
Which of the following options is true related to the protection provided by AWS Shield Standard?

A. It defends against the most common, frequently occurring Network, Transport, and Application layer DDoS attacks.
B. It defends against the most common, frequently occurring Network, Transport layer DDoS attacks.
C. It defends against the most common, frequently occurring DDoS attacks only at the Network layer.
D. It defends against the most common, frequently occurring DDoS attacks only at the Application layer.

Answer: B

Explanation
Correct Option
It defends against the most common, frequently occurring Network, Transport layer DDoS attacks.

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield - Standard and Advanced.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your web site or applications. When you use AWS Shield Standard with Amazon CloudFront and Amazon Route 53, you receive comprehensive availability protection against all known infrastructure (Layer 3 and 4) attacks.

Screenshot from:
https://aws.amazon.com/shield/

Incorrect Options
All other options are incorrect.

Reference:
https://aws.amazon.com/shield/

Question 47:
You have a use case where you need to connect your on-premises IT infrastructure to the multiple 250 VPCs using a consistent high bandwidth connection. Which of the following options would you recommend?

A. AWS Site-to-Site VPN
B. AWS Direct Connect with Transit Gateway
C. Virtual Private Gateway
D. Customer Gateway

Answer: B

Explanation
Correct Option
AWS Direct Connect with Transit Gateway: Using Direct Connect, you can connect all VPCs using one logical connection. So, for example, if we add Transit Gateway with Direct Connect Gateway, just by using one physical link from on-premises to AWS, we can reach up to 5000 VPC using Transit Gateway. A transit gateway allows VPCs to intercommunicate as long as they are connected to the same transit gateway.

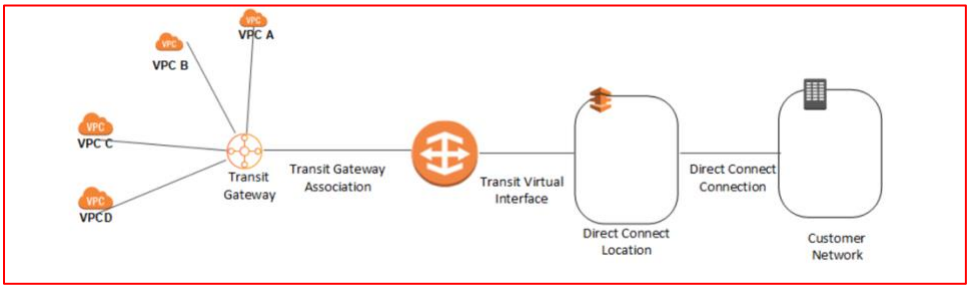

The above screenshot is from:

https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

It shows that how the Direct Connect gateway using Transit Gateway enables you to create a single connection to your Direct Connect connection that all of your VPCs can use, and provides connectivity to on-premises as well.

Incorrect Options
AWS Site-to-Site VPN: AWS Site-to-Site VPN can only connect your IT on-premises to one VPC. Also, AWS Site-to-Site VPN uses the Internet whether connection bandwidth might not be consistent because of the inherent nature of how traffic on the Internet flows.
Virtual Private Gateway: The Virtual Private Gateway is not a full connectivity option between the on-premises and AWS VPC. A virtual gateway allows resources that are outside of your VPC to communicate to resources that are inside of your VPC.
Customer Gateway: The customer gateway is not a full connectivity option between the on-premises and AWS VPC. It is at the customer end of the connectivity. A customer gateway allows resources that are outside of the on-premises environment to communicate to resources that are inside of the on-premises environment.

Reference:
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html
https://blog.opstree.com/2020/09/01/why-we-should-use-transit-direct-connect-gateways/

Question 48:
You are working as a lead software engineer. You have been asked to translate large volumes of text for analysis quickly. Which of the following AWS service can you use for this use case?

A. Amazon Rekognition
B. Amazon Transcribe
C. Amazon Polly
D. Amazon Translate

Answer: D

Explanation
Correct Option
Amazon Translate: Amazon Translate is a neural machine translation that delivers fast, high-quality, affordable, and customizable language translation. Amazon Translate differs from traditional statistical and rule-based translation algorithms. Instead, it uses neural machine translation, which uses deep learning models to provide more accurate and natural-sounding translations. As a result, the Amazon Translate service can help you localize content such as websites and applications for your different types of diverse users. In addition, it can quickly translate large volumes of text for analysis and efficiently enable cross-lingual communication between users.

Incorrect Options

Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.

Amazon Transcribe: Amazon Transcribe service helps you quickly add high-quality speech-to-text capabilities to your applications. For example, you can quickly extract actionable insights from customer conversations. In another use case, content producers can use this service to convert audio and video assets into fully searchable content automatically. For example, you can create subtitles for your broadcast content to increase accessibility and improve customer experience. Amazon Transcribe service can be used in the medical field as well. For example, medical doctors and practitioners can use Amazon Transcribe Medical to quickly document clinical conversations into electronic health record (EHR) systems for analysis.

Amazon Polly: Amazon Polly is a service that turns text into lifelike speech. Amazon Polly's Text-to-Speech (TTS) uses advanced deep learning techniques to synthesize natural-sounding human speech. This natural-sounding human speech helps developers build speech-enabled products -- applications that can talk. Using machine learning, Amazon Polly offers Neural Text-to-Speech (NTTS) voices, delivering advanced improvements in speech quality. Amazon Polly Brand Voice can also create a custom NTTS voice for your organization's exclusive use.

Reference:

https://aws.amazon.com/translate/
https://aws.amazon.com/polly/
https://aws.amazon.com/transcribe/
https://aws.amazon.com/rekognition/

Question 49:

Your company is required to maintain a history of all changes to EC2 to maintain compliance. Which of the following services will you use to record the history of changes?

A. Amazon CloudWatch
B. AWS CloudTrail
C. AWS Config
D. AWS Logs

Answer: C

Explanation

Correct Option

AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations.

For example, if you have an EC2 instance of m5 type, you change it to m5a type. If AWS Config is not turned on, you have lost the record of that change. But if you have turned on AWS Config, you can find the record of that change. In addition, AWS Config also has many pre-built rules that you can use to run. For example, if you would like to know if an Elastic IP Address is attached to an EC2 instance, you can create that rule. That way, if you either run it or let it run on a schedule. The run result will let you know if the Elastic IP Address is associated with any EC2 instance or not. This kind of rule setup helps in compliance checkups. Of course, you can create your custom rule as well.

The best way to remember about AWS Config: if you are looking for resource change history, audit, or compliance -- think AWS Config.

You can do the following using AWS Config:

- Evaluate your AWS resource configurations for the settings you want.
- Get a snapshot of the current settings or configurations of the resources within your AWS account.
- Get the historical configurations of resources.
- Receive notification about whenever a resource is created, modified, or deleted in your AWS account.
- View relationships among resources; for example, you might want to find all resources that use a particular security group.

While AWS Config helps you answer questions like - what did my AWS resource look like at a point in time? You can use AWS CloudTrail to answer -- who made an API call to modify this resource?

How AWS Config Works:

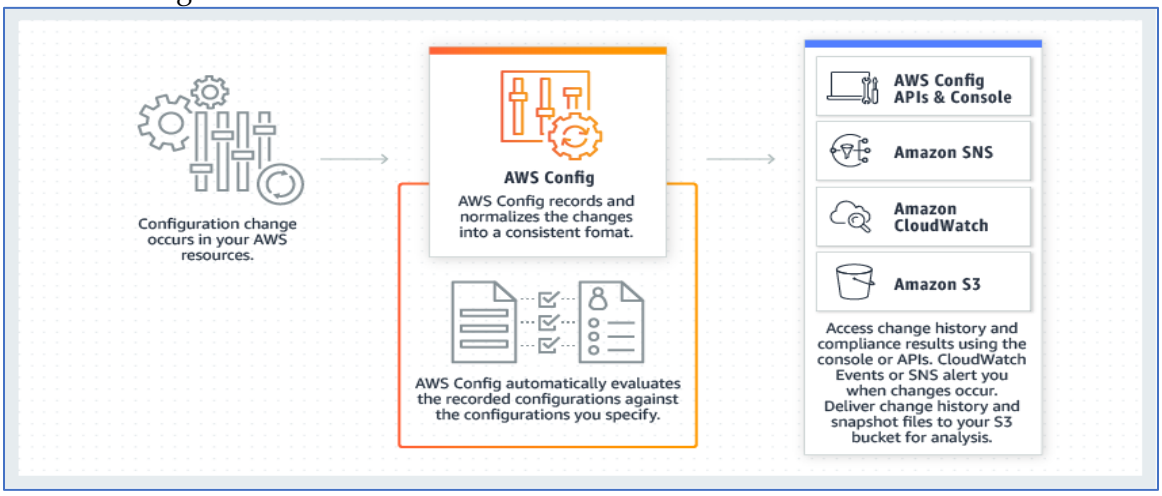

Incorrect Options

Amazon CloudWatch: Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights. There are many components or features of the CloudWatch service. Amazon CloudWatch is an excellent service for building Resilient systems.

If you are looking for resource performance monitoring, events, and alerts – CloudWatch is a go-to service. For example, you can configure a CloudWatch alarm that sends an email message using Amazon SNS when CPU utilization crosses the threshold of 80%.
AWS CloudTrail: AWS CloudTrail provides auditing, security monitoring, and operational troubleshooting by tracking user activity and API usage. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure, giving you control over storage, analysis, and remediation actions. In addition, CloudTrail provides the event history of your AWS account activity, including actions taken through AWS Management.
AWS Logs: This is a made-up option. This is a distraction.

Reference:
https://aws.amazon.com/config/
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html
https://aws.amazon.com/cloudwatch/

Question 50:
Which of the following is not true about VPC Sharing?

A. If you run out of IPv4 addresses, you can create VPC sharing to utilize the resources more efficiently.
B. The only owner of the VPC can change the configuration or setup, such as creating subnets, setting up all the route tables, setting up NACLs, etc.
C. The owner of the VPC can share the subnet with multiple accounts.
D. There will be one billing account for all the resources created in a shared subnet.

Answer: D

Explanation
Correct Option

There will be one billing account for all the resources created in the shared subnet:

- VPC Sharing helps in preserving IP space. If you run out of IPv4 addresses, you can create VPC sharing to utilize the resources more efficiently.
- No VPC Peering is required. If you have fewer VPCs, you have little connectivity between different VPCs.
- Separation of duties. A central team can create and manage VPCs. You will have separation of duties where owner creates the VPCs and users cannot change it.
- Billing and Security. Please continue to have a separation of accounts and billed for the resources they create.

Incorrect Options
All other options are true.

Reference:
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing.html

Question 51:
You are working as a lead software engineer for a media company. You have been asked to quickly document clinical conversations into electronic health record (EHR) systems for analysis. Which of the following AWS service can you use for this use case?

A. Amazon Rekognition
B. Amazon Transcribe
C. Amazon Personalize
D. Amazon Kendra

Answer: B

Explanation

Correct Option

Amazon Transcribe: Amazon Transcribe service helps you quickly add high-quality speech-to-text capabilities to your applications. For example, you can quickly extract actionable insights from customer conversations. In another use case, content producers can use this service to convert audio and video assets into fully searchable content automatically. For example, you can create subtitles for your broadcast content to increase accessibility and improve customer experience. Amazon Transcribe service can be used in the medical field as well. For example, medical doctors and practitioners can use Amazon Transcribe Medical to quickly document clinical conversations into electronic health record (EHR) systems for analysis.

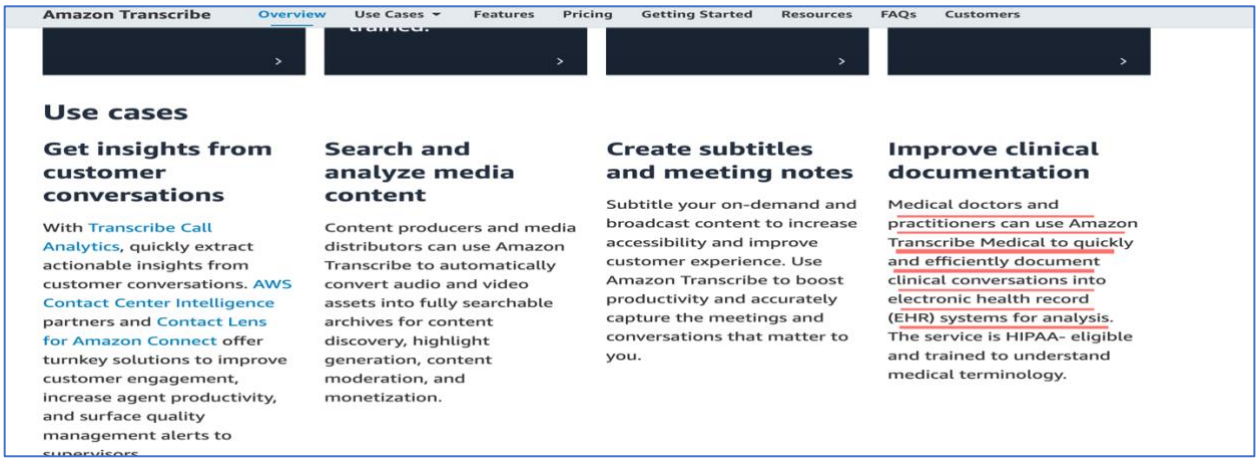

Screenshot from:
https://aws.amazon.com/transcribe/

Incorrect Options

Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.

Amazon Personalize: Amazon Personalize is a fully managed ML service for real-time personalized recommendations. For example, you can use this service for product recommendations, personalized product re-ranking, and customized direct marketing. Amazon Personalize provisions the infrastructure and manages the entire ML pipeline, including pre-processing, features extraction, and applying the best algorithm. Additionally, it then trains, optimizes, and deploys the model. You just need to call API endpoints for the deployed model. All data is encrypted, private, and secure, and is only used to create recommendations for your users.
Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

Reference:
https://aws.amazon.com/transcribe/
https://aws.amazon.com/rekognition/
https://aws.amazon.com/kendra/
https://aws.amazon.com/personalize/

Question 52:
You are working in a DevOps group of your company. There are concerns about AWS cost, and your group has been asked to make sure all Elastic IP Addresses must be used otherwise released. Which of the following services will you use to find out if each Elastic IP Address is associated with an EC2 instance or not?

A. Amazon CloudWatch
B. AWS CloudTrail
C. AWS Config
D. AWS X-Ray

Answer: C

Explanation
Correct Option
AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations.

For example, if you have an EC2 instance of m5 type, you change it to m5a type. If AWS Config is not turned on, you have lost the record of that change. But if you have turned on AWS Config, you can find the record of that change. In addition, AWS Config also has many pre-built rules that you can use to run. For example, if you would like to know if an Elastic IP Address is attached to an EC2 instance, you can create that rule. That way, if you either run it or let it run on a schedule. The run result will let you know if the Elastic IP Address is associated with any EC2

instance or not. This kind of rule setup helps in compliance checkups. Of course, you can create your custom rule as well.

The best way to remember about AWS Config: if you are looking for resource change history, audit, or compliance -- think AWS Config.

You can do the following using AWS Config:

- Evaluate your AWS resource configurations for the settings you want.
- Get a snapshot of the current settings or configurations of the resources within your AWS account.
- Get the historical configurations of resources.
- Receive a notification about whenever a resource is created, modified, or deleted in your AWS account.
- View relationships among resources; for example, you might want to find all resources that use a particular security group.

While AWS Config helps you answer questions like - what did my AWS resource look like at a point in time? You can use AWS CloudTrail to answer -- who made an API call to modify this resource?

How AWS Config Works:

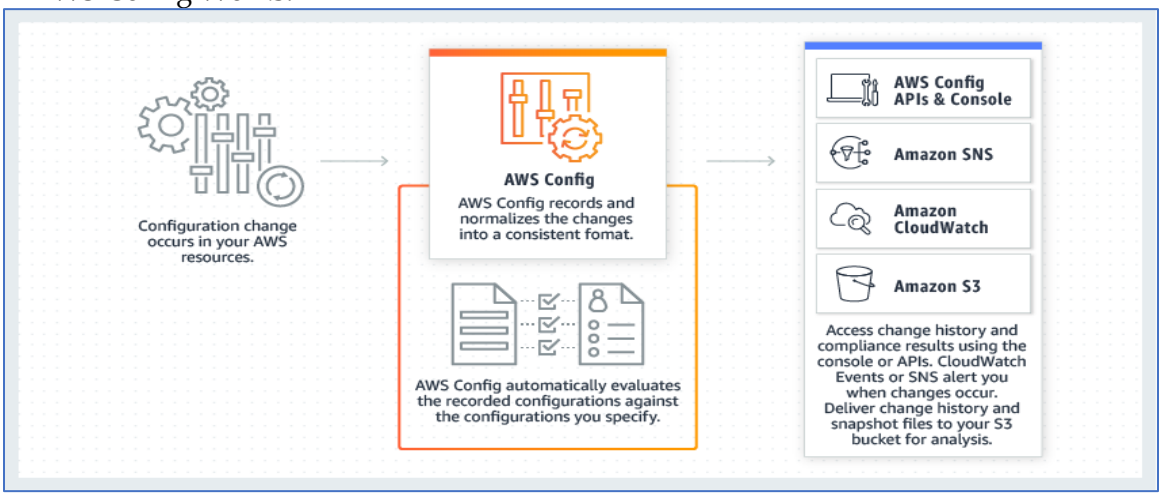

Exam Alert:
You might find questions asking you to select from AWS CloudWatch, AWS CloudTrail, or AWS Config. Just remember the following:

- Resource performance monitoring, events, and alerts -- AWS CloudWatch.
- API calls, account-specific activity, and audit – AWS CloudTrail.
- Resource-specific change history, audit, and compliance -- AWS Config.

Incorrect Options

Amazon CloudWatch: Amazon CloudWatch is an AWS monitoring and observability service. AWS CloudWatch service helps you monitor your applications and resource optimizations, respond to system-wide performance changes and provide a unified view of operation health by providing data and actionable insights. This service cannot be used for the use case mentioned in the question.

AWS CloudTrail: AWS CloudTrail provides auditing, security monitoring, and operational troubleshooting by tracking user activity and API usage. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure, giving you control over storage, analysis, and remediation actions. In addition, CloudTrail provides the event history of your AWS account activity, including actions taken through AWS Management. This service cannot be used for the use case mentioned in the question.

AWS X-Ray: AWS X-Ray helps developers analyze and debug distributed applications, such as applications developed using microservices. AWS X-Ray can help you understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. In addition, x-Ray provides an end-to-end view of requests traveling through your application and shows a map of your application's underlying components. This service cannot be used for the use case mentioned in the question.

Reference:

https://docs.aws.amazon.com/awscloudtrail
https://aws.amazon.com/xray/
https://aws.amazon.com/config/
https://aws.amazon.com/cloudwatch/

Question 53:

You are working as a lead software engineer. You have been asked to implement a feature in an existing application to flag suspicious online payment transactions before processing payments and fulfilling orders. Which of the following AWS service can you use for this use case?

A. Amazon Rekognition
B. Amazon Fraud Detector
C. Amazon Kendra
D. Amazon Textract

Answer: B

Explanation

Correct Option

Amazon Fraud Detector: Amazon Fraud Detector is a fully managed service enabling customers to identify potentially fraudulent activities. For example, you can flag suspicious online payment transactions before processing payments and fulfilling orders. In another example, you can detect new account fraud. You can accurately distinguish between legitimate and high-risk account registrations, so that you can selectively introduce additional checks—such as phone or email verification.

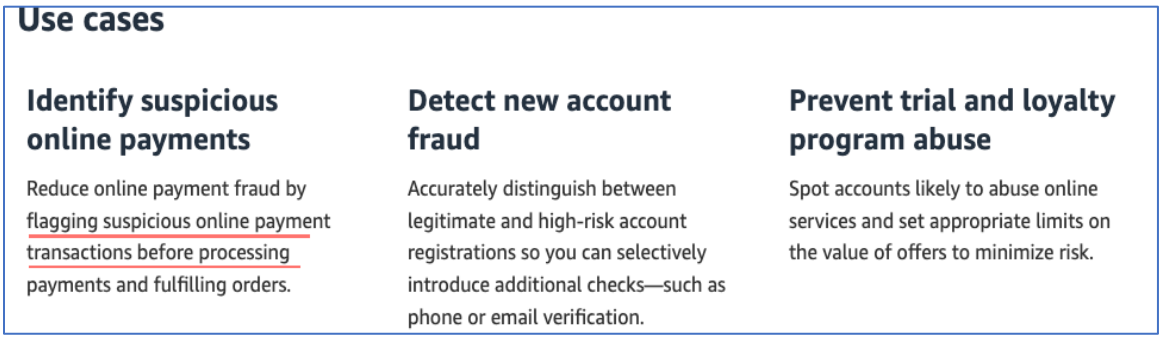

Screenshot from:
https://aws.amazon.com/fraud-detector/

Incorrect Options
Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.

Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.
Amazon Textract: Amazon Textract service enables you to add document text detection and analysis to your applications easily. Using Amazon Textract, customers can automatically extract text and data from millions of scanned documents in just hours. Amazon Textract has many use cases. For example, you can use Amazon Textract to detect typed and handwritten text in various documents. In another use case, using the Amazon Textract Document Analysis API, you can extract text, forms, and tables from structured data documents. You can process invoices and receipts with the AnalyzeExpense API in another use case. Finally, by using the AnalyzeID API, you can process ID documents such as driver's licenses and passports issued by the U.S. government.

Reference:
https://aws.amazon.com/rekognition/
https://aws.amazon.com/kendra/
https://aws.amazon.com/textract/
https://aws.amazon.com/fraud-detector/

Question 54:
You are working as a lead software engineer. You have been asked to implement features in an existing application to add product recommendations, personalized product re-ranking, and customized direct marketing. Which of the following AWS service can you use for this use case?

A. Amazon Rekognition

B. Amazon Fraud Detector
C. Amazon Personalize
D. Amazon Kendra

Answer: C

Explanation
Correct Option
Amazon Personalize: Amazon Personalize is a fully managed ML service for real-time personalized recommendations. For example, you can use this service for product recommendations, personalized product re-ranking, and customized direct marketing. Amazon Personalize provisions the infrastructure and manages the entire ML pipeline, including pre-processing, features extraction, applying the best algorithm. Additionally, it then trains, optimize, and deploy the model. You just need to call API endpoints for the deployed model. All data is encrypted, private, and secure, and is only used to create recommendations for your users.


Amazon Personalize enables developers to build applications with the same machine learning (ML) technology used by Amazon.com for real-time personalized recommendations – no ML expertise required.

Amazon Personalize makes it easy for developers to build applications capable of delivering a wide array of personalization experiences, including specific product recommendations, personalized product re-ranking, and customized direct marketing. Amazon Personalize is a fully managed machine learning service that goes beyond rigid, static rule-based recommendation systems and trains, tunes, and deploys custom ML models to deliver highly customized recommendations to customers across industries such as retail and media and entertainment.

Amazon Personalize provisions the necessary infrastructure and manages the entire ML pipeline, including processing the data, identifying features, using the best algorithms, and training, optimizing, and hosting the models. You will receive results via an Application Programming Interface (API) and only pay for what you use, with no minimum fees or upfront commitments. All data is encrypted to be private and secure, and is only used to create recommendations for your users.

Screenshot from:
https://aws.amazon.com/fraud-detector/

Incorrect Options
Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.
Amazon Fraud Detector: Amazon Fraud Detector is a fully managed service enabling customers to identify potentially fraudulent activities. For example, you can flag suspicious online payment transactions before processing payments and fulfilling orders. In another example, you can detect new account fraud. You can accurately distinguish between legitimate and high-risk account registrations so that you can selectively introduce additional checks—such as phone or email verification.
Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when

they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

Reference:

https://aws.amazon.com/rekognition/
https://aws.amazon.com/kendra/
https://aws.amazon.com/fraud-detector/

Question 55:

Which AWS service can you use to make an SSH connection to an EC2 instance without opening an inbound port?

A. AWS Systems Manager
B. AWS Systems Manager Session Manager
C. AWS CloudTrail
D. AWS Config

Answer: B

Explanation

Correct Option

AWS Systems Manager Session Manager: AWS SSM Session Manager is a fully-managed service that provides an interactive browser-based shell and CLI experience. It helps provide secure and auditable instance management without opening inbound ports, maintaining bastion hosts, and managing SSH keys. Session Manager helps to enable compliance with corporate policies that require controlled access to instances and increase security and suitability of access to the cases while providing simplicity and cross-platform instance access to end-users.

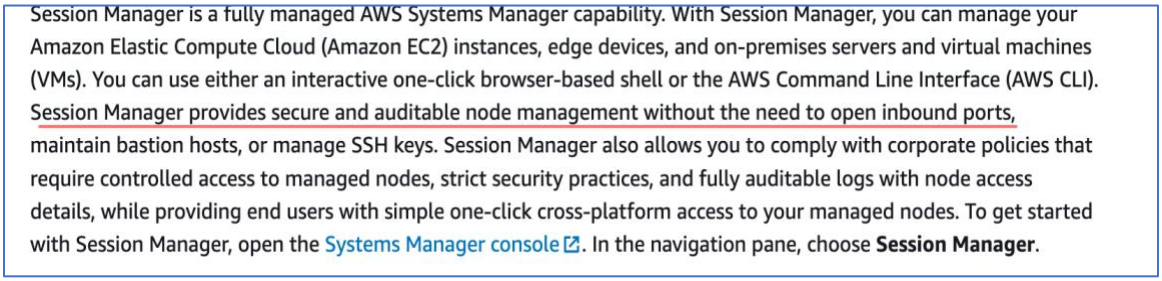
Session Manager is a fully managed AWS Systems Manager capability. With Session Manager, you can manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, edge devices, and on-premises servers and virtual machines (VMs). You can use either an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager also allows you to comply with corporate policies that require controlled access to managed nodes, strict security practices, and fully auditable logs with node access details, while providing end users with simple one-click cross-platform access to your managed nodes. To get started with Session Manager, open the Systems Manager console. In the navigation pane, choose **Session Manager**.

Screenshot from:
https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html

Incorrect Options

AWS Systems Manager: AWS Systems Manager is a set of fully managed services and capabilities that simplify the management of your Windows and Linux instances regardless of whether they are running on EC2 or on-premises. AWS Systems Manager allows performing tasks such as collecting system inventory, applying operating system patches, automation of creating Amazon Machine Images (AMIs), and configuring operating systems and applications at scale.

AWS CloudTrail: AWS CloudTrail provides auditing, security monitoring, and operational troubleshooting by tracking user activity and API usage. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure, giving you control over storage, analysis, and remediation actions. In addition, CloudTrail provides the event history of your AWS account activity, including actions taken through AWS Management.
AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations.

Reference:
https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html
https://graystum.com/aws-ssm-do-you-really-need-ssh/

Question 56:
You are working as a lead software engineer. You have been asked to redesign a search feature of an existing application to add natural language search capabilities so that employees and customers can easily find the right answers to questions when they need them instead of searching through troves of unstructured data. Which of the following AWS service can you use for this use case?

A. Amazon Polly
B. Amazon Kendra
C. Amazon Rekognition
D. Amazon Textract

Answer: B

Explanation
Correct Option
Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

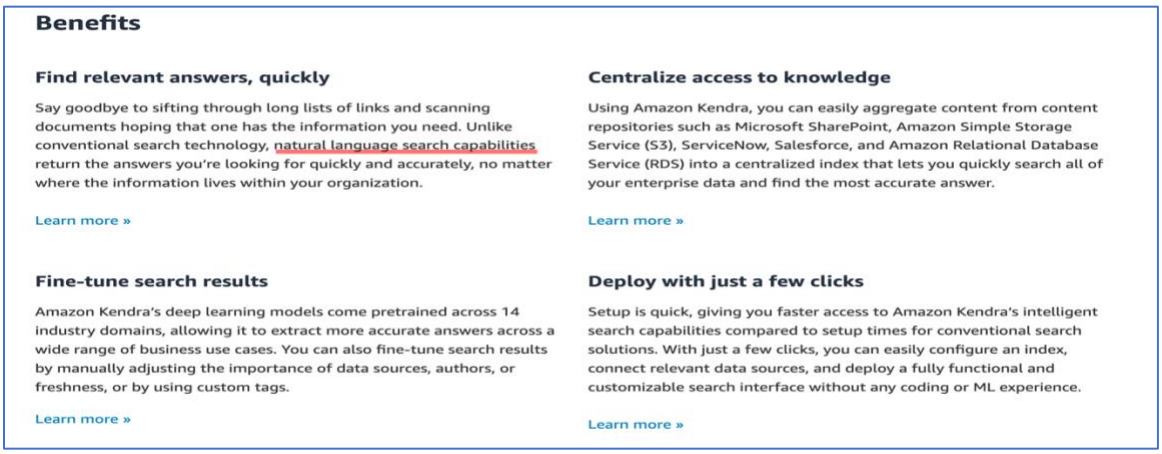

Screenshot from:
https://aws.amazon.com/kendra/

Incorrect Options

Amazon Polly: Amazon Polly is a service that turns text into lifelike speech. Amazon Polly's Text-to-Speech (TTS) uses advanced deep learning techniques to synthesize natural-sounding human speech. This natural-sounding human speech helps developers build speech-enabled products -- applications that can talk. Using machine learning, Amazon Polly offers Neural Text-to-Speech (NTTS) voices, delivering advanced improvements in speech quality. Amazon Polly Brand Voice can also create a custom NTTS voice for your organization's exclusive use.

Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.

Amazon Textract: Amazon Textract service enables you to add document text detection and analysis to your applications easily. Using Amazon Textract, customers can automatically extract text and data from millions of scanned documents in just hours. Amazon Textract has many use cases. For example, you can use Amazon Textract to detect typed and handwritten text in various documents. In another use case, using the Amazon Textract Document Analysis API, you can extract text, forms, and tables from structured data documents. You can process invoices and receipts with the AnalyzeExpense API in another use case. Finally, by using the AnalyzeID API, you can process ID documents such as driver's licenses and passports issued by the U.S. government.

Reference:

https://aws.amazon.com/rekognition/
https://aws.amazon.com/kendra/
https://aws.amazon.com/textract/
https://aws.amazon.com/polly

Question 57:
Which of the following AWS service can you use to eliminate the need to hardcode database credentials in getting a connection from the MySQL database?

A. AWS Shield
B. Amazon IAM
C. AWS Secrets Manager
D. AWS Config

Answer: C

Explanation
Correct Option
AWS Secrets Manager: AWS Secrets Manager, a secrets management service, protects your AWS account's access to applications, services, and IT resources. The service helps you easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs. This call eliminates the need to hardcode sensitive information in plain text. The AWS Secrets Manager also offers secret rotation with built-in integration for Amazon RDS, Redshift, and Amazon DocumentDB.

Incorrect Options
AWS Shield: AWS Shield prevents DDoS (Distributed Denial of Service) attacks on AWS resources for both global and regional resources.
Amazon IAM: Amazon IAM service allows you to securely control access to AWS services and resources. You can create and manage AWS users and groups using IAM. You can set up permissions to allow and deny their access to AWS resources.
AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations.

All these services cannot be used for the use case mentioned in the question.

Reference:
https://aws.amazon.com/secrets-manager/
https://aws.amazon.com/shield
https://aws.amazon.com/config/
https://aws.amazon.com/iam

Question 58:
Which of the following AWS service/feature can you use to scan your AWS infrastructure, compare it with AWS best practices, and provides recommended action?

A. AWS Trusted Advisor
B. AWS Systems Manager
C. AWS Shield Advanced
D. AWS Config

Answer: A

Explanation

Correct Option

AWS Trusted Advisor: AWS Trusted Advisor analyzes and evaluates your AWS account using checks and provides best practice recommendations. These checks help identify ways to optimize your AWS infrastructure, improve security and performance, reduce costs, and monitor service quotas. AWS Trusted Advisor provides best practice recommendations in five categories: Cost Optimization, Performance, Security, Fault Tolerance, and Service Limits. Whether establishing new workflows, developing applications, or as part of ongoing improvement, take advantage of the recommendations provided by Trusted Advisor regularly to help keep your solutions provisioned optimally.

While Trusted advisor checks are based on the support plan the customer has, both Basic and Developer support plans have access to the seven core Trusted Advisor checks. Unlike documentation-based guidance, such as AWS Well-Architected Tool, this tool provides recommendations against AWS Well-Architected Framework best practices and can track against your current AWS architecture.

How Trusted Advisor Works:

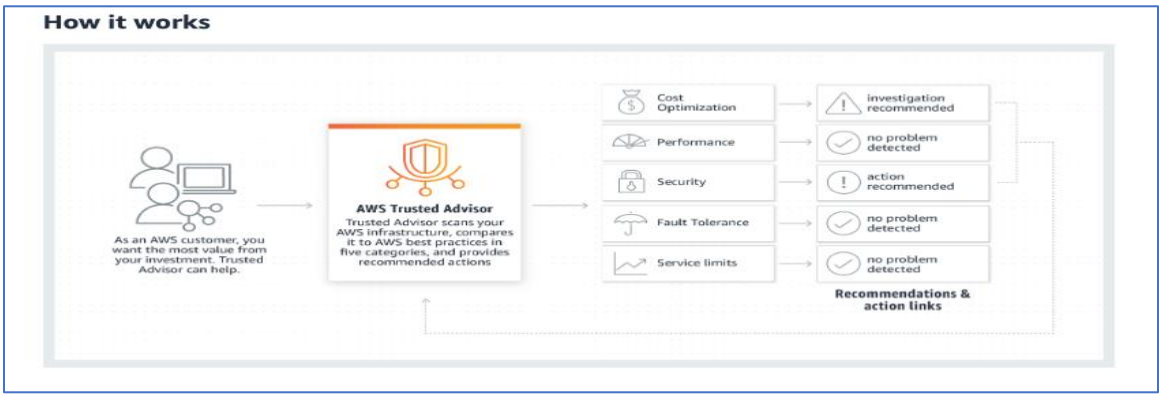

Screenshot from:
https://aws.amazon.com/premiumsupport/technology/trusted-advisor/

AWS Trusted Advisor Recommendations:

Screenshot from:
https://aws.amazon.com/premiumsupport/technology/trusted-advisor/

Incorrect Options

AWS Systems Manager: AWS Systems Manager is a set of fully managed services and capabilities that simplify the management of your Windows and Linux instances regardless of whether they are running on EC2 or on-premises. AWS Systems Manager allows performing tasks such as collecting system inventory, applying operating system patches, automation of creating Amazon Machine Images (AMIs), and configuring operating systems and applications at scale.

AWS Shield Advanced: AWS Shield prevents DDoS (Distributed Denial of Service) attacks on AWS resources for both global and regional resources. AWS Shield is available in two levels: AWS Shield Standard and AWS Shield Advanced. AWS Shield Standard is free and open to all customers. AWS Shield Standard protects all AWS customers against common and most frequently occurring infrastructure (layer 3 and 4) attacks like SYN/UDP floods, reflection attacks, and others to support the high availability of your applications on AWS. AWS Shield Advanced is a paid service. Being a paid service, it provides every feature that AWS Shield Standard. Additionally, it offers additional protection such as visibility and monitoring of DDoS attacks, 24x7 support from the AWS DDoS support team, and AWS WAF subscription at no cost.

AWS Config: AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records changes to your AWS resource configurations. It also allows you to automate the evaluation of recorded configurations against desired configurations.

Reference:

https://aws.amazon.com/premiumsupport/technology/trusted-advisor/
https://aws.amazon.com/config/
https://aws.amazon.com/shield/
https://aws.amazon.com/systems-manager/

Question 59:

Which AWS service can quickly process ID documents such as driver's licenses and passports issued by the U.S. government?

A. Amazon Polly
B. Amazon Kendra
C. Amazon Rekognition
D. Amazon Textract

Answer: D

Explanation

Correct Option

Amazon Textract: Amazon Textract service enables you to add document text detection and analysis to your applications easily. Using Amazon Textract, customers can automatically extract text and data from millions of scanned documents in just hours. Amazon Textract has many use cases. For example, you can use Amazon Textract to detect typed and handwritten text in various documents. In another use case, using the Amazon Textract Document Analysis API, you can extract text, forms, and tables from structured data documents. You can process invoices and receipts with the AnalyzeExpense API in another use case. Finally, by using the AnalyzeID API, you can process ID documents such as driver's licenses and passports issued by the U.S. government.

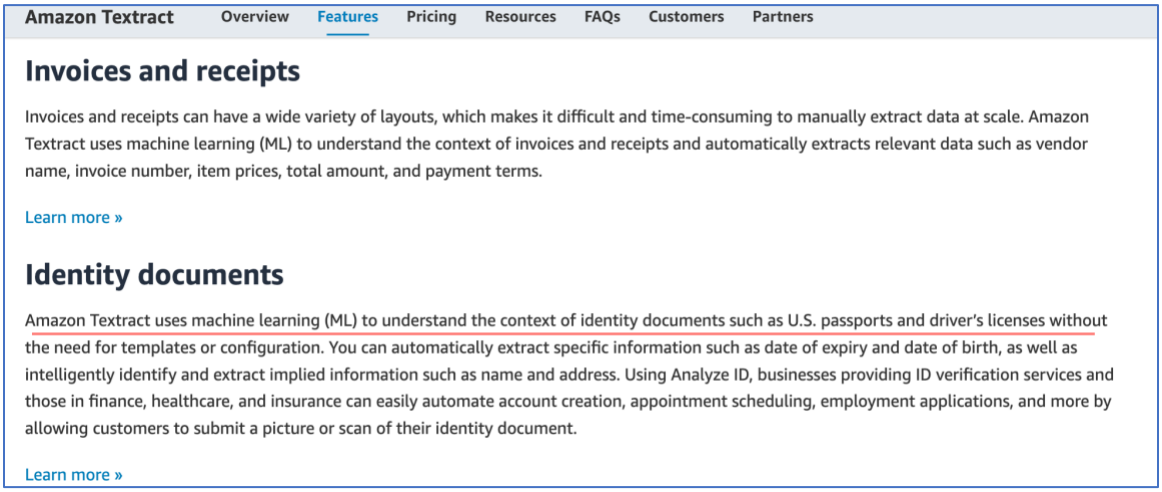

Incorrect Options

Amazon Polly: Amazon Polly is a service that turns text into lifelike speech. Amazon Polly's Text-to-Speech (TTS) uses advanced deep learning techniques to synthesize natural-sounding human speech. This natural-sounding human speech helps developers build speech-enabled products -- applications that can talk. Using machine learning, Amazon Polly offers Neural Text-to-Speech (NTTS) voices, delivering advanced improvements in speech quality. Amazon Polly Brand Voice can also create a custom NTTS voice for your organization's exclusive use.

Amazon Kendra: Amazon Kendra is a fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.

Reference:

https://aws.amazon.com/rekognition/
https://aws.amazon.com/kendra/
https://aws.amazon.com/textract/
https://aws.amazon.com/polly

Question 60:

You have been designing a real-time analytic application, in which if a user submits an order, the order information is sent to the DynamoDB database. Which of the following AWS services can you use to implement this highly available application cost-effectively? (Select Two)

A. Amazon Kinesis

B. Amazon SageMaker
C. AWS Elastic Beanstalk
D. AWS Lambda
E. Amazon EFS

Answer: A, D

Explanation
Correct Options
Amazon Kinesis: You can Amazon Kinesis to ingest stream (order). Amazon Kinesis makes it easy to ingest, process, and analyze real-time, streaming data. These ingest, processes, and analysis of real-time streaming data help you get timely insights so that you can make better management decisions or react quickly. Amazon Kinesis enables you to process and analyze data as it arrives and responds instantly instead of the old classic style of data collection, where you have to wait until all your data is collected before the processing can begin.
AWS Lambda: AWS Lambda can be used to process stream (order) and store to DynamoDB. AWS Lambda enables you to execute code without provisioning or managing servers. Instead, you are charged based on the number of requests for your functions and the duration it takes for your code to execute. AWS Lambda executes code in response to events such as object uploads to S3, updates to DynamoDB tables, or other events such as website clicks.

Incorrect Options

Amazon SageMaker: Amazon SageMaker enables business analysts, data scientists, and ML engineers to build, train, and deploy machine learning (ML) models for different use cases with fully managed infrastructure, tools, and workflows. Business analysts can use SageMaker to make ML predictions using a visual interface with SageMaker Canvas. Data scientists can use SageMaker to prepare data and build, train, and deploy models with SageMaker Studio. ML engineers can use SageMaker to deploy and manage models at scale with SageMaker with MLOps.

AWS Elastic Beanstalk: AWS Elastic Beanstalk is a PaaS (Platform-as-a-Service) type of AWS service that is used for deploying and scaling web applications. This AWS service takes your application code and deploys it while provisioning the compute resources required for the application to run.

Amazon EFS: Amazon EFS is a cloud-native, serverless, and fully managed file system that is accessible from Linux instances via the NFS protocol. It is built to scale on-demand to petabytes without disrupting applications. The EFS scales out and scales in automatically as you add and remove files. As a result, you don't need to provision and manage capacity to accommodate growth. In addition, Amazon EFS is designed to provide massively parallel shared access to thousands of EC2 instances. Thus, enabling your applications to achieve high aggregate throughput and IOPS with consistent low latency.

Reference:
https://aws.amazon.com/kinesis/
https://aws.amazon.com/elasticbeanstalk/

https://aws.amazon.com/efs/
https://aws.amazon.com/sagemaker/

Question 61:
You are working as an AWS consultant for a client. The client would like to do operational planning to predict levels of web traffic, AWS usage, and IoT sensor usage. Which of the following AWS services is the best fit for this use case if you need to implement this use quickly?

A. Amazon Lex
B. Amazon Kendra
C. Amazon Rekognition
D. Amazon Forecast

Answer: D

Explanation
Correct Option
Amazon Forecast: Amazon Forecast uses statistical and machine learning algorithms to deliver highly accurate time-series forecasts – without any machine learning experience. It is a fully managed service. Amazon Forecast provides automation by finding the optimal combination of machine learning algorithms for your datasets. In addition, it offers several filling methods to automatically handle missing values in your datasets.

You can use this service for use cases such as retail demand planning to predict product demand, allowing you to vary inventory and pricing more accurately at different store locations. It can also be used in supply chain planning to forecast the quantity of raw goods, services, or other inputs required by manufacturing. Another use case is a resource planning to predict staffing, advertising, energy consumption, and server capacity requirements. And finally, Amazon Forecast can be used in operational planning to predict levels of web traffic, AWS usage, and IoT sensor usage.

You can use the APIs, AWS Command Line Interface (AWS CLI), Python Software Development Kit (SDK), and Amazon Forecast Console to import time-series datasets, train predictors, and generate forecasts.

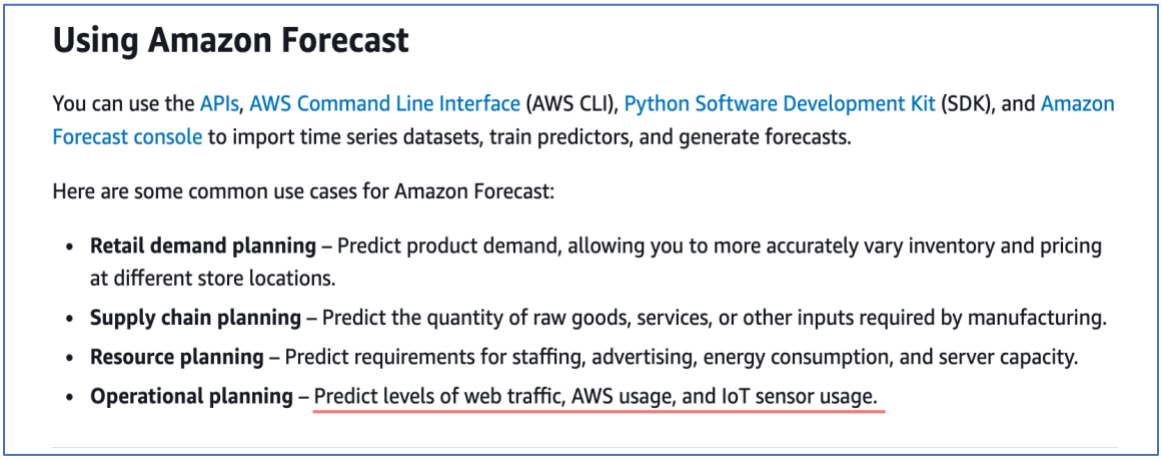

Using Amazon Forecast

You can use the APIs, AWS Command Line Interface (AWS CLI), Python Software Development Kit (SDK), and Amazon Forecast console to import time series datasets, train predictors, and generate forecasts.

Here are some common use cases for Amazon Forecast:

- **Retail demand planning** – Predict product demand, allowing you to more accurately vary inventory and pricing at different store locations.
- **Supply chain planning** – Predict the quantity of raw goods, services, or other inputs required by manufacturing.
- **Resource planning** – Predict requirements for staffing, advertising, energy consumption, and server capacity.
- **Operational planning** – Predict levels of web traffic, AWS usage, and IoT sensor usage.

Screenshot from:
https://docs.aws.amazon.com/forecast/latest/dg/what-is-forecast.html

Incorrect Options

Amazon Lex: Amazon Lex, which uses advanced natural language models, is a fully managed service to build conversational applications, such as conversational chatbots. The conversational chatbot types of applications help improve customer service and increase conversion efficiency. Amazon Lex helps developers to build conversational chatbots quickly -- no deep learning expertise is necessary. To create a bot, you specify the basic conversation flow in the Amazon Lex console, and the service manages the dialogue and dynamically adjusts the responses in the conversation. In addition, it provides pre-built integration with AWS Lambda. For example, you can easily integrate Amazon Lex with Amazon Cognito, AWS Mobile Hub, Amazon CloudWatch, and Amazon DynamoDB. Integration of Amazon Lex with Lambda provides bots access to pre-built serverless enterprise connectors to link to data in SaaS applications, such as Salesforce.

Amazon Kendra: Amazon Kendra is fully managed intelligent search service that adds natural language search capabilities. Amazon Kendra reimagines enterprise search for websites and applications so that employees and customers can easily find the right answers to questions when that they need them. How Kendra does it -- Kendra does it by searching through troves of unstructured data to provide the right answer.

Amazon Rekognition: Amazon Rekognition is a simple and easy service to quickly analyze pictures and videos stored on S3. It is a fully managed computer vision service that helps automate your image and video analysis, thus avoiding manual inspection. In addition, the service employs proven and highly scalable deep learning technology. Using Amazon Rekognition, you can easily add image and video analysis capability to your applications through simple API endpoints.

Reference:

https://docs.aws.amazon.com/forecast/latest/dg/what-is-forecast.html
https://aws.amazon.com/rekognition/
https://aws.amazon.com/kendra/
https://aws.amazon.com/lex/

Question 62:

Which of the following can you use to connect Window File Server from an EC2 Linux instance?

A. AWS Systems Manager
B. Amazon FSx for Windows
C. Amazon API Gateway
D. AWS Direct Connect

Answer: B

Explanation

Correct Option

Amazon FSx for Windows: Amazon FSx for Windows File Server is fully managed, highly reliable, and scalable file storage accessible over the industry-standard Service Message Block (SMB) protocol. It is built on Windows Server, delivering many administrative features. For

example, it includes user quotas and Microsoft Active Directory (AD) integration. In addition, Amazon FSx provides high throughput levels and sub-millisecond latencies. Amazon FSx is accessible from Windows, Linux, and macOS compute instances and devices.

Incorrect Options

AWS Systems Manager: AWS Systems Manager is a set of fully managed services and capabilities that simplify the management of your Windows and Linux instances regardless of whether they are running on EC2 or on-premises. AWS Systems Manager allows performing tasks such as collecting system inventory, applying operating system patches, automation of creating Amazon Machine Images (AMIs), and configuring operating systems and applications at scale.

AWS API Gateway: Amazon API Gateway, a fully managed service, makes it easy for developers to develop, publish, maintain, monitor, and secure APIs at any scale. APIs are considered to be the "front door" for applications to access data, business logic, or functionality from backend services. You can create RESTful APIs and WebSocket APIs that enable real-time communication between applications using API Gateway. API Gateway supports containerized and serverless workloads and web applications as well.

AWS Direct Connect: AWS Direct Connect helps make it easy to set up a dedicated network connection from your on-premises data center to the AWS cloud, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. The AWS Direct Connect is a physical connection. However, AWS Direct Connect does not encrypt your traffic in transit.

Reference:

https://aws.amazon.com/fsx/windows/
https://aws.amazon.com/systems-manager/
https://aws.amazon.com/directconnect/

Question 63:

Which of the following storage service is transient?

A. Amazon EFS
B. Amazon EBS
C. Amazon S3
D. Amazon EC2 Instance Store

Answer: D

Explanation

Correct Option

Amazon EC2 Instance Store: An instance store is temporary block-level storage for an EC2 instance. An instance store provides temporary block-level storage for your EC2 instances. The instance store storage is located on disks that are physically attached to the host computer. An instance store is ideal for the temporary storage of information that frequently changes, such as buffers, caches, and other temporary content. The key point to note about Instance Store is that it is temporary storage. What it means is that the data is lost if the instance experiences failure or when the instance is terminated.

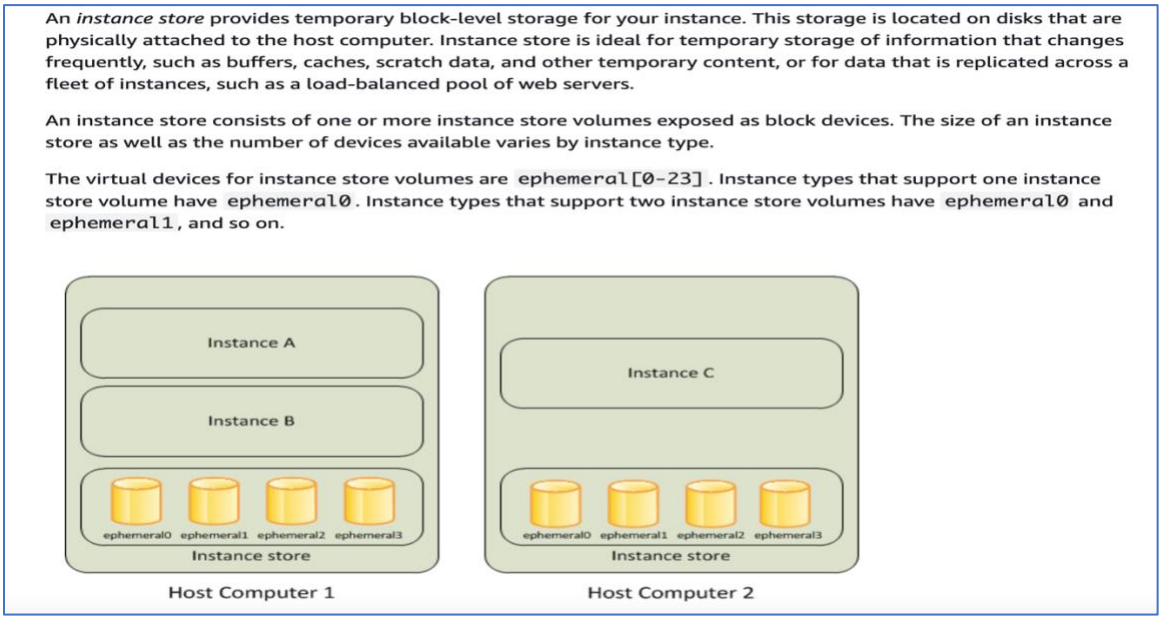

Screenshot from:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html

Incorrect Options

Amazon EFS: Amazon EFS is persistent storage. Amazon EFS is a cloud-native, serverless, and fully managed file system that is accessible from Linux instances via the NFS protocol. It is built to scale on-demand to petabytes without disrupting applications. The EFS scales out and scales in automatically as you add and remove files. As a result, you don't need to provision and manages capacity to accommodate growth. In addition, Amazon EFS is designed to provide massively parallel shared access to thousands of EC2 instances. Thus, enabling your applications to achieve high aggregate throughput and IOPS with consistent low latency.

Amazon EBS: Amazon EBS is an easy-to-use, high-performance, block-storage service designed to store persistent data with Amazon EC2 instances. It is a block-storage service and not a file storage service. It is designed to work with EC2 for both throughput and transaction-intensive workloads at any scale. Many workloads, such as enterprise applications, containerized applications, and many other types of applications, are widely deployed on Amazon EBS. An EBS can only be mounted to one EC2 instance at a time.

Amazon S3: Amazon S3 is an object storage service offering scalability, data availability, security, and performance. Customers having various use cases can use S3 to store, protect, and retrieve any amount of data for different use cases at any time, from anywhere for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.

These are all persistence storage options.

Reference:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html

https://aws.amazon.com/efs/
https://aws.amazon.com/ebs
https://aws.amazon.com/s3

Question 64:
You need to access a file from two EC2 instances running in two separate AZs. Which of the following storage service can you use for this use case?

A. Amazon EBS
B. Amazon S3
C. Amazon EFS
D. Amazon EC2 Instance Store

Answer: C

Explanation
Correct Option
Amazon EFS: Amazon EFS is persistent storage. Amazon EFS is a cloud-native, serverless, and fully managed file system that is accessible from Linux instances via the NFS protocol. It is built to scale on-demand to petabytes without disrupting applications. The EFS scales out and scales in automatically as you add and remove files. As a result, you don't need to provision and manages capacity to accommodate growth. In addition, Amazon EFS is designed to provide massively parallel shared access to thousands of EC2 instances. Thus, enabling your applications to achieve high aggregate throughput and IOPS with consistent low latency.

You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 and other AWS compute instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.

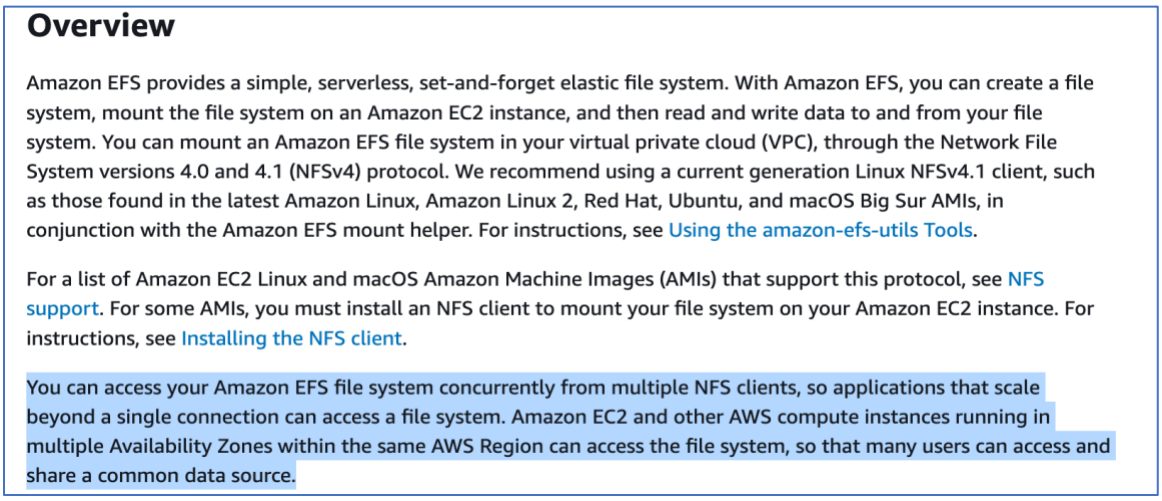

Screenshot from:
https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html

There is no minimum fee or setup charge. You only pay for what you use for the storage, for reading and writing access to data stored in Infrequent Access storage classes, and for any provisioned throughput.

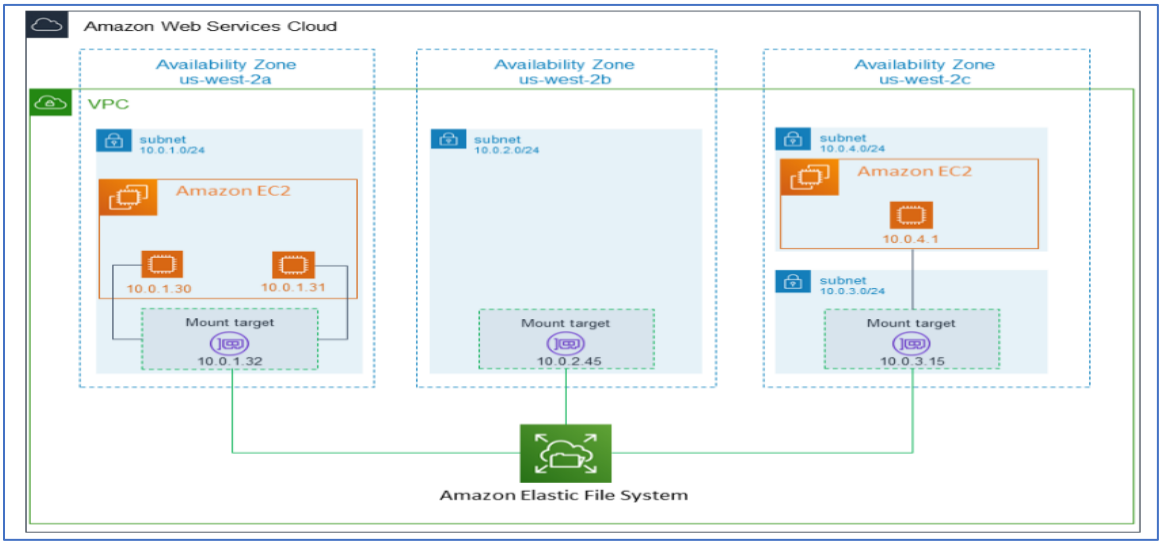

EFS Features:

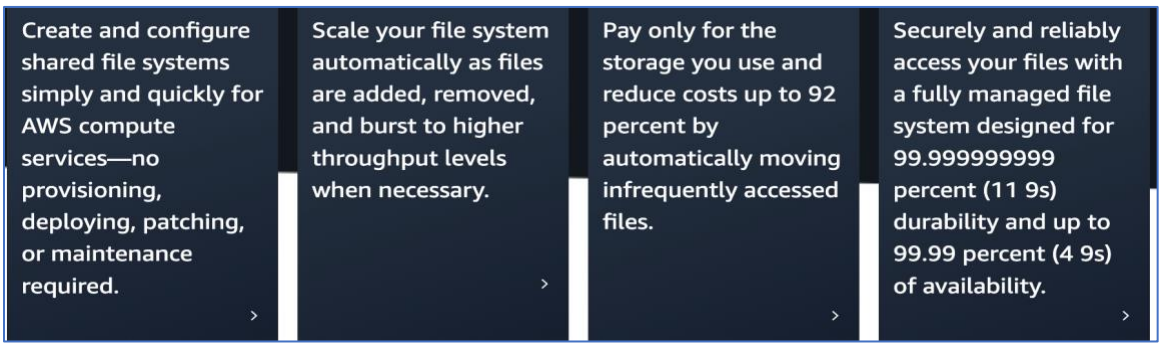

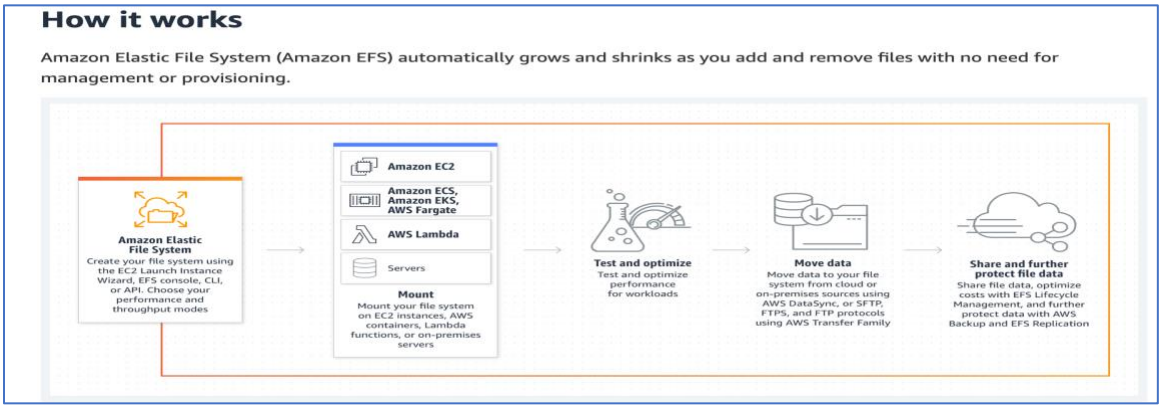

Screenshot from:

https://aws.amazon.com/efs/

Incorrect Options

Amazon EBS: Amazon EBS is an easy-to-use, high-performance, block-storage service designed to store persistent data with Amazon EC2 instances. It is a block-storage service and not a file storage service. It is designed to work with EC2 for both throughput and transaction-intensive workloads at any scale. Many workloads, such as enterprise applications, containerized applications, and many other types of applications, are widely deployed on Amazon EBS. An EBS can only be mounted to one EC2 instance at a time.

Amazon S3: Amazon S3 is an object storage service offering scalability, data availability, security, and performance. Customers having various use cases can use S3 to store, protect, and retrieve any amount of data for different use cases at any time, from anywhere for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.

Amazon EC2 Instance Store: An instance store is temporary block-level storage for an EC2 instance. An instance store provides temporary block-level storage for your EC2 instances. The instance store storage is located on disks that are physically attached to the host computer. An instance store is ideal for the temporary storage of information that frequently changes, such as buffers, caches, and other temporary content. The key point to note about Instance Store is that it is temporary storage. What it means is that the data is lost if the instance experiences failure or when the instance is terminated.

Reference:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html
https://aws.amazon.com/efs/
https://aws.amazon.com/ebs
https://aws.amazon.com/s3
https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html

Question 65:

Which of the following AWS services can you use to enrich events from SaaS applications using AWS AI/ML services to gain valuable insights?

A. Amazon SageMaker
B. Amazon EventBridge
C. AWS Glue
D. Amazon SNS

Answer: B

Explanation

Correct Option

Amazon EventBridge: You can enrich your events from SaaS applications using AWS AI/ML services and gain valuable insights. For example, you can load your data from Shopify to EventBridge to trigger a workflow and use AI services such as Amazon Comprehend to images of new retail products.

Amazon EventBridge, a serverless event bus, makes it easier to build event-driven applications. It helps scale the event-driven applications using events generated from your applications, SaaS applications, and AWS services. Amazon EventBridge delivers real-time data streams from event sources such as AWS services, SaaS apps, microservices, or custom apps to targets like AWS Lambda and other SaaS applications. When building application architectures for event-driven applications, you can set up routing rules to determine where to send your data to react in real-time to the event sources with event publishers and consumers completely decoupled.

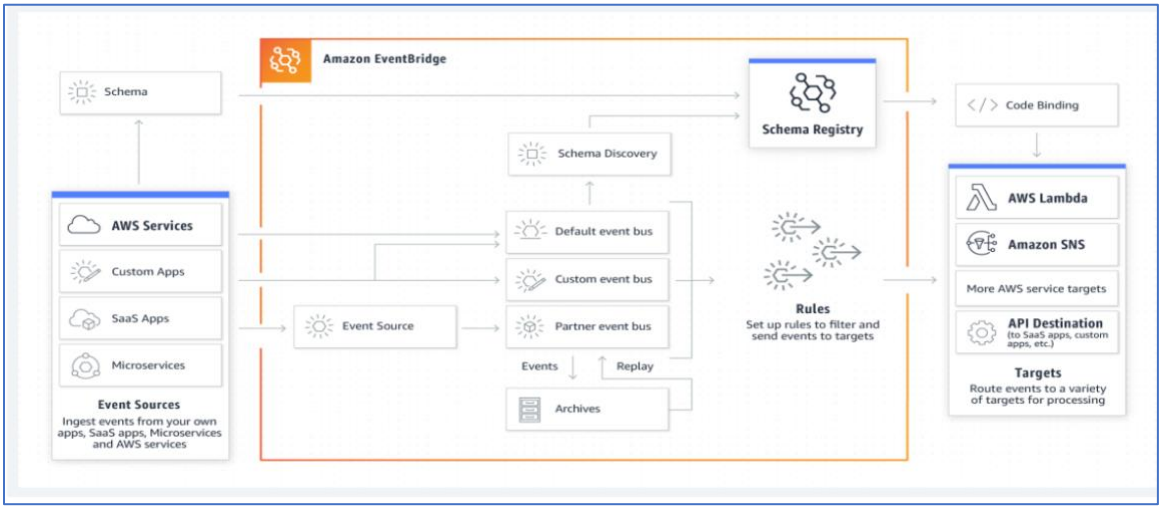

Screenshot from:
https://aws.amazon.com/eventbridge/

Use cases of Amazon EventBridge are re-architecting for speed, monitoring, and auditing, extending functionality via SaaS integrations, and customizing SaaS with AI/ML.

Incorrect Options

Amazon SageMaker: Amazon SageMaker enables business analysts, data scientists, and ML engineers to build, train, and deploy machine learning (ML) models for different use cases with fully managed infrastructure, tools, and workflows. Business analysts can use SageMaker to make ML predictions using a visual interface with SageMaker Canvas. Data scientists can use SageMaker to prepare data and build, train, and deploy models with SageMaker Studio. ML engineers can use SageMaker to deploy and manage models at scale with SageMaker with MLOps.

AWS Glue: AWS Glue is a serverless service for analytics, machine learning, and data engineering. It makes it easy to discover, prepare, and combine data for data engineering or related tasks. AWS Glue can crawl your data sources to identify data formats and suggests schemas to store data. The automatic generation of schema structure from the provided data sources saves time which helps in the faster development of ETL jobs such as data extraction, data pre-processing, data transformations, and data loading on the AWS Glue platform. Using AWS Glue, you can write Spark ETL jobs with data stored in S3, RDS, or on file using Python or Scala and set up the data pipeline in Glue as well to schedule ETL jobs.

Amazon SNS: Amazon SNS is an essentially fully managed pub/sub messaging service of AWS. It is highly available, durable, and secure. Amazon SNS can help you decouple microservices, distributed systems, and serverless applications. It is for both application-to-application (A2A) and application-to-person (A2P) communication. The A2A pub/sub functionality provides topics for high-throughput, push-based, many-to-many messaging between distributed systems, microservices, and event-driven serverless applications. Amazon SNS allows applications to send time-critical messages to multiple subscribers through a "push" mechanism, implying that the receiving applications have to be present and run to receive the notifications.

Reference:

https://aws.amazon.com/eventbridge/
https://aws.amazon.com/sns/
https://aws.amazon.com/glue/
https://aws.amazon.com/sagemaker/