

Congratulation!! You have selected the best success partner to become



Exam overview

Level: Associate

Length: 130 minutes to complete the exam

Cost: 150 USD

Visit [Exam pricing](#) for additional cost information.

Format: 65 questions, either multiple choice or multiple response

Delivery method: Pearson VUE and PSI; testing center or online proctored exam

<https://aws.amazon.com/certification/certified-solutions-architect-associate/>

QUESTION NO: 1

A solutions architect is designing a solution where users will be directed to a backup static error page if the primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53 where their domain is pointing to an Application Load Balancer (ALB).

Which configuration should the solutions architect use to meet the company's needs while minimizing changes and infrastructure overhead?

A.

Point a Route 53 alias record to an Amazon CloudFront distribution with the ALB as one of its origins. Then, create custom error pages for the distribution.

B.

Set up a Route 53 active-passive failover configuration. Direct traffic to a static error page hosted within an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.

C.

Update the Route 53 record to use a latency-based routing policy. Add the backup static error page hosted within an Amazon S3 bucket to the record so the traffic is sent to the most responsive endpoints.

D.

Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance hosting a static error page as endpoints. Route 53 will only send requests to the instance if the health checks fail for the ALB.

Answer: B

Explanation:

Active-passive failover

Use an active-passive failover configuration when you want a primary resource or group of resources to be available the majority of the time and you want a secondary resource or group of resources to be on standby in case all the primary resources become unavailable. When responding to queries, Route 53 includes only the healthy primary resources. If all the primary resources are unhealthy, Route 53 begins to include only the healthy secondary resources in response to DNS queries.

To create an active-passive failover configuration with one primary record and one secondary record, you just create the records and specify Failover for the routing policy. When the primary resource is healthy, Route 53 responds to DNS queries using the primary record. When the primary resource is unhealthy, Route 53 responds to DNS queries using the secondary record.

How Amazon Route 53 averts cascading failures

As a first defense against cascading failures, each request routing algorithm (such as weighted and failover) has a mode of last resort. In this special mode, when all records are considered unhealthy, the Route 53 algorithm reverts to considering all records healthy.

For example, if all instances of an application, on several hosts, are rejecting health check requests, Route 53 DNS servers will choose an answer anyway and return it rather than returning no DNS answer or returning an NXDOMAIN (non-existent domain) response. An application can respond to users but still fail health checks, so this provides some protection against misconfiguration.

Similarly, if an application is overloaded, and one out of three endpoints fails its health checks, so that it's excluded from Route 53 DNS responses, Route 53 distributes responses between the two remaining endpoints. If the remaining endpoints are unable to handle the additional load and they fail, Route 53 reverts to distributing requests to all three endpoints.

Reference:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html>

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-problems.html>

QUESTION NO: 2

A solutions architect is designing a high performance computing (HPC) workload on Amazon EC2. The EC2 instances need to communicate to each other frequently and require network performance with low latency and high throughput.

Which EC2 configuration meets these requirements?

- A.**
Launch the EC2 instances in a cluster placement group in one Availability Zone.
- B.**
Launch the EC2 instances in a spread placement group in one Availability Zone.
- C.**
Launch the EC2 instances in an Auto Scaling group in two Regions and peer the VPCs.
- D.**
Launch the EC2 instances in an Auto Scaling group spanning multiple Availability Zones.

Answer: A

Explanation:

Placement groups

When you launch a new EC2 instance, the EC2 service attempts to place the instance in such a way that all of your instances are spread out across underlying hardware to minimize correlated failures. You can use placement groups to influence the placement of a group of interdependent instances to meet the needs of your workload. Depending on the type of workload.

Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.

Reference:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

QUESTION NO: 3

A company wants to host a scalable web application on AWS. The application will be accessed by users from different geographic regions of the world. Application users will be able to download and upload unique data up to gigabytes in size. The development team wants a cost-effective solution to minimize upload and download latency and maximize performance.

What should a solutions architect do to accomplish this?

A.

Use Amazon S3 with Transfer Acceleration to host the application.

B.

Use Amazon S3 with CacheControl headers to host the application.

C.

Use Amazon EC2 with Auto Scaling and Amazon CloudFront to host the application.

D.

Use Amazon EC2 with Auto Scaling and Amazon ElastiCache to host the application.

Answer: C

Reference:

QUESTION NO: 4

A company is migrating from an on-premises infrastructure to the AWS Cloud. One of the company's applications stores files on a Windows file server farm that uses Distributed File System Replication (DFSR) to keep data in sync. A solutions architect needs to replace the file server farm.

Which service should the solutions architect use?

- A.**
Amazon Elastic File System (Amazon EFS)
- B.**
Amazon FSx
- C.**
Amazon S3
- D.**
AWS Storage Gateway

Answer: B

Explanation:

Migrating Existing Files to Amazon FSx for Windows File Server Using AWS DataSync

We recommend using AWS DataSync to transfer data between Amazon FSx for Windows File Server file systems. DataSync is a data transfer service that simplifies, automates, and accelerates moving and replicating data between on-premises storage systems and other AWS storage services over the internet or AWS Direct Connect. DataSync can transfer your file system data and metadata, such as ownership, time stamps, and access permissions.

Reference:

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html>

QUESTION NO: 5

A company has a legacy application that processes data in two parts. The second part of the process takes longer than the first, so the company has decided to rewrite the application as two microservices running on Amazon ECS that can scale independently.

How should a solutions architect integrate the microservices?

A.

Implement code in microservice 1 to send data to an Amazon S3 bucket. Use S3 event notifications to invoke microservice 2.

B.

Implement code in microservice 1 to publish data to an Amazon SNS topic. Implement code in microservice 2 to subscribe to this topic.

C.

Implement code in microservice 1 to send data to Amazon Kinesis Data Firehose. Implement code in microservice 2 to read from Kinesis Data Firehose.

D.

Implement code in microservice 1 to send data to an Amazon SQS queue. Implement code in microservice 2 to process messages from the queue.

Answer: D

Explanation:

QUESTION NO: 6

A company captures clickstream data from multiple websites and analyzes it using batch processing. The data is loaded nightly into Amazon Redshift and is consumed by business analysts. The company wants to move towards near-real-time data processing for timely insights. The solution should process the streaming data with minimal effort and operational overhead.

Which combination of AWS services are MOST cost-effective for this solution? (Choose two.)

A.

Amazon EC2

B.

AWS Lambda

- C. Amazon Kinesis Data Streams
- D. Amazon Kinesis Data Firehose
- E. Amazon Kinesis Data Analytics

Answer: B,D
Explanation:

Kinesis Data Streams and Kinesis Client Library (KCL) – Data from the data source can be continuously captured and streamed in near real-time using Kinesis Data Streams. With the Kinesis Client Library (KCL), you can build your own application that can preprocess the streaming data as they arrive and emit the data for generating incremental views and downstream analysis. Kinesis Data Analytics – This service provides the easiest way to process the data that is streaming through Kinesis Data Stream or Kinesis Data Firehose using SQL. This enables customers to gain actionable insight in near real-time from the incremental stream before storing it in Amazon S3.

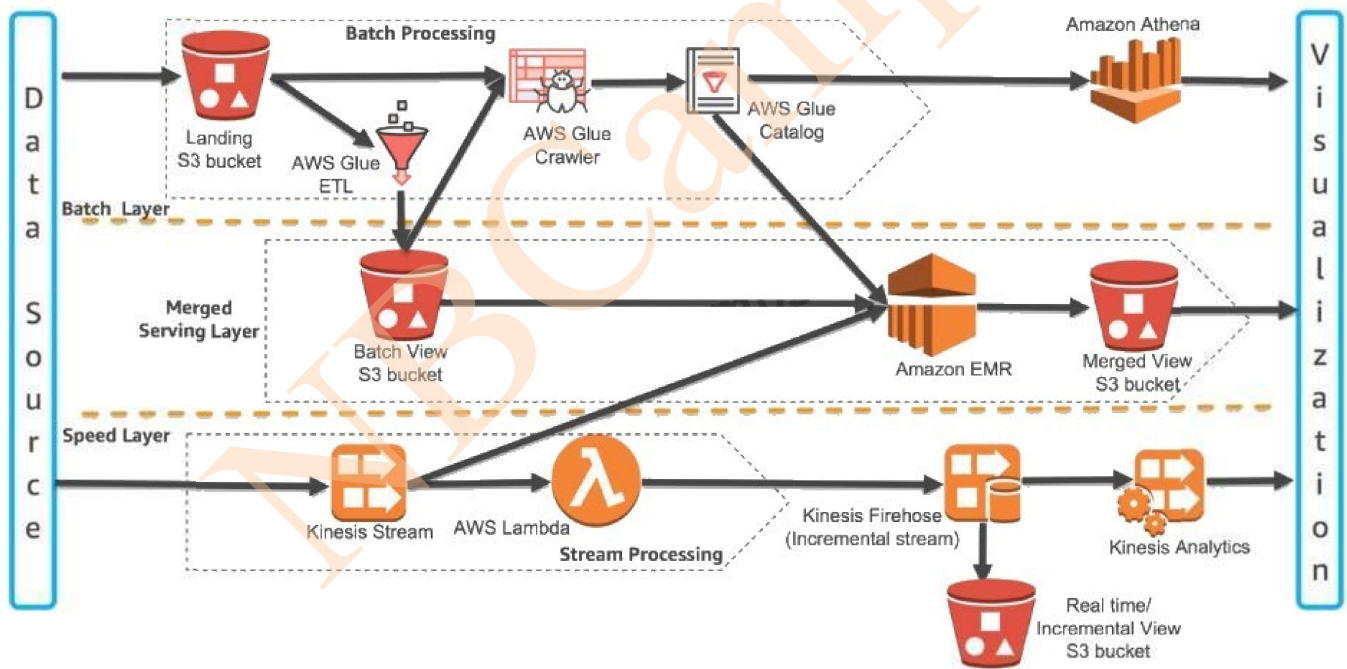


Figure 2: Lambda Architecture Building Blocks on AWS

Reference:

<https://d1.awsstatic.com/whitepapers/lambda-architecture-on-for-batch-aws.pdf>

QUESTION NO: 7

A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch executes. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application.

What should a solutions architect recommend to ensure the application is able to handle the workload and avoid downtime?

- A.**
Configure an Amazon CloudFront distribution in front of the ALB.
- B.**
Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
- C.**
Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
- D.**
Configure Amazon ElastiCache to remove some of the workload from the EC2 instances.

Answer: C

Explanation:

Scheduled Scaling for Amazon EC2 Auto Scaling

Scheduled scaling allows you to set your own scaling schedule. For example, let's say that every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling actions based on the predictable traffic patterns of your web application. Scaling actions are performed automatically as a function of time and date.

Reference:

https://docs.aws.amazon.com/autoscaling/ec2/userguide/schedule_time.html

QUESTION NO: 8

A company runs a multi-tier web application that hosts news content. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones and use an Amazon Aurora database. A solutions architect needs to make the application more resilient to periodic increases in request rates.

Which architecture should the solutions architect implement? (Choose two.)

A.

Add AWS Shield.

B.

Add Aurora Replica.

C.

Add AWS Direct Connect.

D.

Add AWS Global Accelerator.

E.

Add an Amazon CloudFront distribution in front of the Application Load Balancer.

Answer: D,E

Explanation:

AWS Global Accelerator

Acceleration for latency-sensitive applications

Many applications, especially in areas such as gaming, media, mobile apps, and financials, require very low latency for a great user experience. To improve the user experience, Global Accelerator directs user traffic to the application endpoint that is nearest to the client, which reduces internet latency and jitter. Global Accelerator routes traffic to the closest edge location by using Anycast, and then routes it to the closest regional endpoint over the AWS global network. Global Accelerator quickly reacts to changes in network performance to improve your users' application performance.

Amazon CloudFront

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment.

Reference:

<https://docs.aws.amazon.com/global-accelerator/latest/dg/introduction-benefits-of-migrating.html>

QUESTION NO: 9

An application running on AWS uses an Amazon Aurora Multi-AZ deployment for its database. When evaluating performance metrics, a solutions architect discovered that the database reads are causing high I/O and adding latency to the write requests against the database.

What should the solutions architect do to separate the read requests from the write requests?

- A.**
Enable read-through caching on the Amazon Aurora database.
- B.**
Update the application to read from the Multi-AZ standby instance.
- C.**
Create a read replica and modify the application to use the appropriate endpoint.
- D.**
Create a second Amazon Aurora database and link it to the primary database as a read replica.

Answer: C

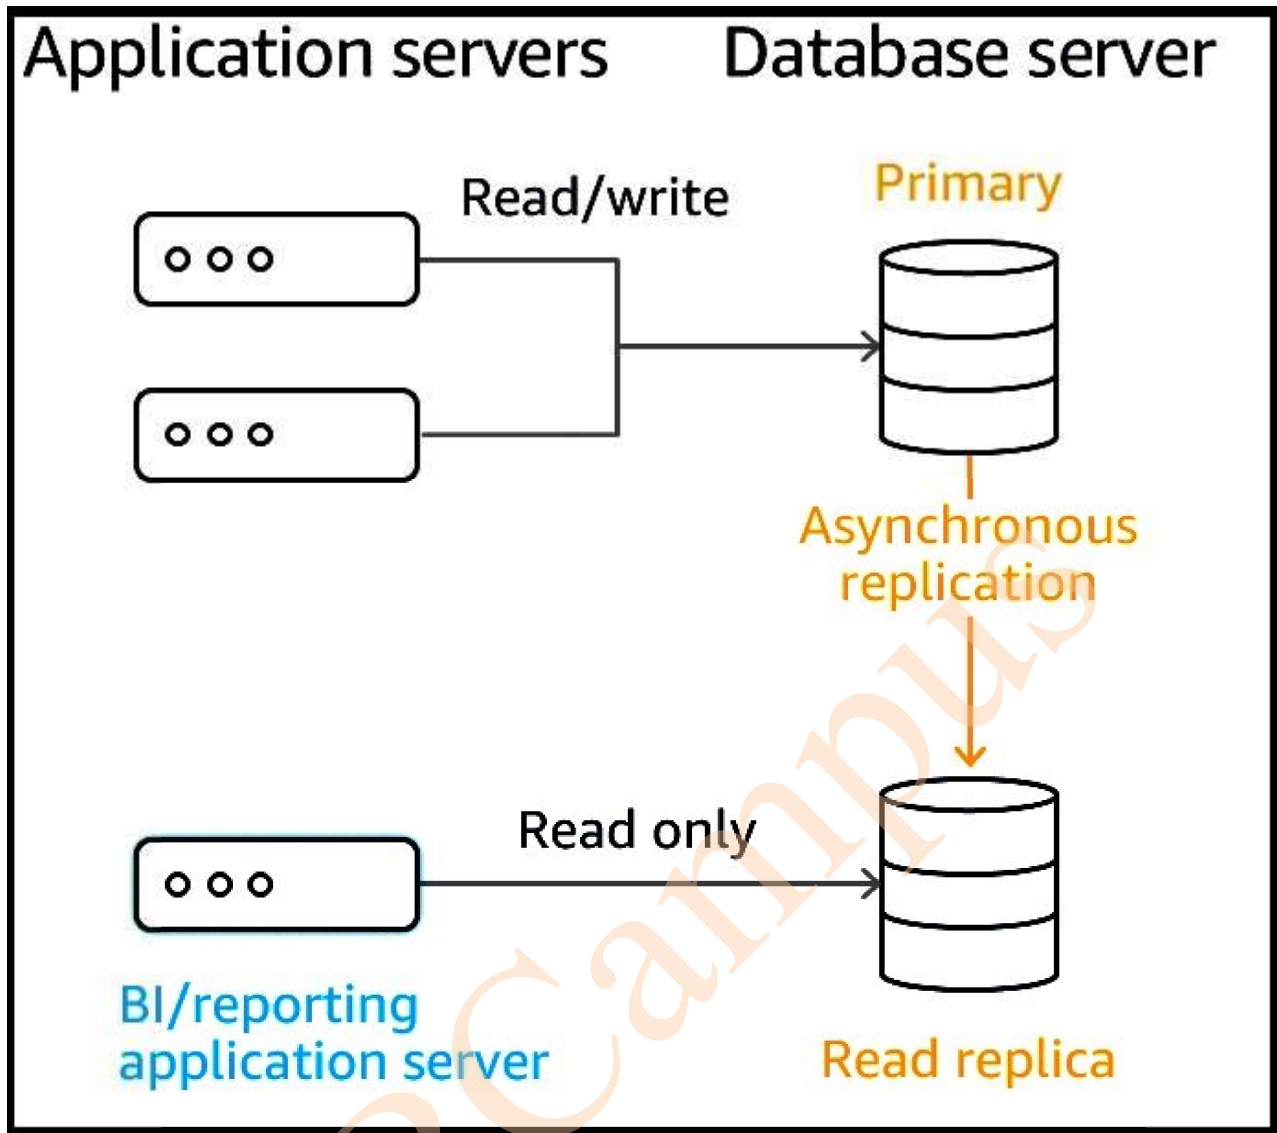
Explanation:

Amazon RDS Read Replicas

Amazon RDS Read Replicas provide enhanced performance and durability for RDS database (DB) instances. They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances. Read replicas are available in Amazon RDS for MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server as well as Amazon Aurora.

For the MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database engines, Amazon RDS creates a second DB instance using a snapshot of the source DB instance. It then uses the engines' native asynchronous replication to update the read replica whenever there is a change to the source DB instance. The read replica operates as a DB instance that allows only read-only connections; applications can connect to a read replica just as they would to any DB instance. Amazon RDS replicates all databases in the source DB instance.

Amazon Aurora further extends the benefits of read replicas by employing an SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora replicas share the same underlying storage as the source instance, lowering costs and avoiding the need to copy data to the replica nodes. For more information about replication with Amazon Aurora, see the [online documentation](#).



Reference:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

<https://aws.amazon.com/rds/features/read-replicas/>

QUESTION NO: 10

A recently acquired company is required to build its own infrastructure on AWS and migrate multiple applications to the cloud within a month. Each application has approximately 50 TB of data to be transferred. After the migration is complete, this company and its parent company will both require secure network connectivity with consistent throughput from their data centers to the applications. A solutions architect must ensure one-time data migration and ongoing network connectivity.