

Amazon.Premium.AWS Certified Solutions Architect - Professional.by.VCEplus.105q

Number: AWS VCEplus
Passing Score: 800
Time Limit: 120 min
File Version: 1.2



Certification: AWS Certified Solutions Architect - Professional

Certification Full Name: AWS Certified Solutions Architect - Professional

Certification Provider: Amazon

Exam Code: AWS Certified Solutions Architect - Professional

Exam Name: AWS Certified Solutions Architect - Professional

Website: www.vceplus.com

Free Exam: <https://vceplus.com/exam-AWS-Certified-Solutions-Architect-Professional/>

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in AWS Certified Solutions Architect - Professional exam products and you get latest questions. We strive to deliver the best AWS Certified Solutions Architect - Professional exam product for top grades in your first attempt.

VCE to PDF Converter : <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Google+ : <https://plus.google.com/+Vcepluscom>

LinkedIn : <https://www.linkedin.com/company/vceplus>

QUESTION 1

Your company policies require encryption of sensitive data at rest. You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance.

Which of these options would allow you to encrypt your data at rest? (Choose 3)

- A. Implement third party volume encryption tools
- B. Implement SSL/TLS for all services running on the server
- C. Encrypt data inside your applications before storing it on EBS
- D. Encrypt data using native data encryption drivers at the file system level
- E. Do nothing as EBS volumes are encrypted by default

Correct Answer: ACD

Section: (none)

Explanation

Explanation/Reference:

**QUESTION 2**

A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

- A. Upload the certificate on an S3 bucket owned by the security officers and accessible only by EC2 Role of the web servers.
- B. Configure the web servers to retrieve the certificate upon boot from an CloudHSM is managed by the security officers.
- C. Configure system permissions on the web servers to restrict access to the certificate only to the authority security officers
- D. Configure IAM policies authorizing access to the certificate store only to the security officers and terminate SSL on an ELB.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You'll terminate the SSL at ELB. and the web request will get unencrypted to the EC2 instance, even if the certs are stored in S3, it has to be configured on the web servers or load balancers somehow, which becomes difficult if the keys are stored in S3. However, keeping the keys in the cert store and using IAM to restrict access gives a clear separation of concern between security officers and developers. Developer's personnel can still configure SSL on ELB without actually

handling the keys.

QUESTION 3

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1 KB of sensor data every minute to a backend hosted on AWS.

During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling.

Which setup will meet the requirements?

- A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
- B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
- C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
- D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

QUESTION 4

A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC.

How should they architect their solution to achieve these goals?

- A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see all traffic across the VPC.
- B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
- C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.
- D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:**QUESTION 5**

A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest Which of the following methods can achieve this? (Choose 3)

- A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
- B. Use Amazon S3 server-side encryption with customer-provided keys.
- C. Use Amazon S3 server-side encryption with EC2 key pair.
- D. Use Amazon S3 bucket policies to restrict access to the data at rest.
- E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
- F. Use SSL to encrypt the data while in transit to Amazon S3.

Correct Answer: ABE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/UsinaKMSEncryption.html>

**QUESTION 6**

Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

- A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idle. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
- B. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
- C. Change the storage class of the S3 objects to Reduced Redundancy Storage. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS Once data is processed, change the storage class of the S3 objects to Glacier.
- D. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 7

You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoDB for their dynamic data and then archiving nightly into S3 for further processing with EMR. They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access.

Which approach provides a cost effective scalable mitigation to this kind of attack?

- A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC they would then establish Internet connectivity into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
- B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet.
- C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running a host-based WAF. They would redirect Route 53 to resolve to the new WAF tier ELB. The WAF tier would then pass the traffic to the current web tier. The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group.
- D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering. This will enable the ELB itself to perform WAF functionality.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 8

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring the following requirements are met.

Provide the ability for real-time analytics of the inbound biometric data.

Ensure processing of the biometric data is highly durable, elastic and parallel.

The results of the analytic processing should be persisted for data mining.

Which architecture outlined below will meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data, analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.

- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EM R to collect the inbound sensor data, analyze the data from EUR with Amazon Kinesis and save me results to DynamoDB.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 9

You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture, Which alternatives should you consider? (Choose 2)

- A. Configure a NAT instance in your VPC Create a default route via the NAT instance and associate it with all subnets Configure a DNS A record that points to the NAT instance public IP address.
- B. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers Configure a Route53 CNAME record to your CloudFront distribution.
- C. Place all your web servers behind ELB Configure a Route53 CNMIE to point to the ELB DNS name.
- D. Assign EIPs to all web servers. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
- E. Configure ELB with an EIP Place all your Web servers behind ELB Configure a Route53 A record that points to the EIP.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 10

Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC.

The optimal setup for persistence and security that meets the above requirements would be the following.

- A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
- B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code Alter its security group to allow access to it

from hosts within your VPC's IP address block.

- C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
- D. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable Alter its security group to allow access to It from hosts in your application subnets.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 11

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructure's ability to handle unexpected increases in traffic.

The application currently consists of 2 tiers: a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database.

Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

- A. Failover environment: Create an S3 bucket and configure it for website hosting. Migrate your DNS to Route53 using zone file import, and leverage Route53 DNS failover to failover to the S3 hosted website.
- B. Hybrid environment: Create an AMI, which can be used to launch web servers in EC2. Create an Auto Scaling group, which uses the AMI to scale the web tier based on incoming traffic. Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted in AWS.
- C. Offload traffic from on-premises environment: Setup a CloudFront distribution, and configure CloudFront to cache objects from a custom origin. Choose to customize your object cache behavior, and select a TTL that objects should exist in cache.
- D. Migrate to AWS: Use VM Import/Export to quickly convert an on-premises web server to an AMI. Create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffic. Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to migrate the database.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 12

You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want

other Internet traffic to use your existing link to an Internet Service Provider.

What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

- A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.
- B. Create a private interface on your AWS Direct Connect link. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC.
- C. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure; advertise specific routes for your network to AWS.
- D. Create a private interface on your AWS Direct connect link. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 13

Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence,

The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.

Do you need to change anything in the architecture to maintain the high availability of the application with the anticipated additional load? Why?

- A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
- B. No, if the cache node fails you can always get the same data from the DB without having any availability impact.
- C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
- D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

ElastiCache for Memcached

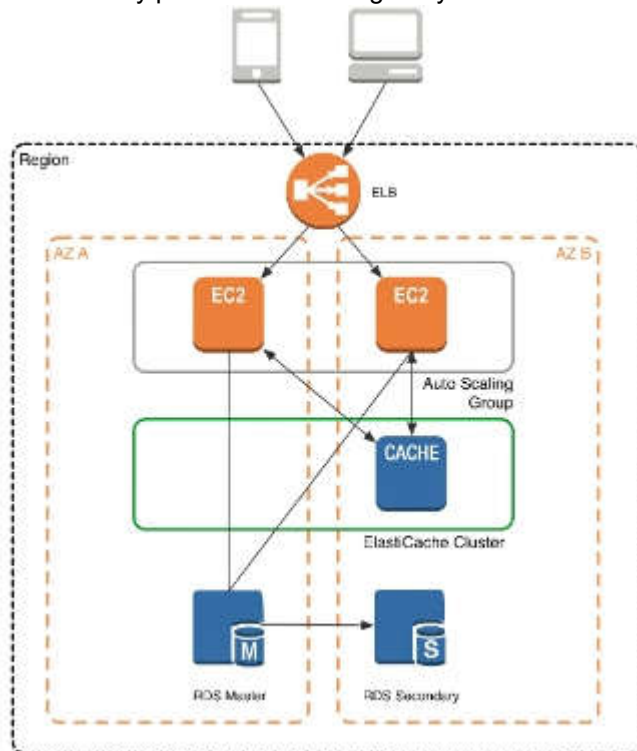
The primary goal of caching is typically to offload reads from your database or other primary data source. In most apps, you have hot spots of data that are regularly queried, but only updated periodically. Think of the front page of a blog or news site, or the top 100 leaderboard in an online game. In this type of case, your app can receive dozens, hundreds, or even thousands of requests for the same data before it's updated again. Having your caching layer handle these queries has several advantages. First, it's considerably cheaper to add an in-memory cache than to scale up to a larger database cluster. Second, an in-memory cache is also easier to scale out, because it's easier to distribute an in-memory cache horizontally than a relational database.

Last, a caching layer provides a request buffer in the event of a sudden spike in usage. If your app or game ends up on the front page of Reddit or the App Store, it's not unheard of to see a spike that is 10 to 100 times your normal application load. Even if you autoscale your application instances, a 10x request spike will likely make your database very unhappy.

Let's focus on ElastiCache for Memcached first, because it is the best fit for a caching-focused solution. We'll revisit Redis later in the paper, and weigh its advantages and disadvantages.

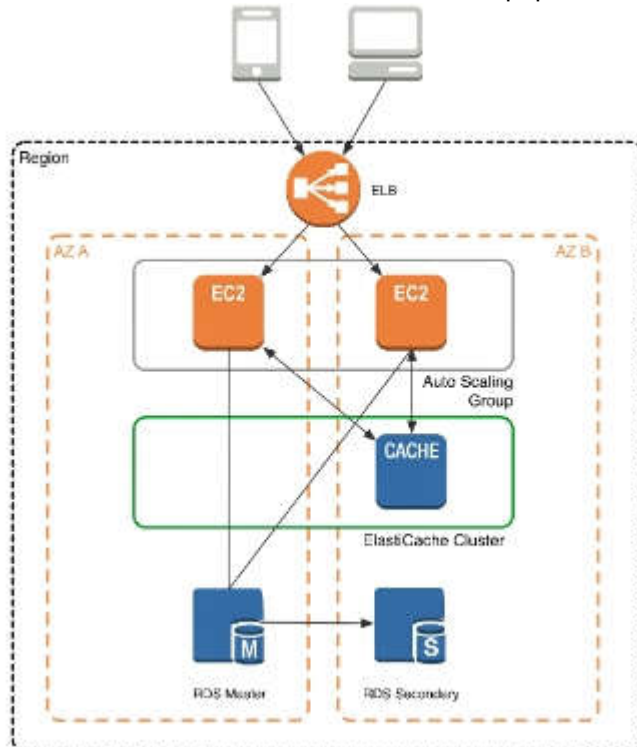
Architecture with ElastiCache for Memcached

When you deploy an ElastiCache Memcached cluster, it sits in your application as a separate tier alongside your database. As mentioned previously, Amazon ElastiCache does not directly communicate with your database tier, or indeed have any particular knowledge of your database. A simplified deployment for a web application looks something like this:

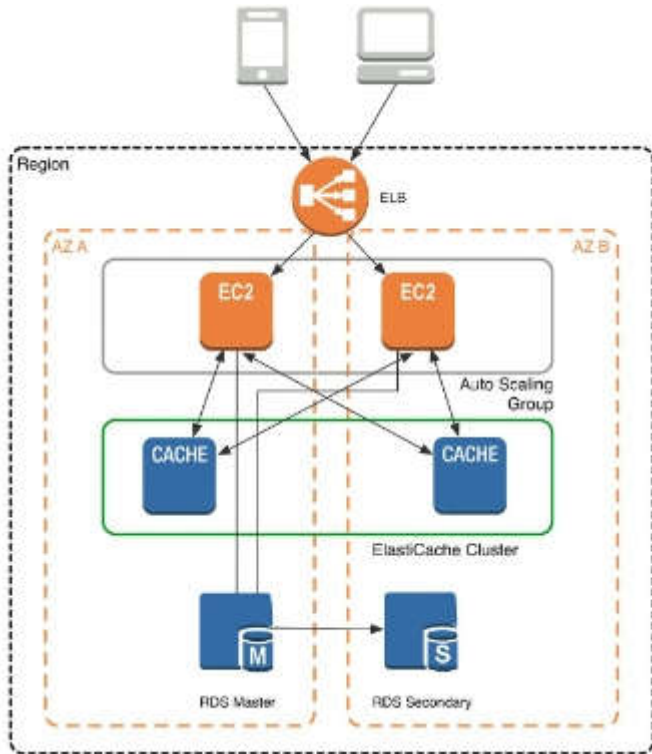


In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given

EC2 instance, that EC2 instance is responsible for communicating with ElastiCache and the database tier. For development purposes, you can begin with a single ElastiCache node to test your application, and then scale to additional cluster nodes by modifying the ElastiCache cluster. As you add additional cache nodes, the EC2 application instances are able to distribute cache keys across multiple ElastiCache nodes. The most common practice is to use client-side sharding to distribute keys across cache nodes, which we will discuss later in this paper.



When you launch an ElastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an ElastiCache cluster in a specific Availability Zone, make sure to specify the Preferred Zone(s) option during cache cluster creation. The Availability Zones that you specify will be where ElastiCache will launch your cache nodes. We recommend that you select Spread Nodes Across Zones, which tells ElastiCache to distribute cache nodes across these zones as evenly as possible. This distribution will mitigate the impact of an Availability Zone disruption on your ElastiCache nodes. The trade-off is that some of the requests from your application to ElastiCache will go to a node in a different Availability Zone, meaning latency will be slightly higher. For more details, refer to Creating a Cache Cluster in the Amazon ElastiCache User Guide. As mentioned at the outset, ElastiCache can be coupled with a wide variety of databases. Here is an example architecture that uses Amazon DynamoDB instead of Amazon RDS and MySQL:



This combination of DynamoDB and ElastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases. In addition, DynamoDB uses a key-value access pattern similar to ElastiCache, which also simplifies the programming model. Instead of using relational SQL for the primary database but then key-value patterns for the cache, both the primary database and cache can be programmed similarly. In this architecture pattern, DynamoDB remains the source of truth for data, but application reads are offloaded to ElastiCache for a speed boost.

QUESTION 14

An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes the customer realizes that data corruption occurred roughly 1.5 hours ago. What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

- A. Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.
- B. Use synchronous database master-slave replication between two availability zones.
- C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored in S3 every 5 minutes.
- D. Take 15 minute DB backups stored in Glacier with transaction logs stored in S3 every 5 minutes.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 15

You are designing the network infrastructure for an application server in Amazon VPC. Users will access all application instances from the Internet, as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link.

How would you design routing to meet the above requirements?

- A. Configure a single routing table with a default route via the Internet gateway. Propagate a default route via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
- B. Configure a single routing table with a default route via the Internet gateway. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
- C. Configure a single routing table with two default routes: one to the Internet via an Internet gateway, the other to the on-premises network via the VPN gateway. Use this routing table across all subnets in the VPC.
- D. Configure two routing tables: one that has a default route via the Internet gateway, and other that has a default route via the VPN gateway. Associate both routing tables with each VPC subnet.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 16

You control access to S3 buckets and objects with:

- A. Identity and Access Management (IAM) Policies
- B. Access Control Lists (ACLs).
- C. Bucket Policies.
- D. All of the above

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:**QUESTION 17**

The AWS IT infrastructure that AWS provides, complies with the following IT security standards, including:

- A. SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II), SOC 2 and SOC 3
- B. FISMA, DIACAP, and FedRAMP
- C. PCI DSS Level 1, ISO 27001, ITAR and FIPS 140-2
- D. HIPAA, Cloud Security Alliance (CSA) and Motion Picture Association of America (MPAA)
- E. All of the above

Correct Answer: ABC

Section: (none)

Explanation

Explanation/Reference:**QUESTION 18**

Auto Scaling requests are signed with a _____ signature calculated from the request and the user's private key.

- A. SSL
- B. AES-256
- C. HMAC-SHA1
- D. X.509

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:**QUESTION 19**

The following policy can be attached to an IAM group. It lets an IAM user in that group access a "home directory" in AWS S3 that matches their user name using the console.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Action": ["s3:*"],  
      "Effect": "Allow",  
      "Resource": ["arn:aws:s3:::bucket-name"],  
      "Condition": {"StringLike": {"s3:prefix": ["home/${aws:username}/*"]}}  
    },  
    {  
      "Action": ["s3:*"],  
      "Effect": "Allow",  
      "Resource": ["arn:aws:s3:::bucket-name/home/${aws:username}/*"]  
    }  
  ]  
}
```

- A. True
- B. False



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 20

What does elasticity mean to AWS?

- A. The ability to scale computing resources up easily, with minimal friction and down with latency
- B. The ability to scale computing resources up and down easily, with minimal friction.
- C. The ability to provision cloud computing resources in expectation of future demand.
- D. The ability to recover from business continuity events with minimal friction.

Correct Answer: B

Section: (none)

Explanation**Explanation/Reference:****QUESTION 21**

The following are AWS Storage services? Choose 2 Answers

- A. AWS Relational Database Service (AWS RDS)
- B. AWS ElastiCache
- C. AWS Glacier
- D. AWS Import/Export

Correct Answer: BD

Section: (none)

Explanation**Explanation/Reference:****QUESTION 22**

How is AWS readily distinguished from other vendors in the traditional IT computing landscape?

- A. Experienced. Scalable and elastic. Secure. Cost-effective. Reliable
- B. Secure. Flexible. Cost-effective. Scalable and elastic. Global
- C. Secure. Flexible. Cost-effective. Scalable and elastic. Experienced
- D. Flexible. Cost-effective. Dynamic. Secure. Experienced.

Correct Answer: C

Section: (none)

Explanation**Explanation/Reference:****QUESTION 23**

You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached. The EC2 instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The four EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes), for a total of 16,000 random IOPS on the instance. The EC2 instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later, in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned

IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPs like the original four, for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all. What is the problem and a valid solution?

- A. The EBS-Optimized throughput limits the total IOPS that can be utilized; use an EBSOptimized instance that provides larger throughput.
- B. Small block sizes cause performance degradation, limiting the I/O throughput; configure the instance device driver and filesystem to use 64KB blocks to increase throughput.
- C. The standard EBS Instance root volume limits the total IOPS rate; change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.
- D. Larger storage volumes support higher Provisioned IOPS rates; increase the provisioned volume storage of each of the 6 EBS volumes to 1TB.
- E. RAID 0 only scales linearly to about 4 devices; use RAID 0 with 4 EBS Provisioned IOPS volumes, but increase each Provisioned IOPS EBS volume to 6,000 IOPS.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 24

Your company is storing millions of sensitive transactions across thousands of 100-GB files that must be encrypted in transit and at rest. Analysts concurrently depend on subsets of files, which can consume up to 5 TB of space, to generate simulations that can be used to steer business decisions. You are required to design an AWS solution that can cost effectively accommodate the long-term storage and in-flight subsets of data.

- A. Use Amazon Simple Storage Service (S3) with server-side encryption, and run simulations on subsets in ephemeral drives on Amazon EC2.
- B. Use Amazon S3 with server-side encryption, and run simulations on subsets in-memory on Amazon EC2.
- C. Use HDFS on Amazon EMR, and run simulations on subsets in ephemeral drives on Amazon EC2.
- D. Use HDFS on Amazon Elastic MapReduce (EMR), and run simulations on subsets in-memory on Amazon Elastic Compute Cloud (EC2).
- E. Store the full data set in encrypted Amazon Elastic Block Store (EBS) volumes, and regularly capture snapshots that can be cloned to EC2 workstations.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 25

Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer

wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?

What is the best approach to meet your customer's requirements?

- A. Send all the log events to Amazon SQS, setup an Auto Scaling group of EC2 servers to consume the logs and apply the heuristics.
- B. Send all the log events to Amazon Kinesis, develop a client process to apply heuristics on the logs
- C. Configure Amazon CloudTrail to receive custom logs, use EMR to apply heuristics the logs
- D. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3, use EMR to apply heuristics on the logs

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The throughput of an Amazon Kinesis stream is designed to scale without limits via increasing the number of shards within a stream. However, there are certain limits you should keep in mind while using Amazon Kinesis Streams:

By default, Records of a stream are accessible for up to 24 hours from the time they are added to the stream. You can raise this limit to up to 7 days by enabling extended data retention.

The maximum size of a data blob (the data payload before Base64-encoding) within one record is 1 megabyte (MB).

Each shard can support up to 1000 PUT records per second.

For more information about other API level limits, see Amazon Kinesis Streams Limits.

QUESTION 26

A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate Its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability.

Which is the most appropriate?

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RDS MySQL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
- E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tor the website on EC2 instances and use Route53 with DNS round-robin.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead:

Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)

Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the relevant JPEGs in S3. Your EC2 boxes just generate links to them.

If your app already serves it's own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily.

If you use S3, don't serve directly from the bucket - Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.

EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck.

Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.

Cache Layer: Adding CloudFront in front of S3 will help people on the other side of the world -- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)

You can also put CloudFront in front of your app, since your archive search results should be fairly static. This will also allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.

Database Layer: A few options:

Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.

Use RDS to run MySQL/Postgres

I'm not as familiar with ElasticSearch / Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.

App Layer:

When creating the app layer from scratch, consider CloudFormation and/or OpsWorks. It's extra stuff to learn, but helps down the road.

Java+Tomcat is right up the alley of ElasticBeanstalk. (Basically EC2 + Autoscale + ELB).

Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc.

Saving money: If you don't mind having downtime:

run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be really robust.

use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.)

Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.

Rewrite the application to use less memory and CPU - that way you can run on fewer/smaller boxes. (May or may not be worth the investment.)

If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though.

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use CloudFormation for all of it.

QUESTION 27

Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed. Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary. You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console.

Which option below will meet the needs for your NOC members?

- A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
- B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
- C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
- D. Use your on-premises SAML 2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 28

You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each account's bill to a Master AWS account using Consolidated Billing. To make sure you keep within budget, you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts.

Identify which option will allow you to achieve this goal.

- A. Create IAM users in the Master account with full Admin permissions. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from the Master account.
- B. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.
- C. Create IAM users in the Master account. Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.

D. Link the accounts using Consolidated Billing. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Bucket Owner Granting Cross-account Permission to objects It Does Not Own

In this example scenario, you own a bucket and you have enabled other AWS accounts to upload objects. That is, your bucket can have objects that other AWS accounts own.

Now, suppose as a bucket owner, you need to grant cross-account permission on objects, regardless of who the owner is, to a user in another account. For example, that user could be a billing application that needs to access object metadata. There are two core issues:

The bucket owner has no permissions on those objects created by other AWS accounts. So for the bucket owner to grant permissions on objects it does not own, the object owner, the AWS account that created the objects, must first grant permission to the bucket owner. The bucket owner can then delegate those permissions.

Bucket owner account can delegate permissions to users in its own account but it cannot delegate permissions to other AWS accounts, because cross-account delegation is not supported.

In this scenario, the bucket owner can create an AWS Identity and Access Management (IAM) role with permission to access objects, and grant another AWS account permission to assume the role temporarily enabling it to access objects in the bucket.

Background: Cross-Account Permissions and Using IAM Roles

IAM roles enable several scenarios to delegate access to your resources, and cross-account access is one of the key scenarios. In this example, the bucket owner, Account A, uses an IAM role to temporarily delegate object access cross-account to users in another AWS account, Account C. Each IAM role you create has two policies attached to it:

A trust policy identifying another AWS account that can assume the role.

An access policy defining what permissions—for example, `s3:GetObject`—are allowed when someone assumes the role. For a list of permissions you can specify in a policy, see [Specifying Permissions in a Policy](#).

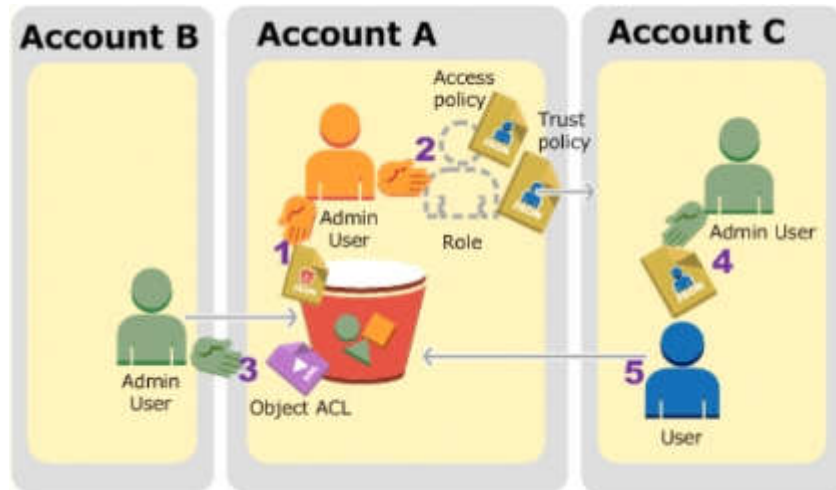
The AWS account identified in the trust policy then grants its user permission to assume the role. The user can then do the following to access objects:

Assume the role and, in response, get temporary security credentials.

Using the temporary security credentials, access the objects in the bucket.

For more information about IAM roles, go to [Roles \(Delegation and Federation\)](#) in [IAM User Guide](#).

The following is a summary of the walkthrough steps:



Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects. Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account A. The access policy attached to the role limits what user in Account C can do when the user accesses Account A. Account B administrator uploads an object to the bucket owned by Account A, granting full-control permission to the bucket owner.

Account C administrator creates a user and attaches a user policy that allows the user to assume the role.

User in Account C first assumes the role, which returns the user temporary security credentials. Using those temporary credentials, the user then accesses objects in the bucket.

For this example, you need three accounts. The following table shows how we refer to these accounts and the administrator users in these accounts. Per IAM guidelines (see About Using an Administrator User to Create Resources and Grant Permissions) we do not use the account root credentials in this walkthrough. Instead, you create an administrator user in each account and use those credentials in creating resources and granting them permissions

AWS Account ID	Account Referred To As	Administrator User in the Account
1111-1111-1111	Account A	AccountAdmin
2222-2222-2222	Account B	AccountBAdmin
3333-3333-3333	Account C	AccountCAdmin

QUESTION 29

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server. Users must be able to access portions of this data while the backups are taking place.

What backup solution would be most appropriate for this use case?

- A. Use Storage Gateway and configure it to use Gateway Cached volumes.
- B. Configure your backup software to use S3 as the target for your data backups.
- C. Configure your backup software to use Glacier as the target for your data backups.
- D. Use Storage Gateway and configure it to use Gateway Stored volumes.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

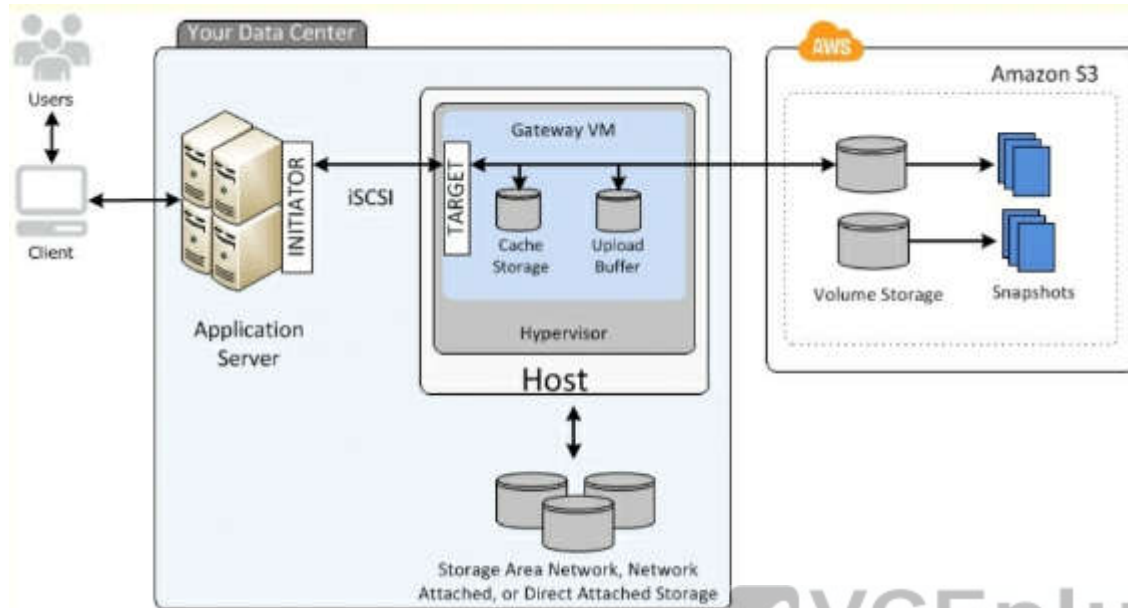
Gateway-Cached Volume Architecture

Gateway-cached volumes let you use Amazon Simple Storage Service (Amazon S3) as your primary data storage while retaining frequently accessed data locally in your storage gateway. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Your gateway stores data that you write to these volumes in Amazon S3 and retains recently read data in your on-premises storage gateway's cache and upload buffer storage.

Gateway-cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-cached volumes can support up to 32 volumes for a total maximum storage volume of 1,024 TiB (1 PiB).

In the gateway-cached volume solution, AWS Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3.

The following diagram provides an overview of the AWS Storage Gateway-cached volume deployment.



After you've installed the AWS Storage Gateway software appliance—the virtual machine (VM)—on a host in your data center and activated it, you can use the AWS Management Console to provision storage volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.

You also allocate disks on-premises for the VM. These on-premises disks serve the following purposes:

Disks for use by the gateway as cache storage - As your applications write data to the storage volumes in AWS, the gateway initially stores the data on the on-premises disks referred to as cache storage before uploading the data to Amazon S3. The cache storage acts as the on-premises durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

The cache storage also lets the gateway store your application's recently accessed data on-premises for low-latency access. If your application requests data, the gateway first checks the cache storage for the data before checking Amazon S3.

You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache storage. Cache storage should also be larger than the upload buffer. This latter guideline helps ensure cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.

Disks for use by the gateway as the upload buffer - To prepare for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an upload buffer. Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS, where it is stored encrypted in Amazon S3.

You can take incremental backups, called snapshots, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or one-time basis. When you delete a snapshot, only the data not needed for any other snapshots is removed.

You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance. All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using server-side

encryption (SSE). However, you cannot access this data with the Amazon S3 API or other tools such as the Amazon S3 console.

QUESTION 30

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization RIs. You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively?

- A. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand ml .large instances when triggered by Cloudwatch. Shut off c3.2xlarge instances.
- B. Configure ELB with two c3.2xlarge instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances. Shut off ml .large instances.
- C. Route traffic to EC2 ml .large and c3.2xlarge instances directly using Route 53 latency based routing and health checks. Shut off ELB.
- D. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



QUESTION 31

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Dunning a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

- A. Latency resource record sets cannot be used in combination with weighted resource record sets.
- B. You did not setup an HTTP health check to one or more of the weighted resource record sets associated with me disabled web servers.
- C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
- D. One of the two working web servers in the other region did not pass its HTTP health check.
- E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example com in the region where you disabled the servers.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

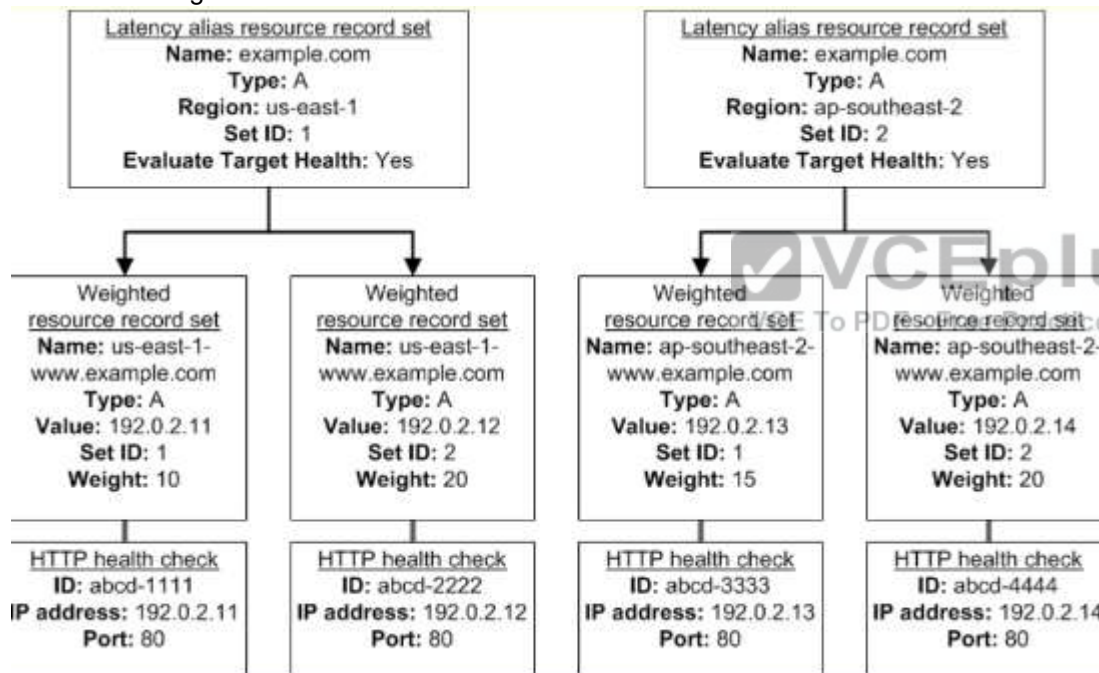
Explanation:

How Health Checks Work in Complex Amazon Route 53 Configurations

Checking the health of resources in complex configurations works much the same way as in simple configurations.

However, in complex configurations, you use a combination of alias resource record sets (including weighted alias, latency alias, and failover alias) and nonalias resource record sets to build a decision tree that gives you greater control over how Amazon Route 53 responds to requests. For more information, see How Health Checks Work in Simple Amazon Route 53 Configurations.

For example, you might use latency alias resource record sets to select a region close to a user and use weighted resource record sets for two or more resources within each region to protect against the failure of a single endpoint or an Availability Zone. The following diagram shows this configuration.

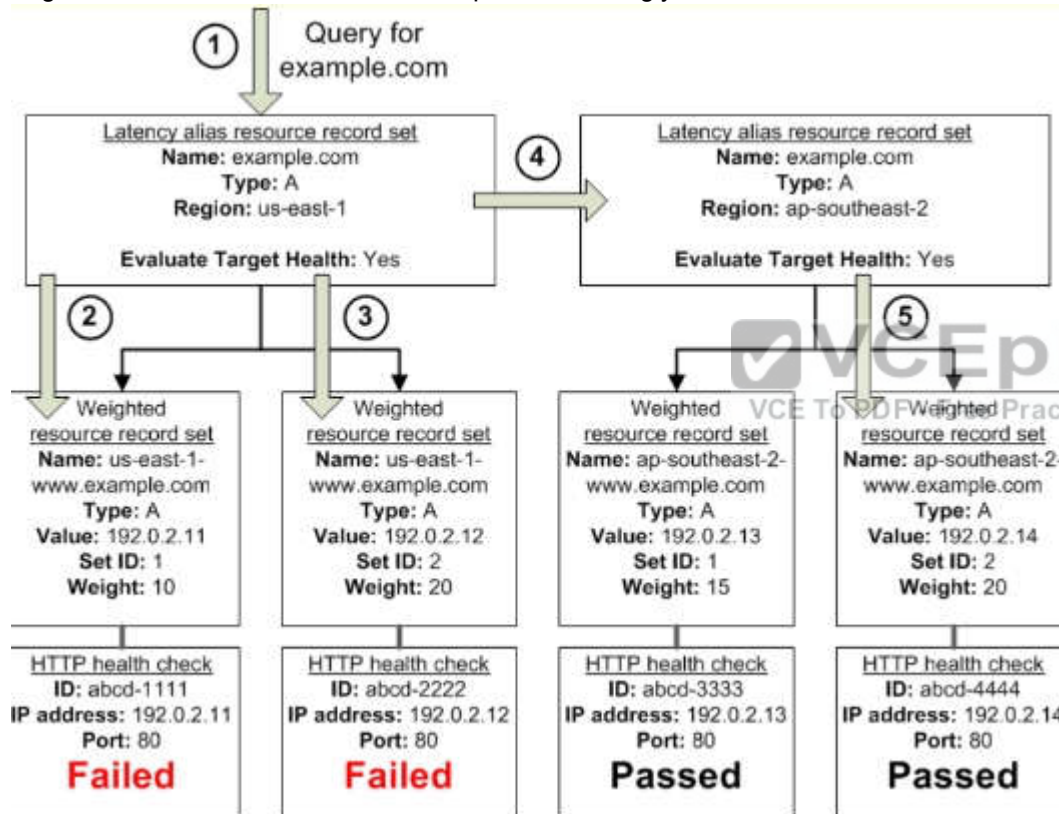


Here's how Amazon EC2 and Amazon Route 53 are configured:

You have Amazon EC2 instances in two regions, us-east-1 and ap-southeast-2. You want Amazon Route 53 to respond to queries by using the resource record sets in the region that provides the lowest latency for your customers, so you create a latency alias resource record set for each region. (You create the latency alias resource record sets after you create resource record sets for the individual Amazon EC2 instances.)

Within each region, you have two Amazon EC2 instances. You create a weighted resource record set for each instance.

The name and the type are the same for both of the weighted resource record sets in each region. When you have multiple resources in a region, you can create weighted or failover resource record sets for your resources. You can also create even more complex configurations by creating weighted alias or failover alias resource record sets that, in turn, refer to multiple resources. Each weighted resource record set has an associated health check. The IP address for each health check matches the IP address for the corresponding resource record set. This isn't required, but it's the most common configuration. For both latency alias resource record sets, you set the value of Evaluate Target Health to Yes. You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.



The preceding diagram illustrates the following sequence of events: Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 selects a weighted resource record set based on weight. Evaluate Target Health is Yes for the latency alias resource record set, so Amazon Route 53 checks the health of the selected weighted resource record set.

The health check failed, so Amazon Route 53 chooses another weighted resource record set based on weight and checks its health. That resource record set also is unhealthy.

Amazon Route 53 backs out of that branch of the tree, looks for the latency alias resource record set with the next-best latency, and chooses the resource record set for ap-southeast-2.

Amazon Route 53 again selects a resource record set based on weight, and then checks the health of the selected resource record set. The health check passed, so Amazon Route 53 returns the applicable value in response to the query. **What Happens When You Associate a Health Check with an Alias Resource Record Set?**

You can associate a health check with an alias resource record set instead of or in addition to setting the value of Evaluate Target Health to Yes. However, it's generally more useful if Amazon Route 53 responds to queries based on the health of the underlying resources—the HTTP servers, database servers, and other resources that your alias resource record sets refer to. For example, suppose the following configuration:

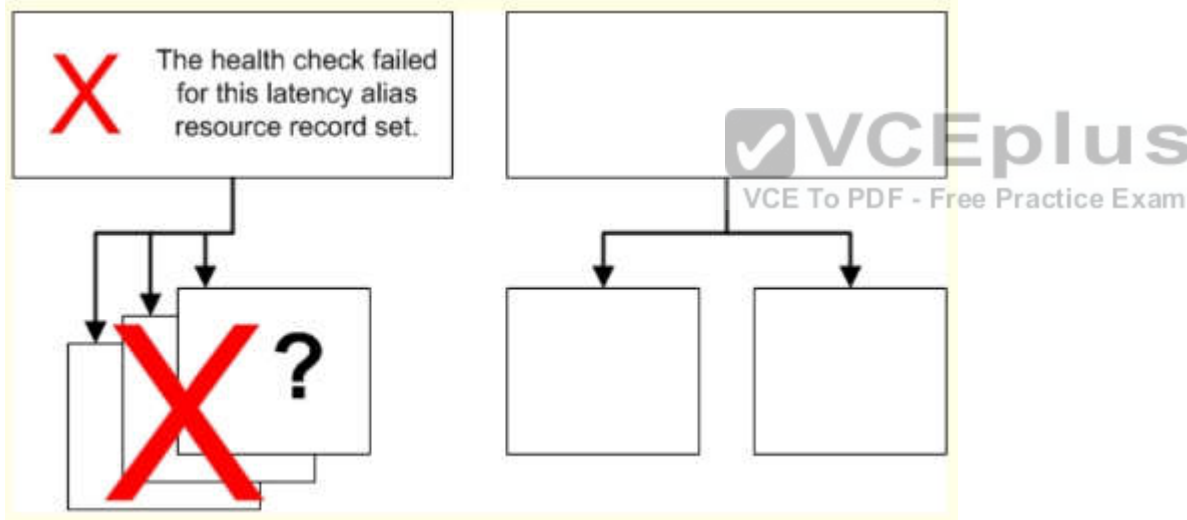
You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.

You set the value of Evaluate Target Health to Yes for the latency alias resource record set.

In this configuration, both of the following must be true before Amazon Route 53 will return the applicable value for a weighted resource record set:

The health check associated with the latency alias resource record set must pass.

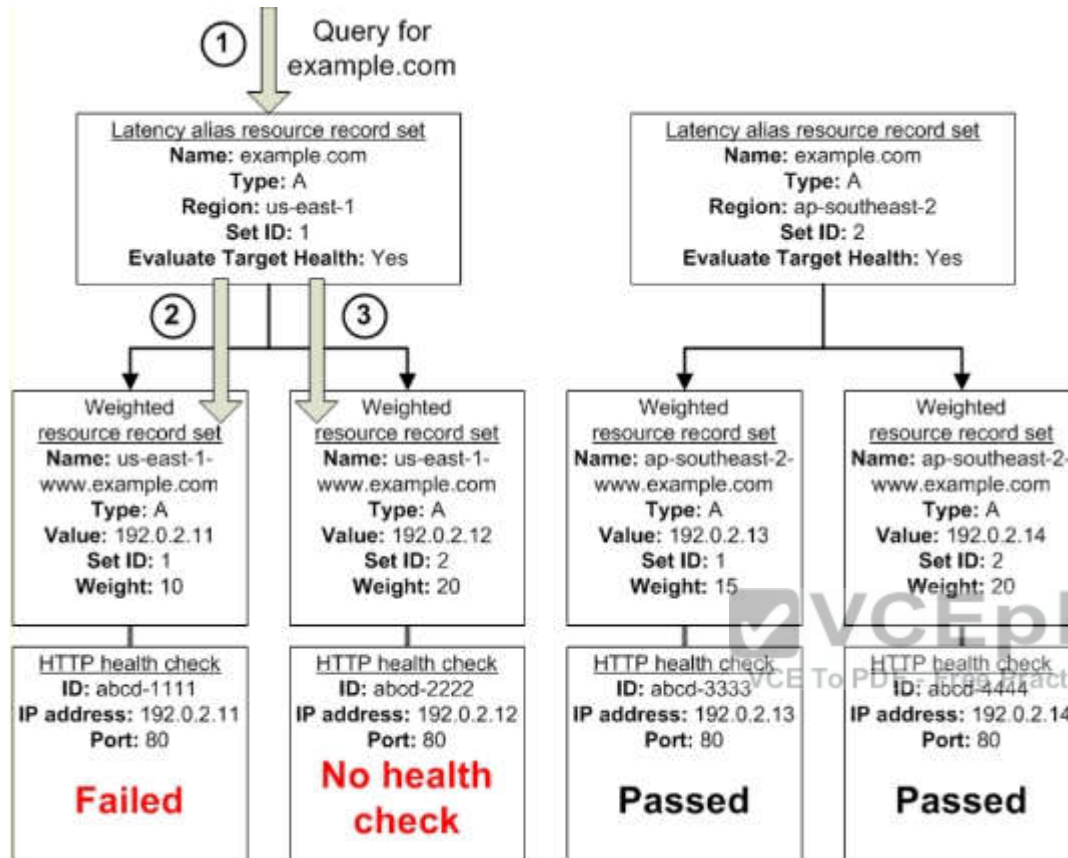
At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.



If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never looks past the failed health check on the alias resource record set.

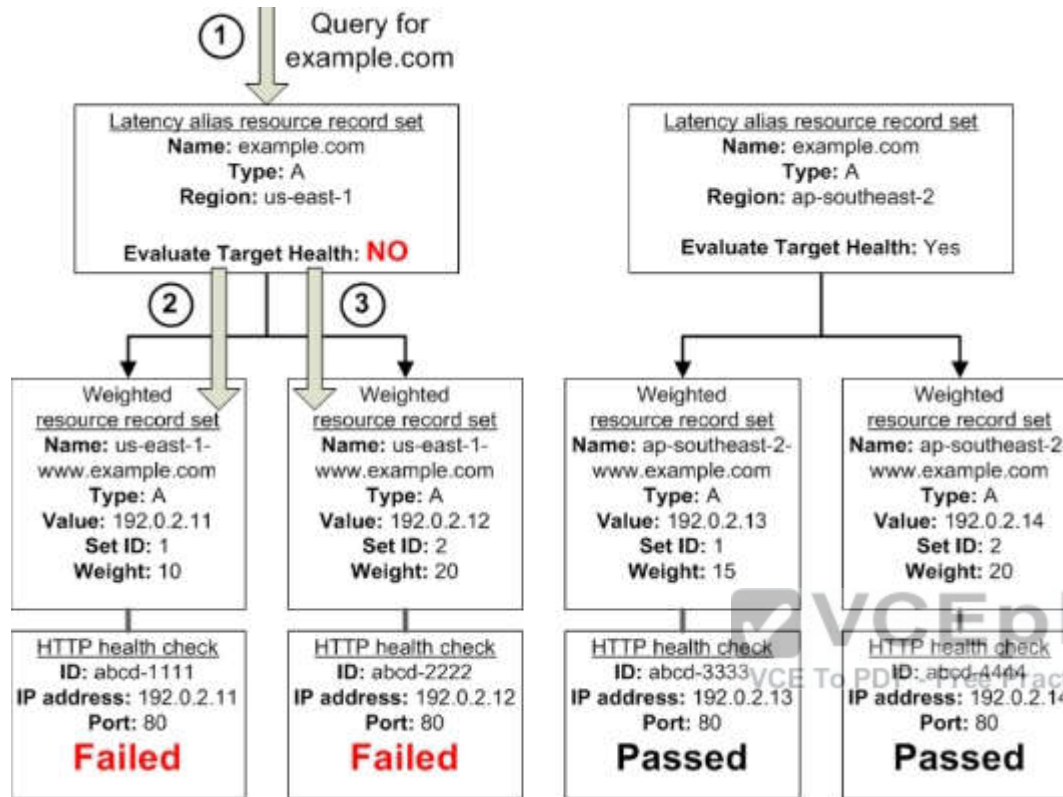
What Happens When You Omit Health Checks?

In a complex configuration, it's important to associate health checks with all of the non-alias resource record sets. Let's return to the preceding example, but assume that a health check is missing on one of the weighted resource record sets in the us-east-1 region:



Here's what happens when you omit a health check on a non-alias resource record set in this configuration: Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 looks up the alias target for the latency alias resource record set, and checks the status of the corresponding health checks. The health check for one weighted resource record set failed, so that resource record set is omitted from consideration. The other weighted resource record set in the alias target for the us-east-1 region has no health check. The corresponding resource might or might not be healthy, but without a health check, Amazon Route 53 has no way to know. Amazon Route 53 assumes that the resource is healthy and returns the applicable value in response to the query. What Happens When You Set Evaluate Target Health to No?

In general, you also want to set Evaluate Target Health to Yes for all of the alias resource record sets. In the following example, all of the weighted resource record sets have associated health checks, but Evaluate Target Health is set to No for the latency alias resource record set for the us-east-1 region:



Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration: Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 determines what the alias target is for the latency alias resource record set, and checks the corresponding health checks. They're both failing. Because the value of Evaluate Target Health is No for the latency alias resource record set for the us-east-1 region, Amazon Route 53 must choose one resource record set in this branch instead of backing out of the branch and looking for a healthy resource record set in the ap-southeast-2 region.

QUESTION 32

Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months. Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure. Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders. How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

- A. Add a business process management application to your Elastic Beanstalk app servers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
- B. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 Use the decider instance to send emails to customers.
- C. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 use SES to send emails to customers.
- D. Use an SQS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute them. Use SES to send emails to customers.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 33

A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically.

What AWS services should be used meet these requirements?

- A. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaimg group monitored with CloudWatch and RDS with read replicas.
- B. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and RDS with read replicas.
- C. Stateful instances for the web and application tier in an autoscaling group monitored with Cloudwatch. And multi-AZ RDS.
- D. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with Cloudwatch and multi-AZ RDS.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 34

You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket.

Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3. You want to configure security to handle potentially millions of users in the most secure manner possible.

What should your server-side application do when a new user registers on the photo-sharing mobile application?

- A. Create an IAM user. Update the bucket policy with appropriate permissions for the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- B. Create an IAM user. Assign appropriate permissions to the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- C. Create a set of long-term credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app and use them to access Amazon S3.
- D. Record the user's information in Amazon RDS and create a role in IAM with appropriate permissions. When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" function. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
- E. Record the user's information in Amazon DynamoDB. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



QUESTION 35

You are tasked with moving a legacy application from a virtual machine running inside your datacenter to an Amazon VPC. Unfortunately, this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse, there's no documentation for it. What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? (Choose 3 answers)

- A. An AWS Direct Connect link between the VPC and the network housing the internal services.
- B. An Internet Gateway to allow a VPN connection.
- C. An Elastic IP address on the VPC instance
- D. An IP address space that does not conflict with the one on-premises
- E. Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
- F. A VM Import of the current virtual machine

Correct Answer: ADF

Section: (none)

Explanation

Explanation/Reference: