

A company needs to migrate a legacy application from an on-premises data center to the AWS Cloud because of hardware capacity constraints. The application runs 24 hours a day, 7 days a week. The application's database storage continues to grow over time.

What should a solutions architect do to meet these requirements MOST cost-effectively?

- A. Migrate the application layer to Amazon EC2 Spot Instances. Migrate the data storage layer to Amazon S3.
- B. Migrate the application layer to Amazon EC2 Reserved Instances. Migrate the data storage layer to Amazon RDS On-Demand Instances.
- C. Migrate the application layer to Amazon EC2 Reserved Instances. Migrate the data storage layer to Amazon Aurora Reserved Instances.
- D. Migrate the application layer to Amazon EC2 On-Demand Instances. Migrate the data storage layer to Amazon RDS Reserved Instances.

Correct Answer: C

Community vote distribution

C (86%)

14%

  **LuckyAro** Highly Voted 1 year, 8 months ago

Selected Answer: C

Amazon EC2 Reserved Instances allow for significant cost savings compared to On-Demand instances for long-running, steady-state workloads like this one. Reserved Instances provide a capacity reservation, so the instances are guaranteed to be available for the duration of the reservation period.

Amazon Aurora is a highly scalable, cloud-native relational database service that is designed to be compatible with MySQL and PostgreSQL. It can automatically scale up to meet growing storage requirements, so it can accommodate the application's database storage needs over time. By using Reserved Instances for Aurora, the cost savings will be significant over the long term.

upvoted 22 times

  **NolaH0la** Highly Voted 1 year, 8 months ago

Option B based on the fact that the DB storage will continue to grow, so on-demand will be a more suitable solution

upvoted 15 times

  **NolaH0la** 1 year, 8 months ago

Since the application's database storage is continuously growing over time, it may be difficult to estimate the appropriate size of the Aurora cluster in advance, which is required when reserving Aurora.

In this case, it may be more cost-effective to use Amazon RDS On-Demand Instances for the data storage layer. With RDS On-Demand Instances, you pay only for the capacity you use and you can easily scale up or down the storage as needed.



upvoted 5 times

  **hristni0** 1 year, 5 months ago

Answer is C. From Aurora Reserved Instances documentation:

If you have a DB instance, and you need to scale it to larger capacity, your reserved DB instance is automatically applied to your scaled DB instance. That is, your reserved DB instances are automatically applied across all DB instance class sizes. Size-flexible reserved DB instances are available for DB instances with the same AWS Region and database engine.

upvoted 2 times

  **Joxtat** 1 year, 8 months ago

The Answer is C.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.html>

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Database STORAGE will grow, not performance need (and required instance size).

upvoted 3 times

  **zdi561** Most Recent 3 months ago

Selected Answer: B

RI's are best for steady-state database workloads.



upvoted 1 times

  **hro** 7 months, 1 week ago

cost-effectively - the answer is C.

The application runs 24 hours a day, 7 days a week. The application's database storage continues to grow over time.

upvoted 3 times

  **MrPCarrot** 8 months, 1 week ago

Answer is C: Amazon EC2 Reserved Instances and Amazon Aurora Reserved Instances = less expensive than RDS.

upvoted 3 times


  **andyngkh86** 9 months, 1 week ago

Amazon Aurora reserved instances is used for the work load on predictable, so answer should be B  
upvoted 1 times

  **Priyapani** 9 months, 3 weeks ago



**Selected Answer: B**

I think it's B as database storage will grow  
upvoted 2 times

  **awsgeek75** 9 months, 3 weeks ago

Application runs 24x7 which means database is also used 24x7. The storage will grow and RDS On-Demand does not have auto-grow storage. You have to configure a storage size for RDS which means it will eventually run out of space. RDS On-Demand just scales CPU, not storage.

Aurora has no storage limitation and can scale storage according to need which is what is required here  
upvoted 4 times

  **Mikado211** 11 months ago

**Selected Answer: C**

24/7 forbids spot instances , so A is excluded  
Cost efficiency require reserved instances , so D is excluded  
Between RDS and Aurora, Aurora is less expensive thanks to the reserved instance, so B is finally excluded

Answer is C  
upvoted 4 times

  **cciesam** 1 year ago

**Selected Answer: B**

I hope it should be B considering Database growth  
upvoted 1 times

  **pentium75** 10 months, 1 week ago

Reserved instance applies to the DB instance size (CPU, RAM etc.), not storage.  
upvoted 3 times

  **Wayne23Fang** 1 year, 1 month ago

My research concludes that From pure price point of view Aurora Reserved might/ usually be slightly more expensive than On-demand RDS. But RDS has less Operation overhead. For the 24x7 nature, I would vote C. But for pure cost-effective, B is less costly.  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

This option involves migrating the application layer to Amazon EC2 Reserved Instances and migrating the data storage layer to Amazon Aurora Reserved Instances. Amazon EC2 Reserved Instances provide a significant discount (up to 75%) compared to On-Demand Instance pricing, making them a cost-effective choice for applications that have steady state or predictable usage. Similarly, Amazon Aurora Reserved Instances provide a significant discount (up to 69%) compared to On-Demand Instance pricing.  
upvoted 2 times

  **ajchi1980** 1 year, 4 months ago

**Selected Answer: C**



To meet the requirements of migrating a legacy application from an on-premises data center to the AWS Cloud in a cost-effective manner, the most suitable option would be:

C. Migrate the application layer to Amazon EC2 Reserved Instances. Migrate the data storage layer to Amazon Aurora Reserved Instances.

Explanation:

Migrating the application layer to Amazon EC2 Reserved Instances allows you to reserve EC2 capacity in advance, providing cost savings compared to On-Demand Instances. This is especially beneficial if the application runs 24/7.

Migrating the data storage layer to Amazon Aurora Reserved Instances provides cost optimization for the growing database storage needs. Amazon Aurora is a fully managed relational database service that offers high performance, scalability, and cost efficiency.  
upvoted 2 times

  **cpen** 1 year, 5 months ago

nnascncnscnknckl  
upvoted 1 times



  **TariqKipkemei** 1 year, 6 months ago

Answer is C  
upvoted 2 times

  **QuangPham810** 1 year, 6 months ago

Answer is C. Refer [https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER\\_WorkingWithReservedDBInstances.html](https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithReservedDBInstances.html) => Size-flexible reserved DB instances

upvoted 2 times

  **Abhineet9148232** 1 year, 7 months ago


**Selected Answer: C**

C: With Aurora Serverless v2, each writer and reader has its own current capacity value, measured in ACUs. Aurora Serverless v2 scales a writer or reader up to a higher capacity when its current capacity is too low to handle the load. It scales the writer or reader down to a lower capacity when its current capacity is higher than needed.

This is sufficient to accommodate the growing data changes.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.how-it-works.html#aurora-serverless-v2.how-it-works.scaling>

upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

**Selected Answer: C**

Typically Amazon RDS cost less than Aurora. But here, it's Aurora reserved.

upvoted 2 times

  **djgodzilla** 10 months ago

although agree and AWS wants you to choose Answer C. You can't convince a cloud accounting analyst that Aurora is cheaper than RDS. no matter what

upvoted 2 times

A university research laboratory needs to migrate 30 TB of data from an on-premises Windows file server to Amazon FSx for Windows File Server. The laboratory has a 1 Gbps network link that many other departments in the university share.

The laboratory wants to implement a data migration service that will maximize the performance of the data transfer. However, the laboratory needs to be able to control the amount of bandwidth that the service uses to minimize the impact on other departments. The data migration must take place within the next 5 days.

Which AWS solution will meet these requirements?

- A. AWS Snowcone
- B. Amazon FSx File Gateway
- C. AWS DataSync
- D. AWS Transfer Family

Correct Answer: C

Community vote distribution

C (97%)

  **kruasan**  2 years ago

Selected Answer: C

AWS DataSync is a data transfer service that can copy large amounts of data between on-premises storage and Amazon FSx for Windows File Server at high speeds. It allows you to control the amount of bandwidth used during data transfer.

- DataSync uses agents at the source and destination to automatically copy files and file metadata over the network. This optimizes the data transfer and minimizes the impact on your network bandwidth.
- DataSync allows you to schedule data transfers and configure transfer rates to suit your needs. You can transfer 30 TB within 5 days while controlling bandwidth usage.
- DataSync can resume interrupted transfers and validate data to ensure integrity. It provides detailed monitoring and reporting on the progress and performance of data transfers.

upvoted 26 times

  **kruasan** 2 years ago

Option A - AWS Snowcone is more suitable for physically transporting data when network bandwidth is limited. It would not complete the transfer within 5 days.

Option B - Amazon FSx File Gateway only provides access to files stored in Amazon FSx and does not perform the actual data migration from on-premises to FSx.

Option D - AWS Transfer Family is for transferring files over FTP, FTPS and SFTP. It may require scripting to transfer 30 TB and monitor progress, and lacks bandwidth controls.

upvoted 16 times

  **Michal\_L\_95**  2 years, 1 month ago

Selected Answer: C

As read a little bit, I assume that B (FSx File Gateway) requires a little bit more configuration rather than C (DataSync). From Stephane Maarek course explanation about DataSync:

An online data transfer service that simplifies, automates, and accelerates copying large amounts of data between on-premises storage systems and AWS Storage services, as well as between AWS Storage services.

You can use AWS DataSync to migrate data located on-premises, at the edge, or in other clouds to Amazon S3, Amazon EFS, Amazon FSx for Windows File Server, Amazon FSx for Lustre, Amazon FSx for OpenZFS, and Amazon FSx for NetApp ONTAP.

upvoted 12 times

  **MatAlves**  7 months, 3 weeks ago

Selected Answer: C

Even if we allocate 60% of total bandwidth for the transfer, that would take 5d2h. Considering that " many other departments in the university share", that wouldn't be feasible.

Ref. <https://expedient.com/knowledgebase/tools-and-calculators/file-transfer-time-calculator/>

On the other hand, snowcone isn't also a great option, because "you will receive the Snowcone device in approximately 4-6 days".

Ref.

<https://aws.amazon.com/snowcone/faqs/#:~:text=You%20will%20receive%20the%20Snowcone,console%20for%20each%20Snowcone%20device.>

upvoted 2 times

  **MatAlves** 7 months, 3 weeks ago

Very unreasonable scenario. The least "bad" option is C, but that will definitely affect users in real production environments.

upvoted 2 times

🗄️ 👤 **Cyberkayu** 1 year, 4 months ago

**Selected Answer: A**

Snow cone can support up to 8TB for HDD and 15TB for each SSD devices. Shipped within 4-6 days. Data migration can begin on next 5 days.

Does not use any amount of bandwidth and impact the production network. Device came with 1G and 10G Base-T Ethernet port. That's the Maximum performance in data transfer. defined in the question.

upvoted 2 times

🗄️ 👤 **JA2018** 5 months, 2 weeks ago

Key in STEM:

The data migration must take place within the next 5 days.

upvoted 1 times

🗄️ 👤 **JA2018** 5 months, 2 weeks ago

Snowcone it's just 15TB per SSD and it takes 4-6 business days to deliver to the customer's location.....

It could takes another 4-6 business days to return the snowcone device back to AWS (to transfer data to the designated S3 bucket) after you completed the on-prems data transfer to the device.

upvoted 1 times

🗄️ 👤 **AZ\_Master** 1 year, 5 months ago

**Selected Answer: C**

Bandwidth control = Data Sync

<https://docs.aws.amazon.com/datasync/latest/userguide/configure-bandwidth.html>

upvoted 5 times

🗄️ 👤 **Ruffyit** 1 year, 5 months ago

Bandwidth Optimization and Control

Transferring hot or cold data should not impede your business. DataSync is equipped with granular controls to optimize bandwidth consumptions.

Throttle transfer speeds up to 10 Gbps during off hours and set limits when network availability is needed elsewhere

upvoted 2 times

🗄️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

C. AWS DataSync

upvoted 2 times

🗄️ 👤 **Nikki013** 1 year, 8 months ago

**Selected Answer: C**

<https://aws.amazon.com/datasync/features/>

upvoted 2 times

🗄️ 👤 **Yousuf\_Ibrahim** 1 year, 6 months ago

Bandwidth Optimization and Control

Transferring hot or cold data should not impede your business. DataSync is equipped with granular controls to optimize bandwidth consumptions. Throttle transfer speeds up to 10 Gbps during off hours and set limits when network availability is needed elsewhere.

upvoted 2 times

🗄️ 👤 **jayce5** 1 year, 10 months ago

**Selected Answer: C**

"Amazon FSx File Gateway" is for storing data, not for migrating. So the answer should be C.

upvoted 3 times

🗄️ 👤 **ACloud\_Guru15** 1 year, 5 months ago

Thanks for the explanation

upvoted 2 times

🗄️ 👤 **shanwford** 2 years ago

**Selected Answer: C**

Snowcone to small and delivertime to long. With DataSync you can set bandwidth limits - so this is fine solution.

upvoted 3 times

🗄️ 👤 **MaxMa** 2 years, 1 month ago

Why not B?

upvoted 1 times

🗄️ 👤 **Guru4Cloud** 1 year, 7 months ago

Transferring will be much longerm rather then 5 days as required.

upvoted 2 times

🗄️ 👤 **AlessandraSAA** 2 years, 1 month ago


A not possible because Snowcone it's just 8TB and it takes 4-6 business days to deliver

B why cannot be <https://aws.amazon.com/storagegateway/file/fsx/>?

C I don't really get this  
D cannot be because not compatible - <https://aws.amazon.com/aws-transfer-family/>  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

With B you cannot "control the amount of bandwidth that the service uses", while C does exactly what is required here.  
upvoted 2 times

  **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: C**

Voting C  
upvoted 2 times

  **Bhawesh** 2 years, 2 months ago

**Selected Answer: C**

C. - DataSync is Correct.  
A. Snowcone is incorrect. The question says data migration must take place within the next 5 days.AWS says: If you order, you will receive the Snowcone device in approximately 4-6 days.  
upvoted 3 times

  **LuckyAro** 2 years, 2 months ago


**Selected Answer: C**

DataSync can be used to migrate data between on-premises Windows file servers and Amazon FSx for Windows File Server with its compatibility for Windows file systems.

The laboratory needs to migrate a large amount of data (30 TB) within a relatively short timeframe (5 days) and limit the impact on other departments' network traffic. Therefore, AWS DataSync can meet these requirements by providing fast and efficient data transfer with network throttling capability to control bandwidth usage.  
upvoted 5 times

  **cloudbusting** 2 years, 2 months ago

<https://docs.aws.amazon.com/datasync/latest/userguide/configure-bandwidth.html>  
upvoted 3 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: C**

<https://aws.amazon.com/datasync/>  
upvoted 3 times

A company wants to create a mobile app that allows users to stream slow-motion video clips on their mobile devices. Currently, the app captures video clips and uploads the video clips in raw format into an Amazon S3 bucket. The app retrieves these video clips directly from the S3 bucket. However, the videos are large in their raw format.

Users are experiencing issues with buffering and playback on mobile devices. The company wants to implement solutions to maximize the performance and scalability of the app while minimizing operational overhead.

Which combination of solutions will meet these requirements? (Choose two.)

- A. Deploy Amazon CloudFront for content delivery and caching.
- B. Use AWS DataSync to replicate the video files across AW'S Regions in other S3 buckets.
- C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats.
- D. Deploy an Auto Sealing group of Amazon EC2 instances in Local Zones for content delivery and caching.
- E. Deploy an Auto Scaling group of Amazon EC2 instances to convert the video files to more appropriate formats.

Correct Answer: A

Community vote distribution

A (55%)

C (43%)

- pentium75

Highly Voted

1 year, 4 months ago

F - Fire the guy who created the current design

upvoted 21 times
- awsgeek75

1 year, 3 months ago

No, make him watch all those videos with buffering!

upvoted 12 times
- Bhawesh

Highly Voted

2 years, 2 months ago

For Minimum operational overhead, the 2 options A,C should be correct.

A. Deploy Amazon CloudFront for content delivery and caching.

C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats.

upvoted 21 times
- surajkrishnamurthy

Most Recent

3 months, 1 week ago

Selected Answer: C

Not able to select Multi Option

Admin please check

upvoted 1 times
- Rcosmos

3 months, 2 weeks ago

Selected Answer: U

A melhor combinação de soluções para atender aos requisitos é:

A. Implante o Amazon CloudFront para entrega de conteúdo e armazenamento em cache.

C. Use o Amazon Elastic Transcoder para converter os arquivos de vídeo em formatos mais apropriados.

upvoted 1 times
- JA2018

5 months, 2 weeks ago

Selected Answer: C

Which combination of solutions will meet these requirements? (Choose two.)

.....

upvoted 1 times
- Mayank0502

10 months, 1 week ago

Selected Answer: C

A & C. Admin has almost every answer wrong

upvoted 3 times
- xyGGXH

1 year, 2 months ago

Selected Answer: A

A&C is correct



upvoted 4 times

  **db95476** 1 year, 4 months ago

**Selected Answer: A**

A and C

upvoted 3 times

  **Ruffyt** 1 year, 5 months ago

For Minimum operational overhead, the 2 options A,C should be correct.

A. Deploy Amazon CloudFront for content delivery and caching.

C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats.

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

For Minimum operational overhead, the 2 options A,C should be correct.

A. Deploy Amazon CloudFront for content delivery and caching.



C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats.

upvoted 2 times

  **Guru4Cloud** 1 year, 7 months ago

examtopics team, please fix this question, please allow to select two answer

upvoted 3 times

  **vincent2023** 3 months, 2 weeks ago

It has been more than 1year and 3 months, but the issue hasn't been fixed.

upvoted 1 times

  **jacob\_ho** 1 year, 8 months ago

Elastic Transcoder has been deprecated, and AWS encourage to use AWS Elemental MediaConvert right now:

<https://aws.amazon.com/blogs/media/how-to-migrate-workflows-from-amazon-elastic-transcoder-to-aws-elemental-mediaconvert/>

upvoted 7 times

  **enc\_0343** 1 year, 10 months ago

**Selected Answer: A**

AC is the correct answer



upvoted 2 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: A**

AC, the only possible answers.

upvoted 2 times

  **Eden** 1 year, 12 months ago

It says choose two so I chose AC

upvoted 2 times

  **WheracanIstart** 2 years, 1 month ago

**Selected Answer: C**

A & C are the right answers.

upvoted 5 times

  **kampatra** 2 years, 1 month ago

**Selected Answer: A**

Correct answer: AC

upvoted 4 times

  **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: C**

A and C. Transcoder does exactly what this needs.

upvoted 3 times

A company is launching a new application deployed on an Amazon Elastic Container Service (Amazon ECS) cluster and is using the Fargate launch type for ECS tasks. The company is monitoring CPU and memory usage because it is expecting high traffic to the application upon its launch. However, the company wants to reduce costs when utilization decreases.

What should a solutions architect recommend?

- A. Use Amazon EC2 Auto Scaling to scale at certain periods based on previous traffic patterns.
- B. Use an AWS Lambda function to scale Amazon ECS based on metric breaches that trigger an Amazon CloudWatch alarm.
- C. Use Amazon EC2 Auto Scaling with simple scaling policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm.
- D. Use AWS Application Auto Scaling with target tracking policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm.

**Correct Answer: D**

*Community vote distribution*

D (100%)

- rrharris**

Highly Voted

1 year, 8 months ago

Answer is D - Auto-scaling with target tracking

upvoted 14 times
- Joxtat**

Highly Voted

1 year, 8 months ago

Selected Answer: D

<https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto-scaling.html>

upvoted 5 times
- phuonglai**

Most Recent

8 months, 3 weeks ago

Selected Answer: D

<https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto-scaling.html>

upvoted 4 times
- awsgeek75**

9 months, 2 weeks ago

Selected Answer: D

<https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto-scaling.html>

upvoted 2 times
- pentium75**

10 months, 1 week ago

Selected Answer: D

This is running on Fargate, so EC2 scaling (A and C) is out. Lambda (B) is too complex.

upvoted 4 times
- TariqKipkemei**

1 year ago

Target tracking will scale in/out the ECS cluster to maintain the average CPU utilization to a set value. e.g. <<<50%>>> Scale out when average CPU utilization is above 50% until average CPU utilization is back to 50%. And scale in when average CPU utilization is below 50% until average CPU utilization is back to 50%.

upvoted 5 times
- TariqKipkemei**

1 year ago

Answer is D

upvoted 2 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: D

Answer is D - Auto-scaling with target tracking

upvoted 2 times
- vincent2023**

3 months, 2 weeks ago



Explanation please !!!

upvoted 1 times
- TariqKipkemei**

1 year, 6 months ago

Answer is D - Application Auto Scaling is a web service for developers and system administrators who need a solution for automatically scaling their scalable resources for individual AWS services beyond Amazon EC2.

upvoted 4 times

  **boxu03** 1 year, 7 months ago

**Selected Answer: D**

should be D  
upvoted 2 times

  **jahmad0730** 1 year, 8 months ago

**Selected Answer: D**

Answer is D  
upvoted 3 times

  **Neha999** 1 year, 8 months ago

D : auto-scaling with target tracking  
upvoted 5 times

A company recently created a disaster recovery site in a different AWS Region. The company needs to transfer large amounts of data back and forth between NFS file systems in the two Regions on a periodic basis.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS DataSync.
- B. Use AWS Snowball devices.
- C. Set up an SFTP server on Amazon EC2.
- D. Use AWS Database Migration Service (AWS DMS).

Correct Answer: A

Community vote distribution

A (100%)

  **LuckyAro** Highly Voted 1 year, 2 months ago

Selected Answer: A



AWS DataSync is a fully managed data transfer service that simplifies moving large amounts of data between on-premises storage systems and AWS services. It can also transfer data between different AWS services, including different AWS Regions. DataSync provides a simple, scalable, and automated solution to transfer data, and it minimizes the operational overhead because it is fully managed by AWS.

upvoted 18 times

  **Ruffyt** Most Recent 5 months, 3 weeks ago

AWS DataSync is a fully managed data transfer service that simplifies moving large amounts of data between on-premises storage systems and AWS services. It can also transfer data between different AWS services, including different AWS Regions. DataSync provides a simple, scalable, and automated solution to transfer data, and it minimizes the operational overhead because it is fully managed by AWS.

upvoted 2 times

  **Salilgen** 4 months, 3 weeks ago

You cannot use DataSync to transfer across Regions with an NFS

<https://docs.aws.amazon.com/datasync/latest/userguide/working-with-locations.html#working-with-locations-cross-regions>

upvoted 1 times

  **TariqKipkemei** 7 months ago

Selected Answer: A

Use AWS DataSync

upvoted 2 times

  **Guru4Cloud** 8 months ago

Selected Answer: A

Use AWS DataSync.

upvoted 2 times

  **kruasan** 1 year ago

Selected Answer: A

- AWS DataSync is a data transfer service optimized for moving large amounts of data between NFS file systems. It can automatically copy files and metadata between your NFS file systems in different AWS Regions.
- DataSync requires minimal setup and management. You deploy a source and destination agent, provide the source and destination locations, and DataSync handles the actual data transfer efficiently in the background.
- DataSync can schedule and monitor data transfers to keep source and destination in sync with minimal overhead. It resumes interrupted transfers and validates data integrity.
- DataSync optimizes data transfer performance across AWS's network infrastructure. It can achieve high throughput with minimal impact to your operations.

upvoted 4 times

  **kruasan** 1 year ago

Option B - AWS Snowball requires physical devices to transfer data. This incurs overhead to transport devices and manually load/unload data. It is not an online data transfer solution.

Option C - Setting up and managing an SFTP server would require provisioning EC2 instances, handling security groups, and writing scripts to automate the data transfer - all of which demand more overhead than DataSync.

Option D - AWS Database Migration Service is designed for migrating databases, not general file system data. It would require converting your NFS data into a database format, incurring additional overhead.

upvoted 4 times

  **ashu089** 1 year, 1 month ago

Selected Answer: A

A only  
upvoted 2 times

  **skiwili** 1 year, 2 months ago

Selected Answer: A

Aaaaaa  
upvoted 2 times

  **NolaH0la** 1 year, 2 months ago

A should be correct  
upvoted 2 times

A company is designing a shared storage solution for a gaming application that is hosted in the AWS Cloud. The company needs the ability to use SMB clients to access data. The solution must be fully managed.

Which AWS solution meets these requirements?

- A. Create an AWS DataSync task that shares the data as a mountable file system. Mount the file system to the application server.
- B. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the application server to the file share.
- C. Create an Amazon FSx for Windows File Server file system. Attach the file system to the origin server. Connect the application server to the file system.
- D. Create an Amazon S3 bucket. Assign an IAM role to the application to grant access to the S3 bucket. Mount the S3 bucket to the application server.

**Correct Answer: C**

*Community vote distribution*

C (100%)

- rrharris**

Highly Voted

1 year, 8 months ago

Answer is C - SMB = storage gateway or FSx

upvoted 9 times
- Neha999**

Highly Voted

1 year, 8 months ago

C L: Amazon FSx for Windows File Server file system

upvoted 7 times
- 8883b6c**

Most Recent

4 months ago

Selected Answer: C

SMB -> FSx

upvoted 1 times
- phuonglai**

8 months, 3 weeks ago

Selected Answer: C

SMB -> FSx

upvoted 4 times
- TariqKipkemei**

1 year ago

Selected Answer: C

SMB = FSx for Windows File Server

upvoted 4 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: C

Answer is C - SMB = storage gateway or FSx

upvoted 3 times
- kruasan**

1 year, 6 months ago

Selected Answer: C

- Amazon FSx for Windows File Server provides a fully managed native Windows file system that can be accessed using the industry-standard SMB protocol. This allows Windows clients like the gaming application to directly access file data.
  - FSx for Windows File Server handles time-consuming file system administration tasks like provisioning, setup, maintenance, file share management, backups, security, and software patching - reducing operational overhead.
  - FSx for Windows File Server supports high file system throughput, IOPS, and consistent low latencies required for performance-sensitive workloads. This makes it suitable for a gaming application.
  - The file system can be directly attached to EC2 instances, providing a performant shared storage solution for the gaming servers.

upvoted 5 times
- kruasan**

1 year, 6 months ago

Option A - DataSync is for data transfer, not providing a shared file system. It cannot be mounted or directly accessed.

Option B - A self-managed EC2 file share would require manually installing, configuring and maintaining a Windows file system and share. This demands significant overhead to operate.

Option D - Amazon S3 is object storage, not a native file system. The data in S3 would need to be converted/formatted to provide file share access, adding complexity. S3 cannot be directly mounted or provide the performance of FSx.

upvoted 5 times

  **elearningtakai** 1 year, 7 months ago

**Selected Answer: C**

Amazon FSx for Windows File Server  
upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago



**Selected Answer: C**

I vote C since FSx supports SMB  
upvoted 2 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: C**

AWS FSx for Windows File Server is a fully managed native Microsoft Windows file system that is accessible through the SMB protocol. It provides features such as file system backups, integrated with Amazon S3, and Active Directory integration for user authentication and access control. This solution allows for the use of SMB clients to access the data and is fully managed, eliminating the need for the company to manage the underlying infrastructure.  
upvoted 3 times

  **Babba** 1 year, 8 months ago

**Selected Answer: C**

C for me  
upvoted 2 times

A company wants to run an in-memory database for a latency-sensitive application that runs on Amazon EC2 instances. The application processes more than 100,000 transactions each minute and requires high network throughput. A solutions architect needs to provide a cost-effective network design that minimizes data transfer charges.

Which solution meets these requirements?

- A. Launch all EC2 instances in the same Availability Zone within the same AWS Region. Specify a placement group with cluster strategy when launching EC2 instances.
- B. Launch all EC2 instances in different Availability Zones within the same AWS Region. Specify a placement group with partition strategy when launching EC2 instances.
- C. Deploy an Auto Scaling group to launch EC2 instances in different Availability Zones based on a network utilization target.
- D. Deploy an Auto Scaling group with a step scaling policy to launch EC2 instances in different Availability Zones.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **kruasan** Highly Voted 1 year, 6 months ago

**Selected Answer: A**

Reasons:

- Launching instances within a single AZ and using a cluster placement group provides the lowest network latency and highest bandwidth between instances. This maximizes performance for an in-memory database and high-throughput application.
- Communications between instances in the same AZ and placement group are free, minimizing data transfer charges. Inter-AZ and public IP traffic can incur charges.
- A cluster placement group enables the instances to be placed close together within the AZ, allowing the high network throughput required. Partition groups span AZs, reducing bandwidth.
- Auto Scaling across zones could launch instances in AZs that increase data transfer charges. It may reduce network throughput, impacting performance.

upvoted 23 times

  **kruasan** 1 year, 6 months ago

In contrast:

Option B - A partition placement group spans AZs, reducing network bandwidth between instances and potentially increasing costs.

Option C - Auto Scaling alone does not guarantee the network throughput and cost controls required for this use case. Launching across AZs could increase data transfer charges.

Option D - Step scaling policies determine how many instances to launch based on metrics alone. They lack control over network connectivity and costs between instances after launch.

upvoted 12 times

  **TariqKipkemei** Highly Voted 1 year ago

**Selected Answer: A**

Cluster placement group packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance.

upvoted 5 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

**Selected Answer: A**

Apart from the fact that BCD distribute the instances across AZ which is bad for inter-node network latency, I think the following article is really useful in understanding A:

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

upvoted 2 times

  **Ruffyit** 11 months, 3 weeks ago

- Launching instances within a single AZ and using a cluster placement group provides the lowest network latency and highest bandwidth between instances. This maximizes performance for an in-memory database and high-throughput application.
- Communications between instances in the same AZ and placement group are free, minimizing data transfer charges. Inter-AZ and public IP traffic can incur charges.
- A cluster placement group enables the instances to be placed close together within the AZ, allowing the high network throughput required. Partition groups span AZs, reducing bandwidth.

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

Launch all EC2 instances in the same Availability Zone within the same AWS Region. Specify a placement group with cluster strategy when launching EC2 instances

upvoted 2 times

  **NoinNothing** 1 year, 6 months ago

**Selected Answer: A**

Cluster - have low latency if its in same AZ and same region so Answer is "A"

upvoted 3 times

  **BeeKayEnn** 1 year, 7 months ago

Answer would be A - As part of selecting all the EC2 instances in the same availability zone, they all will be within the same DC and logically the latency will be very less as compared to the other Availability Zones..

As all the autoscaling nodes will also be on the same availability zones, (as per Placement groups with Cluster mode), this would provide the low-latency network performance

Reference is below:

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

upvoted 4 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: A**

A - Low latency, high net throughput



upvoted 2 times

  **elearningtakai** 1 year, 7 months ago

**Selected Answer: A**

A placement group is a logical grouping of instances within a single Availability Zone, and it provides low-latency network connectivity between instances. By launching all EC2 instances in the same Availability Zone and specifying a placement group with cluster strategy, the application can take advantage of the high network throughput and low latency network connectivity that placement groups provide.

upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

**Selected Answer: A**

Cluster placement groups improves throughput between the instances which means less EC2 instances would be needed thus reducing costs.

upvoted 2 times

  **maciekmaciek** 1 year, 8 months ago

**Selected Answer: A**

A because Specify a placement group

upvoted 2 times

  **KZM** 1 year, 8 months ago

It is option A:



To achieve low latency, high throughput, and cost-effectiveness, the optimal solution is to launch EC2 instances as a placement group with the cluster strategy within the same Availability Zone.

upvoted 3 times

  **ManOnTheMoon** 1 year, 8 months ago



Why not C?

upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

You're thinking operational efficiency. The question asks for cost reduction.

upvoted 4 times

  **rrharris** 1 year, 8 months ago

Answer is A - Clustering

upvoted 3 times

  **Neha999** 1 year, 8 months ago

A : Cluster placement group

upvoted 5 times

A company that primarily runs its application servers on premises has decided to migrate to AWS. The company wants to minimize its need to scale its Internet Small Computer Systems Interface (iSCSI) storage on premises. The company wants only its recently accessed data to remain stored locally.

Which AWS solution should the company use to meet these requirements?

- A. Amazon S3 File Gateway
- B. AWS Storage Gateway Tape Gateway
- C. AWS Storage Gateway Volume Gateway stored volumes
- D. AWS Storage Gateway Volume Gateway cached volumes

Correct Answer: D

Community vote distribution

D (100%)

  **LuckyAro** Highly Voted 2 years, 2 months ago

Selected Answer: D

AWS Storage Gateway Volume Gateway provides two configurations for connecting to iSCSI storage, namely, stored volumes and cached volumes. The stored volume configuration stores the entire data set on-premises and asynchronously backs up the data to AWS. The cached volume configuration stores recently accessed data on-premises, and the remaining data is stored in Amazon S3.

Since the company wants only its recently accessed data to remain stored locally, the cached volume configuration would be the most appropriate. It allows the company to keep frequently accessed data on-premises and reduce the need for scaling its iSCSI storage while still providing access to all data through the AWS cloud. This configuration also provides low-latency access to frequently accessed data and cost-effective off-site backups for less frequently accessed data.

upvoted 41 times

  **tonybuivannghia** 7 months ago

You're correct but a small mistake that the stored volume configuration stores the entire data set on-premises and asynchronously backs up the data to AWS (S3). The cached volume is low latency access to most recent data.

upvoted 2 times

  **smgsi** Highly Voted 2 years, 2 months ago

Selected Answer: D

[https://docs.amazonaws.cn/en\\_us/storagegateway/latest/vgw/StorageGatewayConcepts.html#storage-gateway-cached-concepts](https://docs.amazonaws.cn/en_us/storagegateway/latest/vgw/StorageGatewayConcepts.html#storage-gateway-cached-concepts)

upvoted 10 times

  **TariqKipkemei** Most Recent 1 year, 7 months ago

Selected Answer: D

Frequently accessed data = AWS Storage Gateway Volume Gateway cached volumes

upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

Selected Answer: D

The best AWS solution to meet the requirements is to use AWS Storage Gateway cached volumes (option D).

The key points:

Company migrating on-prem app servers to AWS

Want to minimize scaling on-prem iSCSI storage

Only recent data should remain on-premises

The AWS Storage Gateway cached volumes allow the company to connect their on-premises iSCSI storage to AWS cloud storage. It stores frequently accessed data locally in the cache for low-latency access, while older data is stored in AWS.

upvoted 4 times

  **kruasan** 2 years ago

Selected Answer: D

- Volume Gateway cached volumes store entire datasets on S3, while keeping a portion of recently accessed data on your local storage as a cache. This meets the goal of minimizing on-premises storage needs while keeping hot data local.
- The cache provides low-latency access to your frequently accessed data, while long-term retention of the entire dataset is provided durable and cost-effective in S3.
- You get virtually unlimited storage on S3 for your infrequently accessed data, while controlling the amount of local storage used for cache. This simplifies on-premises storage scaling.

• Volume Gateway cached volumes support iSCSI connections from on-premises application servers, allowing a seamless migration experience. Servers access local cache and S3 storage volumes as iSCSI LUNs.

upvoted 8 times

  **kruasan** 2 years ago

In contrast:  
Option A - S3 File Gateway only provides file interfaces (NFS/SMB) to data in S3. It does not support block storage or cache recently accessed data locally.  
Option B - Tape Gateway is designed for long-term backup and archiving to virtual tape cartridges on S3. It does not provide primary storage volumes or local cache for low-latency access.  
Option C - Volume Gateway stored volumes keep entire datasets locally, then asynchronously back them up to S3. This does not meet the goal of minimizing on-premises storage needs.

upvoted 6 times

  **Steve\_4542636** 2 years, 2 months ago



**Selected Answer: D**

I vote D

upvoted 2 times

  **ManOnTheMoon** 2 years, 2 months ago

Agree with D  
upvoted 2 times

  **Babba** 2 years, 2 months ago



**Selected Answer: D**

recently accessed data to remain stored locally - cached  
upvoted 4 times

  **Bhawesh** 2 years, 2 months ago

**Selected Answer: D**

D. AWS Storage Gateway Volume Gateway cached volumes  
upvoted 4 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: D**

recently accessed data to remain stored locally - cached  
upvoted 4 times

A company has multiple AWS accounts that use consolidated billing. The company runs several active high performance Amazon RDS for Oracle On-Demand DB instances for 90 days. The company’s finance team has access to AWS Trusted Advisor in the consolidated billing account and all other AWS accounts.

The finance team needs to use the appropriate AWS account to access the Trusted Advisor check recommendations for RDS. The finance team must review the appropriate Trusted Advisor check to reduce RDS costs.

Which combination of steps should the finance team take to meet these requirements? (Choose two.)

- A. Use the Trusted Advisor recommendations from the account where the RDS instances are running.
- B. Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time.
- C. Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization.
- D. Review the Trusted Advisor check for Amazon RDS Idle DB Instances.
- E. Review the Trusted Advisor check for Amazon Redshift Reserved Node Optimization.

Correct Answer: *BD*

Community vote distribution

BD (47%)

BC (45%)

7%

- Nietzsche82

Highly Voted

2 years, 2 months ago

Selected Answer: BD

B & D

<https://aws.amazon.com/premiumsupport/knowledge-center/trusted-advisor-cost-optimization/>

upvoted 19 times
- Michal\_L\_95

Highly Voted

2 years, 2 months ago

Selected Answer: BC

I would go with B and C as the company is running for 90 days and C option is basing on 30 days report which would mean that there is higher potential on cost saving rather than on idle instances

upvoted 6 times
- Steve\_4542636

2 years, 2 months ago

C is stating "Reserved Instances" The question states they are using On Demand Instances. Reserved instances are reserved for less money for 1 or 3 years.

upvoted 8 times
- Lalo

1 year, 11 months ago

In the scenario it says for 90 days, therefore the correct answer is D

No C

upvoted 2 times
- pentium75

1 year, 4 months ago

D checks for instances that did not have a connection for a while, which is not the case here, all instances are active.

upvoted 2 times
- Michal\_L\_95

2 years, 1 month ago

Once read the question again, I agree with you.

upvoted 1 times
- pentium75

1 year, 4 months ago

C is about optimizing on-demand instances by turning them into reserved instances.

upvoted 2 times
- Yak\_Yeti

Most Recent

1 week, 1 day ago

Selected Answer: AC

The article emphasizes that Trusted Advisor generates recommendations at the individual AWS account level. While the consolidated billing account provides a centralized view of billing, the specific optimization recommendations for resources like RDS instances will be found within the account where those resources reside.

The article also mentions checks for idle resources like "Amazon EC2 Idle Instances" and "Amazon RDS Idle DB Instances." However, as the original analysis pointed out, the scenario specifies that the RDS instances are active and running a high-performance application, making the "Idle DB Instances" check less relevant for the immediate goal of reducing costs for these specific instances.

upvoted 1 times

🗨️ 👤 **a8a1e0e** 3 months ago

**Selected Answer: AC**

A: AWS Trusted Advisor recommendations for RDS cost optimization are only available within the AWS account where the instances are running. The consolidated billing (payer) account does not provide detailed service-level checks for linked accounts. Therefore, the finance team must access the Trusted Advisor check directly from the AWS account that owns the RDS instances to see relevant recommendations.

C: The company has been using high-performance RDS instances for 90 days, which indicates a long-term workload. Reserved Instances (RIs) offer significant cost savings compared to On-Demand pricing when committing to one or three years. Trusted Advisor identifies instances with consistent usage and recommends RIs to reduce costs, making this the most relevant check for optimizing expenses. Even after 90 days, purchasing an RI for the remaining duration still provides substantial savings.

upvoted 5 times

🗨️ 👤 **MatAlves** 7 months, 3 weeks ago

**Selected Answer: BC**

The answer is either BC or BD, depending on how you interpret "The company runs several active... instances for 90 days for 90 days."

D: it assumes the instances will only run for 90 days, so reserved instances can't be the answer, since it requires 1-3 years utilization.

C: it assumes there is no idle instances since they've been active for the last 90 days.

upvoted 4 times

🗨️ 👤 **aagarwallko** 7 months, 3 weeks ago

B: Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time. This option allows the finance team to see all RDS instance checks across all AWS accounts in one place. Since the company uses consolidated billing, this account will have access to all of the AWS accounts' Trusted Advisor recommendations.

C: Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization. This check can help identify cost savings opportunities for RDS by identifying instances that can be covered by Reserved Instances. This can result in significant savings on RDS costs.

upvoted 2 times

🗨️ 👤 **ike001** 10 months, 2 weeks ago

BD is the answer-Amazon Redshift Reserved Node Optimization and Relational Database Service (RDS) Reserved Instance Optimization check are not available to accounts linked in consolidated billing.

upvoted 3 times

🗨️ 👤 **sandordini** 1 year ago

**Selected Answer: BD**

"you can reserve a DB instance for a one- or three-year term". We only have data for 90 days. I feel it too risky to commit for 1/3 year(s) without information on future usage. If we knew that we expected the same usage pattern for the next 1,2,3 years, I'd agree with C.

upvoted 3 times

🗨️ 👤 **soufiyane** 1 year ago

**Selected Answer: BC**

B) Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time. This option allows the finance team to see all RDS instance checks across all AWS accounts in one place. Since the company uses consolidated billing, this account will have access to all of the AWS accounts' Trusted Advisor recommendations.

C) Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization. This check can help identify cost savings opportunities for RDS by identifying instances that can be covered by Reserved Instances. This can result in significant savings on RDS costs.

upvoted 2 times

🗨️ 👤 **Rhydian25** 10 months, 3 weeks ago

Reserved Instances are for 1 or 3 years. Not for 90 days

upvoted 5 times

🗨️ 👤 **scar0909** 1 year, 1 month ago

**Selected Answer: BC**

<https://aws.amazon.com/premiumsupport/knowledge-center/trusted-advisor-cost-optimization/>

upvoted 2 times

🗨️ 👤 **bujuman** 1 year, 1 month ago

**Selected Answer: BD**

Insights: The company runs several active high performance Amazon RDS for Oracle On-Demand DB instances for 90 days

So it's clear that this company need to check the configuration of any Amazon Relational Database Service (Amazon RDS) for any database (DB) instances that appear to be idle.

upvoted 2 times

🗨️ 👤 **dkw2342** 1 year, 2 months ago



B&C



AWS Trusted Advisor is an online resource to help you reduce cost, increase performance, and improve security by optimizing your AWS environment. (...) Recommendations are based on the previous calendar month's hour-by-hour usage aggregated across all consolidated billing accounts.



<https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/aws-trusted-advisor.html>



Amazon EC2 Reserved Instance Optimization: An important part of using AWS involves balancing your Reserved Instance (RI) purchase against your



On-Demand Instance usage. This check provides recommendations on which RIs will help reduce the costs incurred from using On-Demand Instances. We create these recommendations by analyzing your On-Demand usage for the past 30 days. We then categorizing the usage into eligible categories for reservations.  
<https://docs.aws.amazon.com/awssupport/latest/user/cost-optimization-checks.html#amazon-ec2-reserved-instances-optimization>  
upvoted 2 times



  **Salilgen** 4 months, 3 weeks ago  
The company run instances only for 90 days then reserved instances cannot save costs  
upvoted 1 times



  **NayeraB** 1 year, 2 months ago  
**Selected Answer: BC**  
If you're choosing D for the idle instances, Amazon RDS Reserved Instance Optimization Trusted Advisor check includes recommendations related to underutilized and idle RDS instances. It helps identify instances that are not fully utilized and provides recommendations on how to optimize costs, such as resizing or terminating unused instances, or purchasing reserved instances to match usage patterns more efficiently.  
upvoted 3 times

  **leejwli** 1 year, 3 months ago  
**Selected Answer: BC**  
Reserved Instances can be shared across accounts, and that is the reason why we need to check the consolidated bill.  
upvoted 3 times

  **farnamjam** 1 year, 4 months ago  
**Selected Answer: BC**  
BC  
we don't want to check Idle instances because the instances were active for last 90 days.  
Idle means it was inactive for at least 7 days.  
upvoted 5 times

  **farnamjam** 1 year, 4 months ago  
**Selected Answer: BD**  
BD  
we don't want to check Idle instances because the instances were active for last 90 days.  
Idle means it was inactive for at least 7 days.  
upvoted 2 times

  **MatAlves** 7 months, 3 weeks ago  
Then why you voted for D?  
upvoted 2 times

  **pentium75** 1 year, 4 months ago  
**Selected Answer: BC**  
Reserved Instance Optimization "checks your usage of RDS and provides recommendations on purchase of Reserved Instances to help reduce costs incurred from using RDS On-Demand." In other words, it is not about optimizing reserved instances (as many here think), it about optimizing on-demand instances by converting them to reserved ones.  
  
"Idle DB Instances" check is about databases that have "not had a connection for a prolonged period of time", which we know is not the case here.  
upvoted 4 times

A solutions architect needs to optimize storage costs. The solutions architect must identify any Amazon S3 buckets that are no longer being accessed or are rarely accessed.



Which solution will accomplish this goal with the LEAST operational overhead?


- A. Analyze bucket access patterns by using the S3 Storage Lens dashboard for advanced activity metrics.
- B. Analyze bucket access patterns by using the S3 dashboard in the AWS Management Console.
- C. Turn on the Amazon CloudWatch BucketSizeBytes metric for buckets. Analyze bucket access patterns by using the metrics data with Amazon Athena.
- D. Turn on AWS CloudTrail for S3 object monitoring. Analyze bucket access patterns by using CloudTrail logs that are integrated with Amazon CloudWatch Logs.

**Correct Answer: A**

*Community vote distribution*

A (94%)6%



-   **kpato87**


Highly Voted 

 2 years, 2 months ago

Selected Answer: A

S3 Storage Lens is a fully managed S3 storage analytics solution that provides a comprehensive view of object storage usage, activity trends, and recommendations to optimize costs. Storage Lens allows you to analyze object access patterns across all of your S3 buckets and generate detailed metrics and reports.



upvoted 23 times
-   **bdp123**


Highly Voted 

 2 years, 2 months ago

Selected Answer: A

<https://aws.amazon.com/blogs/aws/s3-storage-lens/>



upvoted 5 times
-   **Gape4**

Most Recent 



 10 months, 2 weeks ago

Selected Answer: A

S3 Storage Lens includes an interactive dashboard which you can find in the S3 console. The dashboard gives you the ability to perform filtering and drill-down into your metrics to really understand how your storage is being used. The metrics are organized into categories like data protection and cost efficiency, to allow you to easily find relevant metrics.

upvoted 2 times
-   **AmijoSando** 11 months, 4 weeks ago



Anyone passed the exam can confirm the right answer ? A or D

upvoted 1 times
-   **xyGGXH** 1 year, 2 months ago



Selected Answer: A

A



S3 Storage Lens is the first cloud storage analytics solution to provide a single view of object storage usage and activity across hundreds, or even thousands, of accounts in an organization, with drill-downs to generate insights at multiple aggregation levels.

upvoted 3 times
-   **Neung983** 1 year, 2 months ago

On the other hand, Option B suggests using the S3 dashboard in the AWS Management Console, which provides a straightforward and user-friendly interface to monitor S3 bucket access patterns. This option may have less operational overhead compared to setting up and managing Storage Lens. Additionally, for simply identifying rarely accessed buckets, the built-in metrics and access analysis provided by the S3 dashboard can often suffice without the need for advanced analytics offered by Storage Lens. Therefore, Option B is considered to have less operational overhead for the specific task described in the question.



upvoted 1 times
-   **jaswantn** 1 year, 2 months ago

But nowhere on S3 Storage Lens dashboard this information is available; that when the bucket is accessed last time. But it gives insight on the bucket's size. with this information we can check if files can be moved to less costly storage class. This way we can reduce storage cost..... The information which is the main requirement of the given scenario, is available when we use Cloudtrail logs ... so i choose option D.

upvoted 1 times
-   **jaswantn** 1 year, 2 months ago

if the bucket is being accessed frequently then we can leave it as it is, otherwise we can move the files to infrequent access storage class thus can save some money.

upvoted 2 times

  **Ruffyit** 1 year, 5 months ago

S3 Storage Lens is a fully managed S3 storage analytics solution that provides a comprehensive view of object storage usage, activity trends, and recommendations to optimize costs. Storage Lens allows you to analyze object access patterns across all of your S3 buckets and generate detailed metrics and reports.

upvoted 2 times

  **TariqKipkemei** 1 year, 7 months ago

**Selected Answer: A**

Amazon S3 Storage Lens was designed to handle this requirement.

upvoted 2 times

  **Wayne23Fang** 1 year, 7 months ago

**Selected Answer: D**

A missed turning on monitoring. It can also help you learn about your customer base and understand your Amazon S3 bill. By default, Amazon S3 doesn't collect server access logs. When you enable logging, Amazon S3 delivers access logs for a source bucket to a target bucket that you choose.

I could not find that S3 storage Lens examples online showing using Lens to identify idle S3 buckets. Instead I find using S3 Access Logging. Hmm.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

How will you find when a bucket was used the last time if you turn on logging NOW?

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

S3 Storage Lens is a cloud-storage analytics feature that provides you with 29+ usage and activity metrics, including object count, size, age, and access patterns. This data can help you understand how your data is being used and identify areas where you can optimize your storage costs. The S3 Storage Lens dashboard provides an interactive view of your storage usage and activity trends. This makes it easy to identify buckets that are no longer being accessed or are rarely accessed.



The S3 Storage Lens dashboard is a fully managed service, so there is no need to set up or manage any additional infrastructure.

upvoted 3 times

  **BigHammer** 1 year, 8 months ago

"S3 Storage Lens" seems to be the popular answer, however, where in Storage Lens can you see if a bucket/object is being USED? I see all kinds of stats, but not that.

upvoted 3 times

  **MatAlves** 7 months, 3 weeks ago

"S3 Storage Lens delivers organization-wide visibility into object storage usage, activity trends, and makes actionable recommendations to optimize costs..."

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

<https://aws.amazon.com/blogs/aws/s3-storage-lens/>

upvoted 3 times

  **kruasan** 2 years ago

**Selected Answer: A**

The S3 Storage Lens dashboard provides visibility into storage metrics and activity patterns to help optimize storage costs. It shows metrics like objects added, objects deleted, storage consumed, and requests. It can filter by bucket, prefix, and tag to analyze specific subsets of data

upvoted 4 times

  **kruasan** 2 years ago

B) The standard S3 console dashboard provides basic info but would require manually analyzing metrics for each bucket. This does not scale well and requires significant overhead.

C) Turning on the BucketSizeBytes metric and analyzing the data in Athena may provide insights but would require enabling metrics, building Athena queries, and analyzing the results. This requires more operational effort than option A.

D) Enabling CloudTrail logging and monitoring the logs in CloudWatch Logs could provide access pattern data but would require setting up CloudTrail, monitoring the logs, and analyzing the relevant info. This option has the highest operational overhead



upvoted 5 times

  **LuckyAro** 2 years, 2 months ago

**Selected Answer: A**

S3 Storage Lens provides a dashboard with advanced activity metrics that enable the identification of infrequently accessed and unused buckets. This can help a solutions architect optimize storage costs without incurring additional operational overhead.

upvoted 4 times

  **Babba** 2 years, 2 months ago

**Selected Answer: A**

it looks like it's A

upvoted 3 times



A company sells datasets to customers who do research in artificial intelligence and machine learning (AI/ML). The datasets are large, formatted files that are stored in an Amazon S3 bucket in the us-east-1 Region. The company hosts a web application that the customers use to purchase access to a given dataset. The web application is deployed on multiple Amazon EC2 instances behind an Application Load Balancer. After a purchase is made, customers receive an S3 signed URL that allows access to the files.

The customers are distributed across North America and Europe. The company wants to reduce the cost that is associated with data transfers and wants to maintain or improve performance.

What should a solutions architect do to meet these requirements?

- A. Configure S3 Transfer Acceleration on the existing S3 bucket. Direct customer requests to the S3 Transfer Acceleration endpoint. Continue to use S3 signed URLs for access control.
- B. Deploy an Amazon CloudFront distribution with the existing S3 bucket as the origin. Direct customer requests to the CloudFront URL. Switch to CloudFront signed URLs for access control.
- C. Set up a second S3 bucket in the eu-central-1 Region with S3 Cross-Region Replication between the buckets. Direct customer requests to the closest Region. Continue to use S3 signed URLs for access control.
- D. Modify the web application to enable streaming of the datasets to end users. Configure the web application to read the data from the existing S3 bucket. Implement access control directly in the application.

**Correct Answer: B**

*Community vote distribution*

B (100%)

 **LuckyAro** Highly Voted 1 year, 8 months ago

**Selected Answer: B**

To reduce the cost associated with data transfers and maintain or improve performance, a solutions architect should use Amazon CloudFront, a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds.

Deploying a CloudFront distribution with the existing S3 bucket as the origin will allow the company to serve the data to customers from edge locations that are closer to them, reducing data transfer costs and improving performance.

Directing customer requests to the CloudFront URL and switching to CloudFront signed URLs for access control will enable customers to access the data securely and efficiently.

upvoted 14 times

 **awsgeek75** Highly Voted 9 months, 2 weeks ago

**Selected Answer: B**

A: Speeds uploads

C: Increases the cost rather than reducing it

D: Stopped reading after "Modify the web application..."

upvoted 10 times

 **Ruffyt** Most Recent 11 months, 3 weeks ago

To reduce the cost associated with data transfers and maintain or improve performance, a solutions architect should use Amazon CloudFront, a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds.

upvoted 2 times

 **TariqKipkemei** 1 year ago

**Selected Answer: B**

Technically both option B and C will work. But because cost is a factor then Amazon CloudFront should be the preferred option.

upvoted 2 times

 **react97** 1 year, 1 month ago

**Selected Answer: B**

B.



1. Amazon CloudFront caches content at edge locations -- reducing the need for frequent data transfer from S3 bucket -- thus significantly lowering data transfer costs (as compared to directly serving data from S3 bucket to customers in different regions)

2. CloudFront delivers content to users from the nearest edge location -- minimizing latency -- improves performance for customers

A - focus on accelerating uploads to S3 which may not necessarily improve the performance needed for serving datasets to customers

C - helps with redundancy and data availability but does not necessarily offer cost savings for data transfer.  
D - complex to implement, does not address data transfer cost

upvoted 6 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/PrivateContent.html>

upvoted 4 times

  **Bhawesh** 1 year, 8 months ago

**Selected Answer: B**

B. Deploy an Amazon CloudFront distribution with the existing S3 bucket as the origin. Direct customer requests to the CloudFront URL. Switch to CloudFront signed URLs for access control.

<https://www.examttopics.com/discussions/amazon/view/68990-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 3 times

A company is using AWS to design a web application that will process insurance quotes. Users will request quotes from the application. Quotes must be separated by quote type, must be responded to within 24 hours, and must not get lost. The solution must maximize operational efficiency and must minimize maintenance.

Which solution meets these requirements?

- A. Create multiple Amazon Kinesis data streams based on the quote type. Configure the web application to send messages to the proper data stream. Configure each backend group of application servers to use the Kinesis Client Library (KCL) to pool messages from its own data stream.
- B. Create an AWS Lambda function and an Amazon Simple Notification Service (Amazon SNS) topic for each quote type. Subscribe the Lambda function to its associated SNS topic. Configure the application to publish requests for quotes to the appropriate SNS topic.
- C. Create a single Amazon Simple Notification Service (Amazon SNS) topic. Subscribe Amazon Simple Queue Service (Amazon SQS) queues to the SNS topic. Configure SNS message filtering to publish messages to the proper SQS queue based on the quote type. Configure each backend application server to use its own SQS queue.
- D. Create multiple Amazon Kinesis Data Firehose delivery streams based on the quote type to deliver data streams to an Amazon OpenSearch Service cluster. Configure the application to send messages to the proper delivery stream. Configure each backend group of application servers to search for the messages from OpenSearch Service and process them accordingly.

Correct Answer: C

Community vote distribution

C (100%)

  **LuckyAro** Highly Voted 2 years, 2 months ago

Selected Answer: C

Quote types need to be separated: SNS message filtering can be used to publish messages to the appropriate SQS queue based on the quote type, ensuring that quotes are separated by type.

Quotes must be responded to within 24 hours and must not get lost: SQS provides reliable and scalable queuing for messages, ensuring that quotes will not get lost and can be processed in a timely manner. Additionally, each backend application server can use its own SQS queue, ensuring that quotes are processed efficiently without any delay.

Operational efficiency and minimizing maintenance: Using a single SNS topic and multiple SQS queues is a scalable and cost-effective approach, which can help to maximize operational efficiency and minimize maintenance. Additionally, SNS and SQS are fully managed services, which means that the company will not need to worry about maintenance tasks such as software updates, hardware upgrades, or scaling the infrastructure.

upvoted 19 times

  **Vlad** Highly Voted 2 years, 2 months ago

C is the best option

upvoted 9 times

  **akshay243007** Most Recent 9 months, 1 week ago

Selected Answer: C

SQS + SNS = fanout

upvoted 3 times

  **Uzbekistan** 1 year, 2 months ago

Option C would be the most suitable solution to meet the requirements while maximizing operational efficiency and minimizing maintenance.

Explanation:

Amazon SNS (Simple Notification Service) allows for the creation of a single topic to which multiple subscribers can be attached. In this scenario, each quote type can be considered a subscriber. Amazon SQS (Simple Queue Service) queues can be subscribed to the SNS topic, and SNS message filtering can be used to direct messages to the appropriate SQS queue based on the quote type. This setup ensures that quotes are separated by quote type and that they are not lost. Each backend application server can then poll its own SQS queue to retrieve and process messages. This architecture is efficient, scalable, and requires minimal maintenance, as it leverages managed AWS services without the need for complex custom code or infrastructure setup.

upvoted 4 times

  **awsgeek75** 1 year, 4 months ago

Selected Answer: C

I originally went for D due to searching requirements but Open Search is for analytics and logs and nothing to do with data coming from streams as in this question.

upvoted 2 times

🗨️ 👤 **Ruffyit** 1 year, 5 months ago

Quote types need to be separated: SNS message filtering can be used to publish messages to the appropriate SQS queue based on the quote type, ensuring that quotes are separated by type.

Quotes must be responded to within 24 hours and must not get lost: SQS provides reliable and scalable queuing for messages, ensuring that quotes will not get lost and can be processed in a timely manner. Additionally, each backend application server can use its own SQS queue, ensuring that quotes are processed efficiently without any delay.

Operational efficiency and minimizing maintenance: Using a single SNS topic and multiple SQS queues is a scalable and cost-effective approach, which can help to maximize operational efficiency and minimize maintenance. Additionally, SNS and SQS are fully managed services, which means that the company will not need to worry about maintenance tasks such as software updates, hardware upgrades, or scaling the

upvoted 3 times

🗨️ 👤 **tekjm** 1 year, 6 months ago

Keyword is "...and must not get lost" = SQS

upvoted 3 times

🗨️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

Create a single SNS topic

Subscribe separate SQS queues per quote type

Use SNS message filtering to send messages to proper queue

Backend servers poll their respective SQS queue

The key points:

Quote requests must be processed within 24 hrs without loss

Need to maximize efficiency and minimize maintenance

Requests separated by quote type

upvoted 2 times

🗨️ 👤 **lexotan** 2 years ago

**Selected Answer: C**

This wrong answers from examtopic are getting me so frustrated. Which one is the correct answer then?

upvoted 6 times

🗨️ 👤 **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: C**

This is the SNS fan-out technique where you will have one SNS service to many SQS services

<https://docs.aws.amazon.com/sns/latest/dg/sns-sqs-as-subscriber.html>

upvoted 7 times

🗨️ 👤 **UnluckyDucky** 2 years, 1 month ago

SNS Fan-out fans message to all subscribers, this uses SNS filtering to publish the message only to the right SQS queue (not all of them).

upvoted 3 times

🗨️ 👤 **Yechi** 2 years, 2 months ago

**Selected Answer: C**

<https://aws.amazon.com/getting-started/hands-on/filter-messages-published-to-topics/>

upvoted 8 times

A company has an application that runs on several Amazon EC2 instances. Each EC2 instance has multiple Amazon Elastic Block Store (Amazon EBS) data volumes attached to it. The application’s EC2 instance configuration and data need to be backed up nightly. The application also needs to be recoverable in a different AWS Region.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Write an AWS Lambda function that schedules nightly snapshots of the application’s EBS volumes and copies the snapshots to a different Region.
- B. Create a backup plan by using AWS Backup to perform nightly backups. Copy the backups to another Region. Add the application’s EC2 instances as resources.
- C. Create a backup plan by using AWS Backup to perform nightly backups. Copy the backups to another Region. Add the application’s EBS volumes as resources.
- D. Write an AWS Lambda function that schedules nightly snapshots of the application's EBS volumes and copies the snapshots to a different Availability Zone.

**Correct Answer: B**

*Community vote distribution*

B (94%)

6%

- khasport**

Highly Voted

1 year, 8 months ago

B is answer so the requirement is "The application’s EC2 instance configuration and data need to be backed up nightly" so we need "add the application’s EC2 instances as resources". This option will backup both EC2 configuration and data

upvoted 21 times
- TungPham**

Highly Voted

1 year, 8 months ago

**Selected Answer: B**

<https://aws.amazon.com/vi/blogs/aws/aws-backup-ec2-instances-efs-single-file-restore-and-cross-region-backup/>  
When you back up an EC2 instance, AWS Backup will protect all EBS volumes attached to the instance, and it will attach them to an AMI that stores all parameters from the original EC2 instance except for two

upvoted 16 times
- raymondfekry**

Most Recent

10 months, 2 weeks ago

**Selected Answer: B**

Question says: " The application’s EC2 instance configuration and data need to be backed up", thus C is not correct, B is

upvoted 3 times
- Ruffyit**

11 months, 3 weeks ago

<https://aws.amazon.com/vi/blogs/aws/aws-backup-ec2-instances-efs-single-file-restore-and-cross-region-backup/>  
When you back up an EC2 instance, AWS Backup will protect all EBS volumes attached to the instance, and it will attach them to an AMI that stores all parameters from the original EC2 instance except for two

upvoted 2 times
- TariqKipkemei**

1 year ago

**Selected Answer: B**

As part of configuring a backup plan you need to enable (opt-in) resource types that will be protected by the backup plan. For this case EC2.  
<https://aws.amazon.com/getting-started/hands-on/amazon-ec2-backup-and-restore-using-aws-backup/#:~:text=the%20services%20used%20with-,AWS%20Backup,-a.%20In%20the%20navigation>

upvoted 2 times
- Guru4Cloud**

1 year, 2 months ago

**Selected Answer: B**

B is the most appropriate solution because it allows you to create a backup plan to automate the backup process of EC2 instances and EBS volumes, and copy backups to another region. Additionally, you can add the application's EC2 instances as resources to ensure their configuration and data are backed up nightly.

upvoted 2 times
- Geekboii**

1 year, 7 months ago

i would say B

upvoted 2 times
- Geekboii**

1 year, 7 months ago

i would say B

upvoted 2 times

  **AlmeroSenior** 1 year, 8 months ago

**Selected Answer: B**

AWS KB states if you select the EC2 instance , associated EBS's will be auto covered .

<https://aws.amazon.com/blogs/aws/aws-backup-ec2-instances-efs-single-file-restore-and-cross-region-backup/>

upvoted 3 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: B**

B is the most appropriate solution because it allows you to create a backup plan to automate the backup process of EC2 instances and EBS volumes, and copy backups to another region. Additionally, you can add the application's EC2 instances as resources to ensure their configuration and data are backed up nightly.

A and D involve writing custom Lambda functions to automate the snapshot process, which can be complex and require more maintenance effort. Moreover, these options do not provide an integrated solution for managing backups and recovery, and copying snapshots to another region.



Option C involves creating a backup plan with AWS Backup to perform backups for EBS volumes only. This approach would not back up the EC2 instances and their configuration

upvoted 3 times

  **Mia2009687** 1 year, 3 months ago

The data is stored in the EBS storage volume, EC2 won't hold the data, I think we need "Add the application's EBS volumes as resources."

upvoted 2 times

  **everfly** 1 year, 8 months ago

**Selected Answer: C**

The application's EC2 instance configuration and data are stored on EBS volume right?

upvoted 2 times

  **pentium75** 10 months, 1 week ago

No, this is not how EC2 works.

upvoted 1 times

  **awsgeek75** 9 months, 2 weeks ago

No, ECS config is the config you provide when launching the EC2 instance. EBS is a resource for EC2 as a part of configuration. When you backup EC2, it will backup the instance which resulted from the configuration and that will include the EBS volumes that are attached to the instance.

upvoted 3 times

  **Rehan33** 1 year, 8 months ago

The data is store on EBS volume so why we are not using EBS as a source instead of EC2

upvoted 1 times

  **thewalker** 9 months, 3 weeks ago

Also, if EBS volumes are added or removed as the requirement, not need to update the AWS Config.

upvoted 1 times

  **obatunde** 1 year, 8 months ago

Because "The application's EC2 instance configuration and data need to be backed up nightly"


upvoted 6 times

  **fulingyu288** 1 year, 8 months ago

**Selected Answer: B**

Use AWS Backup to create a backup plan that includes the EC2 instances, Amazon EBS snapshots, and any other resources needed for recovery. The backup plan can be configured to run on a nightly schedule.

upvoted 2 times

  **zTopic** 1 year, 8 months ago

**Selected Answer: B**

The application's EC2 instance configuration and data need to be backed up nightly > > B

upvoted 2 times

  **NolaHOla** 1 year, 8 months ago

But isn't the data needed to be backed up on the EBS ?

upvoted 1 times

A company is building a mobile app on AWS. The company wants to expand its reach to millions of users. The company needs to build a platform so that authorized users can watch the company’s content on their mobile devices.



What should a solutions architect recommend to meet these requirements?


- A. Publish content to a public Amazon S3 bucket. Use AWS Key Management Service (AWS KMS) keys to stream content.
- B. Set up IPsec VPN between the mobile app and the AWS environment to stream content.
- C. Use Amazon CloudFront. Provide signed URLs to stream content.
- D. Set up AWS Client VPN between the mobile app and the AWS environment to stream content.

**Correct Answer: C**

*Community vote distribution*

C (100%)



-  **Steve\_4542636**

Highly Voted 



 1 year, 8 months ago

Selected Answer: C



Enough with CloudFront already.


upvoted 29 times
-  **awsgeek75** 9 months, 2 weeks ago

This whole exam seems like a sales pitch for CloudFront and SQS... lol!

upvoted 7 times
-  **TariqKipkemei** 1 year, 6 months ago

Hahaha..cloudfront too hyped :)



upvoted 3 times
-  **LuckyAro**

Highly Voted 



 1 year, 8 months ago


Selected Answer: C

Amazon CloudFront is a content delivery network (CDN) that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. CloudFront supports signed URLs that provide authorized access to your content. This feature allows the company to control who can access their content and for how long, providing a secure and scalable solution for millions of users.

upvoted 9 times
-  **mwwt2022** 10 months, 1 week ago

great explanation!



upvoted 2 times
-  **lostmagnet001**

Most Recent 



 9 months ago

Selected Answer: C

CF always for reaching places



upvoted 3 times
-  **Ruffyit** 11 months, 3 weeks ago

Use Amazon CloudFront. Provide signed URLs to stream content.

upvoted 2 times
-  **Guru4Cloud** 1 year, 2 months ago



Selected Answer: C

Use Amazon CloudFront. Provide signed URLs to stream content.



upvoted 2 times
-  **antropaws** 1 year, 5 months ago

Selected Answer: C

C is correct.

upvoted 3 times
-  **kprakashbehera** 1 year, 7 months ago

Cloudfront is the correct solution.

upvoted 5 times
-  **datz** 1 year, 7 months ago

Feel your pain :D hahaha

upvoted 2 times



**jennyka76** 1 year, 8 months ago

C

<https://www.amazonaws.cn/en/cloudfront/>

upvoted 2 times

A company has an on-premises MySQL database used by the global sales team with infrequent access patterns. The sales team requires the database to have minimal downtime. A database administrator wants to migrate this database to AWS without selecting a particular instance type in anticipation of more users in the future.

Which service should a solutions architect recommend?

- A. Amazon Aurora MySQL
- B. Amazon Aurora Serverless for MySQL
- C. Amazon Redshift Spectrum
- D. Amazon RDS for MySQL

**Correct Answer:** B

Community vote distribution

B (100%)

- cloudbusting

Highly Voted

1 year, 8 months ago

"without selecting a particular instance type" = serverless

upvoted 30 times
- elearningtakai

Highly Voted

1 year, 7 months ago

Selected Answer: B

With Aurora Serverless for MySQL, you don't need to select a particular instance type, as the service automatically scales up or down based on the application's needs.

upvoted 11 times
- awsgeek75

Most Recent

9 months, 2 weeks ago

Selected Answer: B

The DBA had one job and he doesn't want to do it... so B it is

upvoted 9 times
- Ruffyit

11 months, 3 weeks ago

without selecting a particular instance type = Amazon Aurora Serverless for MySQL

upvoted 4 times
- TariqKipkemei

1 year ago

Selected Answer: B

without selecting a particular instance type = Amazon Aurora Serverless for MySQL

upvoted 2 times
- Guru4Cloud

1 year, 2 months ago

Selected Answer: B

B. Amazon Aurora Serverless for MySQL

upvoted 2 times
- Diqian

1 year, 2 months ago

What's the difference between A and B. I think Aurora is serverless, isn't it?

upvoted 1 times
- Valder21

1 year, 1 month ago

seems serverless is an option of amazon aurora. Not a very good naming scheme.

upvoted 2 times
- Srikanth0057

1 year, 7 months ago

Selected Answer: B

Bbbbbbb

upvoted 2 times
- Steve\_4542636

1 year, 8 months ago

Selected Answer: B

<https://aws.amazon.com/rds/aurora/serverless/>

upvoted 2 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: B**

Amazon Aurora Serverless for MySQL is a fully managed, auto-scaling relational database service that scales up or down automatically based on the application demand. This service provides all the capabilities of Amazon Aurora, such as high availability, durability, and security, without requiring the customer to provision any database instances.

With Amazon Aurora Serverless for MySQL, the sales team can enjoy minimal downtime since the database is designed to automatically scale to accommodate the increased traffic. Additionally, the service allows the customer to pay only for the capacity used, making it cost-effective for infrequent access patterns.

Amazon RDS for MySQL could also be an option, but it requires the customer to select an instance type, and the database administrator would need to monitor and adjust the instance size manually to accommodate the increasing traffic.

upvoted 3 times

  **Drayen25** 1 year, 8 months ago

Minimal downtime points directly to Aurora Serverless

upvoted 3 times

A company experienced a breach that affected several applications in its on-premises data center. The attacker took advantage of vulnerabilities in the custom applications that were running on the servers. The company is now migrating its applications to run on Amazon EC2 instances. The company wants to implement a solution that actively scans for vulnerabilities on the EC2 instances and sends a report that details the findings.



Which solution will meet these requirements?


- A. Deploy AWS Shield to scan the EC2 instances for vulnerabilities. Create an AWS Lambda function to log any findings to AWS CloudTrail.
- B. Deploy Amazon Macie and AWS Lambda functions to scan the EC2 instances for vulnerabilities. Log any findings to AWS CloudTrail.
- C. Turn on Amazon GuardDuty. Deploy the GuardDuty agents to the EC2 instances. Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings.
- D. Turn on Amazon Inspector. Deploy the Amazon Inspector agent to the EC2 instances. Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings.

**Correct Answer: D**

*Community vote distribution*



D (98%)



-   **siyam008**


Highly Voted 

 1 year, 8 months ago



Selected Answer: D


AWS Shield for DDOS  
Amazon Macie for discover and protect sensitive data  
Amazon GuardDuty for intelligent thread discovery to protect AWS account  
Amazon Inspector for automated security assessment. like known Vulnerability  
upvoted 58 times
-   **robotgeek** 3 months, 2 weeks ago

So 58 upvotes and you are talking about "known Vulnerability" for a "custom app"... ok  
upvoted 1 times
-   **benacert**

Highly Voted 



 10 months, 4 weeks ago



Whenever I feel vulnerable, I use AWS Inspector..  
upvoted 17 times
-   **zinabu**

Most Recent 



 6 months, 2 weeks ago

Selected Answer: D

Amazon Inspector for automated security assessment. like known Vulnerability  
upvoted 4 times
-   **Ruffyit** 11 months, 3 weeks ago

AWS Shield for DDOS  
Amazon Macie for discover and protect sensitive data  
Amazon GuardDuty for intelligent thread discovery to protect AWS account  
Amazon Inspector for automated security assessment. like known Vulnerability  
upvoted 5 times
-   **TariqKipkemei** 1 year ago

Selected Answer: D

vulnerabilities = Amazon Inspector  
malicious activity = Amazon GuardDuty  
upvoted 11 times
-   **Guru4Cloud** 1 year, 2 months ago

Selected Answer: D

Enable Amazon Inspector  
Deploy Inspector agents to EC2 instances  
Use Lambda to generate and distribute vulnerability reports  
The key points:  
  
Migrate on-prem apps with vulnerabilities to EC2  
Need active scanning of EC2 instances for vulnerabilities  
Require reports on findings  
upvoted 4 times

🗨️ 👤 **kruasan** 1 year, 6 months ago

**Selected Answer: D**

Amazon Inspector:

- Performs active vulnerability scans of EC2 instances. It looks for software vulnerabilities, unintended network accessibility, and other security issues.
- Requires installing an agent on EC2 instances to perform scans. The agent must be deployed to each instance.
- Provides scheduled scan reports detailing any findings of security risks or vulnerabilities. These reports can be used to patch or remediate issues.
- Is best suited for proactively detecting security weaknesses and misconfigurations in your AWS environment.

upvoted 4 times

🗨️ 👤 **kruasan** 1 year, 6 months ago

Amazon GuardDuty:

- Monitors for malicious activity like unusual API calls, unauthorized infrastructure deployments, or compromised EC2 instances. It uses machine learning and behavioral analysis of logs.
- Does not require installing any agents. It relies on analyzing AWS CloudTrail, VPC Flow Logs, and DNS logs.
- Alerts you to any detected threats, suspicious activity or policy violations in your AWS accounts. These alerts warrant investigation but may not always require remediation.
- Is focused on detecting active threats, unauthorized behavior, and signs of a compromise in your AWS environment.
- Can also detect some vulnerabilities and misconfigurations but coverage is not as broad as a dedicated service like Inspector.

upvoted 7 times

🗨️ 👤 **datz** 1 year, 7 months ago

**Selected Answer: D**

Amazon Inspector is a vulnerability scanning tool that you can use to identify potential security issues within your EC2 instances.

It is a kind of automated security assessment service that checks the network exposure of your EC2 or latest security state for applications running into your EC2 instance. It has ability to auto discover your AWS workload and continuously scan for the open loophole or vulnerability.

upvoted 2 times

🗨️ 👤 **shanwford** 1 year, 7 months ago

**Selected Answer: D**

Amazon Inspector is a vulnerability scanning tool that you can use to identify potential security issues within your EC2 instances. Guard Duty continuously monitors your entire AWS account via Cloud Trail, Flow Logs, DNS Logs as Input.

upvoted 2 times

🗨️ 👤 **GalileoEC2** 1 year, 7 months ago

**Selected Answer: C**

:) C is the correct

<https://cloudkatha.com/amazon-guardduty-vs-inspector-which-one-should-you-use/>

upvoted 1 times

🗨️ 👤 **jayantp04** 10 months, 2 weeks ago

Document itself saying that

Amazon Inspector is a vulnerability scanning tool

hence correct Answer is D

upvoted 2 times

🗨️ 👤 **MssP** 1 year, 7 months ago

Please, read the link you sent: Amazon Inspector is a vulnerability scanning tool that you can use to identify potential security issues within your EC2 instances. GuardDuty is very critical part to identify threats, based on that findings you can setup automated preventive actions or remediation's. So Answer is D.

upvoted 2 times

🗨️ 👤 **GalileoEC2** 1 year, 7 months ago

**Selected Answer: C**

<https://cloudkatha.com/amazon-guardduty-vs-inspector-which-one-should-you-use/>

upvoted 1 times

🗨️ 👤 **LuckyAro** 1 year, 8 months ago

**Selected Answer: D**

Amazon Inspector is a security assessment service that helps to identify security vulnerabilities and compliance issues in applications deployed on Amazon EC2 instances. It can be used to assess the security of applications that are deployed on Amazon EC2 instances, including those that are custom-built.

To use Amazon Inspector, the Amazon Inspector agent must be installed on the EC2 instances that need to be assessed. The agent collects data about the instances and sends it to Amazon Inspector for analysis. Amazon Inspector then generates a report that details any security vulnerabilities that were found and provides guidance on how to remediate them.

By configuring an AWS Lambda function, the company can automate the generation and distribution of reports that detail the findings. This means that reports can be generated and distributed as soon as vulnerabilities are detected, allowing the company to take action quickly.

upvoted 2 times

🗨️ 👤 **pbpally** 1 year, 8 months ago

**Selected Answer: D**

I'm a little confused on how someone came up with C, it is definitely D.  
upvoted 2 times

  **obatunde** 1 year, 8 months ago

**Selected Answer: D**

Amazon Inspector  
upvoted 3 times



  **obatunde** 1 year, 8 months ago

Amazon Inspector is an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities and unintended network exposure. <https://aws.amazon.com/inspector/features/?nc=sn&loc=2>  
upvoted 4 times

  **Palanda** 1 year, 8 months ago



**Selected Answer: D**

I think D  
upvoted 1 times

  **minglu** 1 year, 8 months ago

**Selected Answer: D**

Inspector for EC2  
upvoted 1 times

  **skiwili** 1 year, 8 months ago

**Selected Answer: D**

Ddddddd  
upvoted 1 times

A company uses an Amazon EC2 instance to run a script to poll for and process messages in an Amazon Simple Queue Service (Amazon SQS) queue. The company wants to reduce operational costs while maintaining its ability to process a growing number of messages that are added to the queue.

What should a solutions architect recommend to meet these requirements?

- A. Increase the size of the EC2 instance to process messages faster.
- B. Use Amazon EventBridge to turn off the EC2 instance when the instance is underutilized.
- C. Migrate the script on the EC2 instance to an AWS Lambda function with the appropriate runtime.
- D. Use AWS Systems Manager Run Command to run the script on demand.

**Correct Answer: C**

*Community vote distribution*

C (91%)9%

- kpato87**

Highly Voted

2 years, 2 months ago

Selected Answer: C

By migrating the script to AWS Lambda, the company can take advantage of the auto-scaling feature of the service. AWS Lambda will automatically scale resources to match the size of the workload. This means that the company will not have to worry about provisioning or managing instances as the number of messages increases, resulting in lower operational costs

upvoted 13 times
- Guru4Cloud**

Highly Voted

1 year, 8 months ago

Selected Answer: C

The key points are:

Currently using an EC2 instance to poll SQS and process messages  
Want to reduce costs while handling growing message volume  
By migrating the polling script to a Lambda function, the company can avoid the cost of running a dedicated EC2 instance. Lambda functions scale automatically to handle message spikes. And Lambda billing is based on actual usage, resulting in cost savings versus provisioned EC2 capacity.

upvoted 8 times
- TariqKipkemei**

Most Recent

1 year, 6 months ago

Selected Answer: C

reduce operational costs = serverless = Lambda functions

upvoted 2 times
- Steve\_4542636**

2 years, 2 months ago

Selected Answer: C

Lambda costs money only when it's processing, not when idle

upvoted 3 times
- ManOnTheMoon**

2 years, 2 months ago

Agree with C

upvoted 1 times
- khasport**

2 years, 2 months ago

the answer is C. With this option, you can reduce operational cost as the question mentioned

upvoted 1 times
- LuckyAro**

2 years, 2 months ago

Selected Answer: C

AWS Lambda is a serverless compute service that allows you to run your code without provisioning or managing servers. By migrating the script to an AWS Lambda function, you can eliminate the need to maintain an EC2 instance, reducing operational costs. Additionally, Lambda automatically scales to handle the increasing number of messages in the SQS queue.

upvoted 1 times
- zTopic**

2 years, 2 months ago

Selected Answer: C

It Should be C.  
Lambda allows you to execute code without provisioning or managing servers, so it is ideal for running scripts that poll for and process messages in an Amazon SQS queue. The scaling of the Lambda function is automatic, and you only pay for the actual time it takes to process the messages.

upvoted 3 times



  **Bhawesh** 2 years, 2 months ago

**Selected Answer: D**

To reduce the operational overhead, it should be:



D. Use AWS Systems Manager Run Command to run the script on demand.

upvoted 3 times

  **lucdt4** 1 year, 11 months ago


No, replace EC2 instead by using lambda to reduce costs

upvoted 1 times

  **pentium75** 1 year, 4 months ago

So every time an item is added to the queue, you log into AWS Systems Manager through your browser, select "Run Command" and select your instance and enter the command to run the script?

upvoted 4 times

  **ike001** 10 months, 2 weeks ago

very sarcastic question :)

upvoted 1 times

A company uses a legacy application to produce data in CSV format. The legacy application stores the output data in Amazon S3. The company is deploying a new commercial off-the-shelf (COTS) application that can perform complex SQL queries to analyze data that is stored in Amazon Redshift and Amazon S3 only. However, the COTS application cannot process the .csv files that the legacy application produces.

The company cannot update the legacy application to produce data in another format. The company needs to implement a solution so that the COTS application can use the data that the legacy application produces.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Glue extract, transform, and load (ETL) job that runs on a schedule. Configure the ETL job to process the .csv files and store the processed data in Amazon Redshift.
- B. Develop a Python script that runs on Amazon EC2 instances to convert the .csv files to .sql files. Invoke the Python script on a cron schedule to store the output files in Amazon S3.
- C. Create an AWS Lambda function and an Amazon DynamoDB table. Use an S3 event to invoke the Lambda function. Configure the Lambda function to perform an extract, transform, and load (ETL) job to process the .csv files and store the processed data in the DynamoDB table.
- D. Use Amazon EventBridge to launch an Amazon EMR cluster on a weekly schedule. Configure the EMR cluster to perform an extract, transform, and load (ETL) job to process the .csv files and store the processed data in an Amazon Redshift table.

Correct Answer: A

Community vote distribution

A (98%)

- awsgeek75

Highly Voted

9 months, 2 weeks ago

Selected Answer: A

Time to sell some Glue.

I believe these kind of questions are there to indoctrinate us into acknowledging how blessed we are to have managed services like AWS Glue when you look at other horrible and painful options

upvoted 19 times
- elearningtakai

Highly Voted

1 year, 7 months ago

Selected Answer: A

A, AWS Glue is a fully managed ETL service that can extract data from various sources, transform it into the required format, and load it into a target data store. In this case, the ETL job can be configured to read the CSV files from Amazon S3, transform the data into a format that can be loaded into Amazon Redshift, and load it into an Amazon Redshift table.

B requires the development of a custom script to convert the CSV files to SQL files, which could be time-consuming and introduce additional operational overhead. C, while using serverless technology, requires the additional use of DynamoDB to store the processed data, which may not be necessary if the data is only needed in Amazon Redshift. D, while an option, is not the most efficient solution as it requires the creation of an EMR cluster, which can be costly and complex to manage.

upvoted 8 times
- pentium75

Most Recent

10 months, 1 week ago

Selected Answer: A

B - Developing a script is surely not minimizing operational effort

C - Stores data in DynamoDB where the new app cannot use it

D - Could work but is total overkill (EMR is for Big Data analysis, not for simple ETL)

upvoted 4 times
- Ruffyit

11 months, 3 weeks ago

A-ETL is serverless & best suited with the requirement who primary job is ETL

B-Usage of Ec2 adds operational overhead & incur costs

C-DynamoDB(NoSql) does suit the requirement as company is performing SQL queries

D-EMR adds operational overhead & incur costs

upvoted 3 times
- ACloud\_Guru15

12 months ago

Selected Answer: A

A-ETL is serverless & best suited with the requirement who primary job is ETL

B-Usage of Ec2 adds operational overhead & incur costs

C-DynamoDB(NoSql) does suit the requirement as company is performing SQL queries

D-EMR adds operational overhead & incur costs

upvoted 2 times

🗄️ 👤 **TariqKipkemei** 1 year ago

**Selected Answer: A**

Data transformation = AWS Glue  
upvoted 2 times

🗄️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

Create an AWS Glue ETL job to process the CSV files  
Configure the job to run on a schedule  
Output the transformed data to Amazon Redshift  
The key points:

Legacy app generates CSV files in S3  
New app requires data in Redshift or S3  
Need to transform CSV to support new app with minimal ops overhead  
upvoted 2 times

🗄️ 👤 **kraken21** 1 year, 7 months ago

**Selected Answer: A**

Glue is server less and has less operational head than EMR so A.  
upvoted 2 times

🗄️ 👤 **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

o meet the requirement with the least operational overhead, a serverless approach should be used. Among the options provided, option C provides a serverless solution using AWS Lambda, S3, and DynamoDB. Therefore, the solution should be to create an AWS Lambda function and an Amazon DynamoDB table. Use an S3 event to invoke the Lambda function. Configure the Lambda function to perform an extract, transform, and load (ETL) job to process the .csv files and store the processed data in the DynamoDB table.  
Option A is also a valid solution, but it may involve more operational overhead than Option C. With Option A, you would need to set up and manage an AWS Glue job, which would require more setup time than creating an AWS Lambda function. Additionally, AWS Glue jobs have a minimum execution time of 10 minutes, which may not be necessary or desirable for this use case. However, if the data processing is particularly complex or requires a lot of data transformation, AWS Glue may be a more appropriate solution.  
upvoted 1 times

🗄️ 👤 **pentium75** 10 months, 1 week ago

Creating and maintaining a Lambda function is more "operational overhead" than using a ready-made service such as Glue. But more important, answer C says "store the processed data in the DynamoDB table" while the application can "analyze data that is stored in Amazon Redshift and Amazon S3 only".  
upvoted 2 times

🗄️ 👤 **MssP** 1 year, 7 months ago

Important point: The COTS performs complex SQL queries to analyze data in Amazon Redshift. If you use DynamoDB -> No SQL querires.  
Option A makes more sense.  
upvoted 4 times

🗄️ 👤 **LuckyAro** 1 year, 8 months ago

**Selected Answer: A**

A would be the best solution as it involves the least operational overhead. With this solution, an AWS Glue ETL job is created to process the .csv files and store the processed data directly in Amazon Redshift. This is a serverless approach that does not require any infrastructure to be provisioned, configured, or maintained. AWS Glue provides a fully managed, pay-as-you-go ETL service that can be easily configured to process data from S3 and load it into Amazon Redshift. This approach allows the legacy application to continue to produce data in the CSV format that it currently uses, while providing the new COTS application with the ability to analyze the data using complex SQL queries.  
upvoted 4 times

🗄️ 👤 **jennyka76** 1 year, 8 months ago

A  
<https://docs.aws.amazon.com/glue/latest/dg/aws-glue-programming-etl-format-csv-home.html>  
I AGREE AFTER READING LINK  
upvoted 2 times

🗄️ 👤 **cloudbusting** 1 year, 8 months ago

A: <https://docs.aws.amazon.com/glue/latest/dg/aws-glue-programming-etl-format.html>  
upvoted 2 times

A company recently migrated its entire IT environment to the AWS Cloud. The company discovers that users are provisioning oversized Amazon EC2 instances and modifying security group rules without using the appropriate change control process. A solutions architect must devise a strategy to track and audit these inventory and configuration changes.

Which actions should the solutions architect take to meet these requirements? (Choose two.)



- A. Enable AWS CloudTrail and use it for auditing.
- B. Use data lifecycle policies for the Amazon EC2 instances.
- C. Enable AWS Trusted Advisor and reference the security dashboard.
- D. Enable AWS Config and create rules for auditing and compliance purposes.
- E. Restore previous resource configurations with an AWS CloudFormation template.

**Correct Answer:** AD

*Community vote distribution*

AD (91%)

9%



-   **LuckyAro** Highly Voted 1 year, 8 months ago

**Selected Answer: AD**

A. Enable AWS CloudTrail and use it for auditing. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS Command Line Interface (CLI), and AWS SDKs and APIs. By enabling CloudTrail, the company can track user activity and changes to AWS resources, and monitor compliance with internal policies and external regulations.



D. Enable AWS Config and create rules for auditing and compliance purposes. AWS Config provides a detailed inventory of the AWS resources in your account, and continuously records changes to the configurations of those resources. By creating rules in AWS Config, the company can automate the evaluation of resource configurations against desired state, and receive alerts when configurations drift from compliance.

Options B, C, and E are not directly relevant to the requirement of tracking and auditing inventory and configuration changes.

upvoted 14 times
-   **zdi561** Most Recent 3 months ago

**Selected Answer: CD**



Amazon Web Services (AWS) CloudTrail is enabled by default for all AWS accounts.

upvoted 1 times
-   **Ruffyit** 11 months, 3 weeks ago

A. Enable AWS CloudTrail and use it for auditing. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS Command Line Interface (CLI), and AWS SDKs and APIs. By enabling CloudTrail, the company can track user activity and changes to AWS resources, and monitor compliance with internal policies and external regulations.

D. Enable AWS Config and create rules for auditing and compliance purposes. AWS Config provides a detailed inventory of the AWS resources in your account, and continuously records changes to the configurations of those resources. By creating rules in AWS Config, the company can automate the evaluation of resource configurations against desired state, and receive alerts when configurations drift from compliance.



Options B, C, and E are not directly relevant to the requirement of tracking and auditing inventory and configuration changes.

upvoted 2 times
-   **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: AD**

A. Enable AWS CloudTrail and use it for auditing. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS Command Line Interface (CLI), and AWS SDKs and APIs. By enabling CloudTrail, the company can track user activity and changes to AWS resources, and monitor compliance with internal policies and external regulations.

D. Enable AWS Config and create rules for auditing and compliance purposes. AWS Config provides a detailed inventory of the AWS resources in your account, and continuously records changes to the configurations of those resources. By creating rules in AWS Config, the company can automate the evaluation of resource configurations against desired state, and receive alerts when configurations drift from compliance.

upvoted 2 times
-   **mrsoa** 1 year, 3 months ago

**Selected Answer: CD**

I am gonna go with CD



AWS Cloudtrail is already enabled so no need to enable it and for the auding we are gonna use AWS config Answer D

C because Trusted advisor checks the security groups

upvoted 1 times

  **pentium75** 10 months, 1 week ago

"AWS CloudTrail is already enabled" says who?  
upvoted 3 times

  **awsgeek75** 9 months, 2 weeks ago

CloudTrail is not enabled by default or in the question scenario. Even if it was, Trusted Advisor would just give you recommendations and usage reports. It won't audit anything for you  
upvoted 1 times



  **kruasan** 1 year, 6 months ago

**Selected Answer: AD**

A) Enable AWS CloudTrail and use it for auditing.  
AWS CloudTrail provides a record of API calls and can be used to audit changes made to EC2 instances and security groups. By analyzing CloudTrail logs, the solutions architect can track who provisioned oversized instances or modified security groups without proper approval.  
D) Enable AWS Config and create rules for auditing and compliance purposes.  
AWS Config can record the configuration changes made to resources like EC2 instances and security groups. The solutions architect can create AWS Config rules to monitor for non-compliant changes, like launching certain instance types or opening security group ports without permission. AWS Config would alert on any violations of these rules.  
upvoted 3 times

  **kruasan** 1 year, 6 months ago

The other options would not fully meet the auditing and change tracking requirements:  
B) Data lifecycle policies control when EC2 instances are backed up or deleted but do not audit configuration changes.  
C) AWS Trusted Advisor security checks may detect some compliance violations after the fact but do not comprehensively log changes like AWS CloudTrail and AWS Config do.  
E) CloudFormation templates enable rollback but do not provide an audit trail of changes. The solutions architect would not know who made unauthorized modifications in the first place.  
upvoted 3 times

  **skiwili** 1 year, 8 months ago

**Selected Answer: AD**

Yes A and D  
upvoted 2 times

  **jennyka76** 1 year, 8 months ago

AGREE WITH ANSWER - A & D  
CloudTrail and Config  
upvoted 2 times

  **Neha999** 1 year, 8 months ago

CloudTrail and Config  
upvoted 3 times

A company has hundreds of Amazon EC2 Linux-based instances in the AWS Cloud. Systems administrators have used shared SSH keys to manage the instances. After a recent audit, the company’s security team is mandating the removal of all shared keys. A solutions architect must design a solution that provides secure access to the EC2 instances.

Which solution will meet this requirement with the LEAST amount of administrative overhead?

- A. Use AWS Systems Manager Session Manager to connect to the EC2 instances.
- B. Use AWS Security Token Service (AWS STS) to generate one-time SSH keys on demand.
- C. Allow shared SSH access to a set of bastion instances. Configure all other instances to allow only SSH access from the bastion instances.
- D. Use an Amazon Cognito custom authorizer to authenticate users. Invoke an AWS Lambda function to generate a temporary SSH key.

**Correct Answer:** A

*Community vote distribution*

A (85%)

Other

- Vlad**

Highly Voted

1 year, 8 months ago

Answer is A  
Using AWS Systems Manager Session Manager to connect to the EC2 instances is a secure option as it eliminates the need for inbound SSH ports and removes the requirement to manage SSH keys manually. It also provides a complete audit trail of user activity. This solution requires no additional software to be installed on the EC2 instances.  
upvoted 12 times
- pentium75**

Highly Voted

10 months, 1 week ago

**Selected Answer: A**

A - Systems Manager Session Manager has EXACTLY that purpose, 'providing secure access to EC2 instances'  
B - STS can generate temporary IAM credentials or access keys but NOT SSH keys  
C - Does not 'remove all shared keys' as requested  
D - Cognito is not meant for internal users, and whole setup is complex  
upvoted 7 times

**pentium75**

Most Recent

10 months, 1 week ago

**Selected Answer: A**

B - Querying is just a feature of Redshift but primarily it's a Data Warehouse - the question says nothing that historical data would have to be stored or accessed or analyzed  
upvoted 2 times

**Ruffyit**

11 months, 3 weeks ago

The key reasons why:  
  
STS can generate short-lived credentials that provide temporary access to the EC2 instances for administering them.  
The credentials can be generated on-demand each time access is needed, eliminating the risks of using permanent shared SSH keys.  
No infrastructure like bastion hosts needs to be maintained.  
The on-premises administrators can use the familiar SSH tools with the temporary keys.  
upvoted 2 times

**TariqKipkemei**

1 year ago

**Selected Answer: A**

Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys.  
upvoted 2 times

**Guru4Cloud**

1 year, 1 month ago

**Selected Answer: B**

The key reasons why:  
  
STS can generate short-lived credentials that provide temporary access to the EC2 instances for administering them.  
The credentials can be generated on-demand each time access is needed, eliminating the risks of using permanent shared SSH keys.  
No infrastructure like bastion hosts needs to be maintained.  
The on-premises administrators can use the familiar SSH tools with the temporary keys.  
upvoted 1 times

**pentium75**

10 months, 1 week ago



STS provides temporary IAM credentials, not SSH keys  
upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

Using AWS Security Token Service (AWS STS) to generate one-time SSH keys on demand is a secure and efficient way to provide access to the EC2 instances without the need for shared SSH keys. STS is a fully managed service that can be used to generate temporary security credentials, allowing systems administrators to connect to the EC2 instances without having to share SSH keys. The temporary credentials can be generated on demand, reducing the administrative overhead associated with managing SSH access

upvoted 2 times

  **ofinto** 1 year, 1 month ago

Can you please provide documentation about generating a one-time SSH with STS?

upvoted 2 times

  **kruasan** 1 year, 6 months ago

**Selected Answer: A**

AWS Systems Manager Session Manager provides secure shell access to EC2 instances without the need for SSH keys. It meets the security requirement to remove shared SSH keys while minimizing administrative overhead.

upvoted 2 times

  **Guru4Cloud** 1 year, 1 month ago

If the systems administrators need to access the EC2 instances from an on-premises environment, using Session Manager may not be the ideal solution.

upvoted 1 times

  **kruasan** 1 year, 6 months ago

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, 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.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

upvoted 3 times

  **kruasan** 1 year, 6 months ago

Who should use Session Manager?

Any AWS customer who wants to improve their security and audit posture, reduce operational overhead by centralizing access control on managed nodes, and reduce inbound node access.

Information Security experts who want to monitor and track managed node access and activity, close down inbound ports on managed nodes, or allow connections to managed nodes that don't have a public IP address.

Administrators who want to grant and revoke access from a single location, and who want to provide one solution to users for Linux, macOS, and Windows Server managed nodes.

Users who want to connect to a managed node with just one click from the browser or AWS CLI without having to provide SSH keys.

upvoted 3 times

  **Stanislav4907** 1 year, 7 months ago

**Selected Answer: C**

You guys seriously don't want to go to SMSM for Avery Single EC2. You have to create solution not used services for one time access. Bastion will give you option to manage 1000s EC2 machines from 1. Plus you can use Ansible from it.



upvoted 2 times

  **UnluckyDucky** 1 year, 7 months ago

Session Manager is the best practice and recommended way by Amazon to manage your instances. Bastion hosts require remote access therefore exposing them to the internet.



The most secure way is definitely session manager therefore answer A is correct imho.

upvoted 4 times

  **Zox42** 1 year, 7 months ago

Question:" the company's security team is mandating the removal of all shared keys", answer C can't be right because it says:"Allow shared SSH access to a set of bastion instances".

upvoted 7 times

  **Steve\_4542636** 1 year, 8 months ago

**Selected Answer: A**

I vote a

upvoted 2 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: A**

AWS Systems Manager Session Manager provides secure and auditable instance management without the need for any inbound connections or open ports. It allows you to manage your instances through an interactive one-click browser-based shell or through the AWS CLI. This means that you don't have to manage any SSH keys, and you don't have to worry about securing access to your instances as access is controlled through IAM policies.

upvoted 5 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

upvoted 3 times

  **jahmad0730** 1 year, 8 months ago

**Selected Answer: A**

Answer must be A

upvoted 3 times

  **jennyka76** 1 year, 8 months ago

ANSWER - A



AWS SESSION MANAGER IS CORRECT LEAST EFFORTS TO ACCESS LINUX SYSTEM IN AWS CONSOLE AND YOUR ARE ALREADY LOGIN TO AWS. SO NO NEED FOR THE TOKEN OR OTHER STUFF DONE IN THE BACKGROUND BY AWS. MAKES SENSE.

upvoted 3 times

  **cloudbusting** 1 year, 8 months ago

Answer is A

upvoted 4 times

  **zTopic** 1 year, 8 months ago

**Selected Answer: A**

Answer is A

upvoted 2 times

A company is using a fleet of Amazon EC2 instances to ingest data from on-premises data sources. The data is in JSON format and ingestion rates can be as high as 1 MB/s. When an EC2 instance is rebooted, the data in-flight is lost. The company’s data science team wants to query ingested data in near-real time.

Which solution provides near-real-time data querying that is scalable with minimal data loss?

- A. Publish data to Amazon Kinesis Data Streams, Use Kinesis Data Analytics to query the data.
- B. Publish data to Amazon Kinesis Data Firehose with Amazon Redshift as the destination. Use Amazon Redshift to query the data.
- C. Store ingested data in an EC2 instance store. Publish data to Amazon Kinesis Data Firehose with Amazon S3 as the destination. Use Amazon Athena to query the data.
- D. Store ingested data in an Amazon Elastic Block Store (Amazon EBS) volume. Publish data to Amazon ElastiCache for Redis. Subscribe to the Redis channel to query the data.

**Correct Answer:** A

*Community vote distribution*

A (74%)B (26%)

- LuckyAro**

Highly Voted

2 years, 2 months ago

Selected Answer: A

A: is the solution for the company's requirements. Publishing data to Amazon Kinesis Data Streams can support ingestion rates as high as 1 MB/s and provide real-time data processing. Kinesis Data Analytics can query the ingested data in real-time with low latency, and the solution can scale as needed to accommodate increases in ingestion rates or querying needs. This solution also ensures minimal data loss in the event of an EC2 instance reboot since Kinesis Data Streams has a persistent data store for up to 7 days by default.

upvoted 15 times
- bogobob**

Highly Voted

1 year, 5 months ago

Selected Answer: B

The fact they specifically mention "near real-time" twice tells me the correct answer is KDF. On top of which its easier to setup and maintain. KDS is really only needed if you need real-time. Also using redshift will mean permanent data retention. The data in A could be lost after a year. Redshift queries are slow but you're still querying near real-time data

upvoted 6 times
- Ernestokoro**

1 year, 4 months ago

You are very correct. see supporting link <https://jayendrapatil.com/aws-kinesis-data-streams-vs-kinesis-firehose/#:~:text=vs%20Kine...-,Purpose,into%20AWS%20products%20for%20processing>.

upvoted 2 times
- Dantecito**

Most Recent

2 months, 4 weeks ago

Selected Answer: A

Also A because "The company’s data science team wants to query ingested data in near-real time" and kinesis data analytics is doing that. With option B we query the tada once is in the data lake.

upvoted 1 times
- Dantecito**

2 months, 4 weeks ago

Selected Answer: A

Option a for me. I will assume that the EC2 instance is responsible for placing the data into the correct database. Therefore, we just need to wait for the instance to reboot. Since we can store the data in Kinesis Data Streams for up to 7 days, we are covered in that regard.

For option B, we are forced to use Amazon Redshift, and the problem statement does not mention anything related to that. I would choose option A because it does not affect the original requirements, and it helps us avoid in-flight data loss as requested.

upvoted 1 times
- wwwxxch**

4 months ago

Selected Answer: B

near-real time --> Kinesis Data Firehose  
And retention day of Kinesis Data Streams cannot be longer than 365 days

upvoted 1 times
- EllenLiu**

4 months, 2 weeks ago

Selected Answer: A

A: focus on performing complex data processing without in-flight data lost, not mention data persistence  
B: focus on data persist for later analysis

upvoted 1 times

  **LeonSauveterre** 5 months, 1 week ago

**Selected Answer: A**

Instance store & ElasiCache are all temporary storages, which cannot address data loss. That rules out C & D.

B: Kinesis Data Firehose is optimized for batch processing rather than real-time querying. It can indeed deliver data to S3 or Redshift, but there's a good chance the delay between ingestion and query availability cannot meet the "near-real-time" requirement.

upvoted 1 times

  **Lin878** 11 months ago

**Selected Answer: B**

[https://aws.amazon.com/pm/kinesis/?gclid=CjwKCAjwvIWzBhAlEiwAHHWgvRQuJmBubZDnO2GasDWwc2iBapfVD6GBelgj2JV6qkldm-K\\_CmMzmxoCdCwQAvD\\_BwE&trk=ee1218b7-7c10-4762-97df-274836a44566&sc\\_channel=ps&ef\\_id=CjwKCAjwvIWzBhAlEiwAHHWgvRQuJmBubZDnO2GasDWwc2iBapfVD6GBelgj2JV6qkldm-K\\_CmMzmxoCdCwQAvD\\_BwE:G:s&s\\_kwcid=AL!4422!3!651510255264!p!!g!!kinesis%20stream!19836376690!149589222920](https://aws.amazon.com/pm/kinesis/?gclid=CjwKCAjwvIWzBhAlEiwAHHWgvRQuJmBubZDnO2GasDWwc2iBapfVD6GBelgj2JV6qkldm-K_CmMzmxoCdCwQAvD_BwE&trk=ee1218b7-7c10-4762-97df-274836a44566&sc_channel=ps&ef_id=CjwKCAjwvIWzBhAlEiwAHHWgvRQuJmBubZDnO2GasDWwc2iBapfVD6GBelgj2JV6qkldm-K_CmMzmxoCdCwQAvD_BwE:G:s&s_kwcid=AL!4422!3!651510255264!p!!g!!kinesis%20stream!19836376690!149589222920)

upvoted 2 times

  **ray320x** 1 year, 3 months ago

Option A is actually correct. The question ask for minimal data loss and that query of data should be near real time, not the ingestion. Kinesis data analytics is near real time.

Recent changes to Redshift actually make B correct as well, but A is also correct.

upvoted 2 times

  **dkw2342** 1 year, 2 months ago

Streaming ingestion provides low-latency, high-speed ingestion of stream data from Amazon Kinesis Data Streams and Amazon Managed Streaming for Apache Kafka into an Amazon Redshift provisioned or Amazon Redshift Serverless materialized view.[1]

Option B mentions Kinesis Data Firehose (now just Firehose), so this won't work.

Option A is the correct answer.

[1]<https://docs.aws.amazon.com/redshift/latest/dg/materialized-view-streaming-ingestion.html>

upvoted 2 times

  **farnamjam** 1 year, 3 months ago

**Selected Answer: A**

Comparison to other options:

B. Kinesis Data Firehose with Redshift: While Redshift is scalable, it doesn't offer real-time querying capabilities. Data needs to be loaded into Redshift from Firehose, introducing latency.

C. EC2 instance store with Kinesis Data Firehose and S3: Storing data in an EC2 instance store is not persistent and data will be lost during reboots. EBS volumes are more appropriate for persistent storage, but the architecture becomes more complex.

D. EBS volume with ElastiCache and Redis: While ElastiCache offers fast in-memory storage, it's not designed for high-volume data ingestion like 1 MB/s. It might struggle with scalability and persistence.

upvoted 3 times

  **Firdous586** 1 year, 3 months ago

I don't understand why people are giving wrong information in the QUESTION its clearly mentioned near Real Time

Kinesis Data Streams is for Real time

Where are Kinesis Datafirehose is for Near real time there for answer is B only

upvoted 5 times

  **Marco\_St** 1 year, 4 months ago

**Selected Answer: A**

Read the question: near real-time querying of data.... it is more about real-time data query once the data is ingested, It does not mention how long time the data needs to be stored. A is better option. B introduces delay of data buffer before it can be queried in redshift

upvoted 1 times

  **practice\_makes\_perfect** 1 year, 5 months ago

**Selected Answer: B**



A is not correct because Kinesis can only store data up to 1 year. The solution need to support querying ALL data instead of "recent" data.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

Says who? They want to "query ingested data in near-real time", it does not say anything about historical data.

upvoted 2 times

  **Ruffyit** 1 year, 5 months ago

A: is the solution for the company's requirements. Publishing data to Amazon Kinesis Data Streams can support ingestion rates as high as 1 MB/s and provide real-time data processing. Kinesis Data Analytics can query the ingested data in real-time with low latency, and the solution can scale as needed to accommodate increases in ingestion rates or querying needs. This solution also ensures minimal data loss in the event of an EC2 instance reboot since Kinesis Data Streams has a persistent data store for up to 7 days by default.

upvoted 2 times

🗨️ 👤 **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: A**

Publish data to Amazon Kinesis Data Streams, Use Kinesis Data Analytics to query the data  
upvoted 3 times

🗨️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

- Provide near-real-time data ingestion into Kinesis Data Streams with the ability to handle the 1 MB/s ingestion rate. Data would be stored redundantly across shards.
- Enable near-real-time querying of the data using Kinesis Data Analytics. SQL queries can be run directly against the Kinesis data stream.
- Minimize data loss since data is replicated across shards. If an EC2 instance is rebooted, the data stream is still accessible.
- Scale seamlessly to handle varying ingestion and query rates.

upvoted 4 times

🗨️ 👤 **Nikki013** 1 year, 8 months ago

**Selected Answer: A**

Answer is A as it will provide a more streamlined solution.

Using B (Firehose + Redshift) will involve sending the data to an S3 bucket first and then copying the data to Redshift which will take more time.

<https://docs.aws.amazon.com/firehose/latest/dev/what-is-this-service.html>

upvoted 5 times

What should a solutions architect do to ensure that all objects uploaded to an Amazon S3 bucket are encrypted?

- A. Update the bucket policy to deny if the PutObject does not have an s3:x-amz-acl header set.
- B. Update the bucket policy to deny if the PutObject does not have an s3:x-amz-acl header set to private.
- C. Update the bucket policy to deny if the PutObject does not have an aws:SecureTransport header set to true.
- D. Update the bucket policy to deny if the PutObject does not have an x-amz-server-side-encryption header set.

**Correct Answer: D**

*Community vote distribution*

D (100%)

  **bdp123** Highly Voted 1 year, 8 months ago

**Selected Answer: D**

<https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/#:~:text=Solution%20overview>  
upvoted 15 times

  **Grace83** 1 year, 7 months ago

Thank you!  
upvoted 1 times

  **Guru4Cloud** Highly Voted 1 year, 2 months ago

**Selected Answer: D**


The x-amz-server-side-encryption header is used to specify the encryption method that should be used to encrypt objects uploaded to an Amazon S3 bucket. By updating the bucket policy to deny if the PutObject does not have this header set, the solutions architect can ensure that all objects uploaded to the bucket are encrypted.  
upvoted 6 times

  **awsgeek75** Most Recent 10 months ago

**Selected Answer: D**

Related reading because (as of Jan 2023) S3 buckets have encryption enabled by default.  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingServerSideEncryption.html>


"If you require your data uploads to be encrypted using only Amazon S3 managed keys, you can use the following bucket policy. For example, the following bucket policy denies permissions to upload an object unless the request includes the x-amz-server-side-encryption header to request server-side encryption:"  
upvoted 4 times

  **kruasan** 1 year, 6 months ago

To encrypt an object at the time of upload, you need to add a header called x-amz-server-side-encryption to the request to tell S3 to encrypt the object using SSE-C, SSE-S3, or SSE-KMS. The following code example shows a Put request using SSE-S3.  
<https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/>  
upvoted 4 times



  **kruasan** 1 year, 6 months ago

The other options would not enforce encryption:  
A) Requiring an s3:x-amz-acl header does not mandate encryption. This header controls access permissions.  
B) Requiring an s3:x-amz-acl header set to private also does not enforce encryption. It only enforces private access permissions.  
C) Requiring an aws:SecureTransport header ensures uploads use SSL but does not specify that objects must be encrypted. Encryption is not required when using SSL transport.  
upvoted 5 times

  **kruasan** 1 year, 6 months ago

**Selected Answer: D**

To encrypt an object at the time of upload, you need to add a header called x-amz-server-side-encryption to the request to tell S3 to encrypt the object using SSE-C, SSE-S3, or SSE-KMS. The following code example shows a Put request using SSE-S3.  
<https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/>  
upvoted 2 times

  **Sbbh** 1 year, 7 months ago

Confusing question. It doesn't state clearly if the object needs to be encrypted at-rest or in-transit  
upvoted 6 times

  **Guru4Cloud** 1 year, 2 months ago

That's true  
upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

**Selected Answer: D**

I vote d

upvoted 2 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: D**



To ensure that all objects uploaded to an Amazon S3 bucket are encrypted, the solutions architect should update the bucket policy to deny any PutObject requests that do not have an x-amz-server-side-encryption header set. This will prevent any objects from being uploaded to the bucket unless they are encrypted using server-side encryption.

upvoted 4 times

  **jennyka76** 1 year, 8 months ago

answer - D

upvoted 2 times

  **zTopic** 1 year, 8 months ago

**Selected Answer: D**

Answer is D

upvoted 2 times

  **Neorem** 1 year, 8 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/amazon-s3-policy-keys.html>

upvoted 2 times

A solutions architect is designing a multi-tier application for a company. The application's users upload images from a mobile device. The application generates a thumbnail of each image and returns a message to the user to confirm that the image was uploaded successfully.

The thumbnail generation can take up to 60 seconds, but the company wants to provide a faster response time to its users to notify them that the original image was received. The solutions architect must design the application to asynchronously dispatch requests to the different application tiers.

What should the solutions architect do to meet these requirements?

- A. Write a custom AWS Lambda function to generate the thumbnail and alert the user. Use the image upload process as an event source to invoke the Lambda function.
- B. Create an AWS Step Functions workflow. Configure Step Functions to handle the orchestration between the application tiers and alert the user when thumbnail generation is complete.
- C. Create an Amazon Simple Queue Service (Amazon SQS) message queue. As images are uploaded, place a message on the SQS queue for thumbnail generation. Alert the user through an application message that the image was received.
- D. Create Amazon Simple Notification Service (Amazon SNS) notification topics and subscriptions. Use one subscription with the application to generate the thumbnail after the image upload is complete. Use a second subscription to message the user's mobile app by way of a push notification after thumbnail generation is complete.

**Correct Answer: C**

*Community vote distribution*

C (94%)

6%

- Steve\_4542636**

Highly Voted

1 year, 8 months ago

Selected Answer: C

I've noticed there are a lot of questions about decoupling services and SQS is almost always the answer.

upvoted 32 times
- Neha999**

Highly Voted

1 year, 8 months ago

D

SNS fan out

upvoted 13 times
- Manimgh**

Most Recent

2 weeks, 6 days ago

Selected Answer: C

C is correct -> take up to 60 seconds = SQS

upvoted 1 times
- LoXoL**

8 months, 4 weeks ago

They don't look like real answers from the official exam...

upvoted 1 times
- awsgeek75**

9 months, 2 weeks ago

Selected Answer: C

Each option is badly worded:

A: "generate the thumbnail and alert the user" doesn't sound sequential so could alert the user during, before or after the thumbnail generation whichever way you interpret it.

B: this is sequential and won't alert until the steps are complete

D: Could work without with the risk of notification loss so C is better but this is also ok

upvoted 2 times
- awsgeek75**

10 months ago

Selected Answer: C

Safe answer is C but B is so badly worded that it can mean anything to confuse people. Step functions to use tiers. What if on of the step is to inform to the user and move on to next step. Anyway, I'll chose C for the exam as it is cleaner.

upvoted 2 times
- wsdasdasdqwda**

1 year ago

... asynchronously dispatch ... => Amazon SQS

upvoted 7 times

🗨️ 👤 **TariqKipkemei** 1 year ago

**Selected Answer: C**

Asynchronous, Decoupling = Amazon Simple Queue Service  
upvoted 5 times

🗨️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

SQS is a fully managed message queuing service that can be used to decouple different parts of an application.  
upvoted 2 times

🗨️ 👤 **Zox42** 1 year, 7 months ago

**Selected Answer: C**

Answers B and D alert the user when thumbnail generation is complete. Answer C alerts the user through an application message that the image was received.  
upvoted 5 times

🗨️ 👤 **Sbbh** 1 year, 7 months ago

B:  
Use cases for Step Functions vary widely, from orchestrating serverless microservices, to building data-processing pipelines, to defining a security-incident response. As mentioned above, Step Functions may be used for synchronous and asynchronous business processes.  
upvoted 1 times

🗨️ 👤 **AlessandraSAA** 1 year, 8 months ago

why not B?  
upvoted 4 times

🗨️ 👤 **Wael216** 1 year, 8 months ago

**Selected Answer: C**

Creating an Amazon Simple Queue Service (SQS) message queue and placing messages on the queue for thumbnail generation can help separate the image upload and thumbnail generation processes.  
upvoted 2 times

🗨️ 👤 **vindahake** 1 year, 8 months ago

C  
The key here is "a faster response time to its users to notify them that the original image was received." i.e user needs to be notified when image was received and not after thumbnail was created.  
upvoted 3 times

🗨️ 👤 **AlmeroSenior** 1 year, 8 months ago

**Selected Answer: C**

A looks like the best way , but its essentially replacing the mentioned app , that's not the ask  
upvoted 2 times

🗨️ 👤 **Mickey321** 1 year, 8 months ago

Selected Answer: A  
<https://docs.aws.amazon.com/lambda/latest/dg/with-s3-tutorial.html>  
upvoted 1 times

🗨️ 👤 **bdp123** 1 year, 8 months ago

**Selected Answer: C**

C is the only one that makes sense  
upvoted 2 times

A company’s facility has badge readers at every entrance throughout the building. When badges are scanned, the readers send a message over HTTPS to indicate who attempted to access that particular entrance.

A solutions architect must design a system to process these messages from the sensors. The solution must be highly available, and the results must be made available for the company’s security team to analyze.



Which system architecture should the solutions architect recommend?


- A. Launch an Amazon EC2 instance to serve as the HTTPS endpoint and to process the messages. Configure the EC2 instance to save the results to an Amazon S3 bucket.
- B. Create an HTTPS endpoint in Amazon API Gateway. Configure the API Gateway endpoint to invoke an AWS Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- C. Use Amazon Route 53 to direct incoming sensor messages to an AWS Lambda function. Configure the Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- D. Create a gateway VPC endpoint for Amazon S3. Configure a Site-to-Site VPN connection from the facility network to the VPC so that sensor data can be written directly to an S3 bucket by way of the VPC endpoint.

**Correct Answer:** B

*Community vote distribution*

B (100%)



-   **kruasan**


Highly Voted 

 1 year ago

Selected Answer: B

- Option A would not provide high availability. A single EC2 instance is a single point of failure.
  - Option B provides a scalable, highly available solution using serverless services. API Gateway and Lambda can scale automatically, and DynamoDB provides a durable data store.
  - Option C would expose the Lambda function directly to the public Internet, which is not a recommended architecture. API Gateway provides an abstraction layer and additional features like access control.
  - Option D requires configuring a VPN to AWS which adds complexity. It also saves the raw sensor data to S3, rather than processing it and storing the results.

upvoted 22 times
-   **TariqKipkemei**

Highly Voted 



 6 months, 3 weeks ago


Selected Answer: B

Highly available = Serverless

The readers send a message over HTTPS = HTTPS endpoint in Amazon API Gateway

Process these messages from the sensors = AWS Lambda function

upvoted 8 times
-   **Guru4Cloud**

Most Recent 

 8 months ago

Selected Answer: B



The correct answer is B. Create an HTTPS endpoint in Amazon API Gateway. Configure the API Gateway endpoint to invoke an AWS Lambda function to process the messages and save the results to an Amazon DynamoDB table.

Here are the reasons why:

API Gateway is a highly scalable and available service that can be used to create and expose RESTful APIs.



Lambda is a serverless compute service that can be used to process events and data.

DynamoDB is a NoSQL database that can be used to store data in a scalable and highly available way.

upvoted 4 times
-   **Steve\_4542636** 1 year, 2 months ago



Selected Answer: B

I vote B

upvoted 2 times
-   **KZM** 1 year, 2 months ago

It is option "B"



Option "B" can provide a system with highly scalable, fault-tolerant, and easy to manage.

upvoted 2 times
-   **LuckyAro** 1 year, 2 months ago

**Selected Answer: B**

Deploy Amazon API Gateway as an HTTPS endpoint and AWS Lambda to process and save the messages to an Amazon DynamoDB table. This option provides a highly available and scalable solution that can easily handle large amounts of data. It also integrates with other AWS services, making it easier to analyze and visualize the data for the security team.

upvoted 4 times

  **zTopic** 1 year, 2 months ago

**Selected Answer: B**

B is Correct

upvoted 4 times

A company wants to implement a disaster recovery plan for its primary on-premises file storage volume. The file storage volume is mounted from an Internet Small Computer Systems Interface (iSCSI) device on a local storage server. The file storage volume holds hundreds of terabytes (TB) of data.

The company wants to ensure that end users retain immediate access to all file types from the on-premises systems without experiencing latency.

Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

- A. Provision an Amazon S3 File Gateway as a virtual machine (VM) that is hosted on premises. Set the local cache to 10 TB. Modify existing applications to access the files through the NFS protocol. To recover from a disaster, provision an Amazon EC2 instance and mount the S3 bucket that contains the files.
- B. Provision an AWS Storage Gateway tape gateway. Use a data backup solution to back up all existing data to a virtual tape library. Configure the data backup solution to run nightly after the initial backup is complete. To recover from a disaster, provision an Amazon EC2 instance and restore the data to an Amazon Elastic Block Store (Amazon EBS) volume from the volumes in the virtual tape library.
- C. Provision an AWS Storage Gateway Volume Gateway cached volume. Set the local cache to 10 TB. Mount the Volume Gateway cached volume to the existing file server by using iSCSI, and copy all files to the storage volume. Configure scheduled snapshots of the storage volume. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.
- D. Provision an AWS Storage Gateway Volume Gateway stored volume with the same amount of disk space as the existing file storage volume. Mount the Volume Gateway stored volume to the existing file server by using iSCSI, and copy all files to the storage volume. Configure scheduled snapshots of the storage volume. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.

**Correct Answer: D**

*Community vote distribution*

D (83%)

C (17%)

- Grace83**

Highly Voted

2 years, 1 month ago

D is the correct answer

Volume Gateway CACHED Vs STORED  
Cached = stores a subset of frequently accessed data locally  
Stored = Retains the ENTIRE ("all file types") in on prem data centre  
upvoted 28 times
- netcj**

Highly Voted

1 year, 7 months ago

Selected Answer: D

"users retain immediate access to all file types"  
immediate cannot be cached -> D  
upvoted 6 times
- dkw2342**

Most Recent

1 year, 2 months ago

Bad question. No RTO/RPO, so impossible to properly answer. They probably want to hear option D.

Depending on RPO, option B is also an adequate solution (data remains immediately accessible without experiencing latency via existing infrastructure, backup to cloud for DR). Also, this option requires LESS changes to existing infra than A. Only argument against B is that VTLs are usually used for legacy DR solutions, not for new ones, where object storage such as S3 is usually supported natively.  
upvoted 2 times
- MrPCarrot**

1 year, 2 months ago

Answer is C go argue somewhere.  
upvoted 2 times
- awsgeek75**

1 year, 4 months ago

Selected Answer: D

A,B are wrong types of gateways for hundreds of TB of data that needs immediate access on-prem. C limits to 10TB. D provides access to all the files.  
upvoted 2 times
- pentium75**

1 year, 4 months ago

**Selected Answer: D**

"Immediate access to all file types from the on-premises systems without experiencing latency" requirement is not met by C. Also the solution is meant for DR purposes, the primary storage for the data should remain on premises.

upvoted 3 times

  **daniel1** 1 year, 6 months ago

**Selected Answer: C**



From chatGPT4  
Considering the requirements of minimal infrastructure change, immediate file access, and low-latency, Option C: Provisioning an AWS Storage Gateway Volume Gateway (cached volume) with a 10 TB local cache, seems to be the most fitting solution. This setup aligns with the existing iSCSI setup and provides a local cache for low-latency access, while also configuring scheduled snapshots for disaster recovery. In the event of a disaster, restoring a snapshot to an Amazon EBS volume and attaching it to an Amazon EC2 instance as described in this option would align with the recovery objective.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

ChatGPT is wrong. "Immediate access to all file types from the on-premises systems without experiencing latency" needs "stored volume" type. With "cached volume" not all data will be available locally.

upvoted 9 times

  **LoXoL** 1 year, 2 months ago

pentium75 is right.

upvoted 2 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: D**

End users retain immediate access to all file types = Volume Gateway stored volume

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

ddddddddd

upvoted 3 times

  **alexandercamachop** 1 year, 11 months ago

**Selected Answer: D**


Correct answer is Volume Gateway Stored which keeps all data on premises.  
To have immediate access to the data. Cached is for frequently accessed data only.

upvoted 3 times

  **omoakin** 1 year, 11 months ago



CCCCCCCCCCCCCCCC

upvoted 1 times

  **24b2e9e** 10 months, 2 weeks ago

The stored volume configuration stores the entire data set on-premises and asynchronously backs up the data to AWS. The cached volume configuration stores recently accessed data on-premises, and the remaining data is stored in Amazon S3  
-that is why D is right

upvoted 2 times

  **lucdt4** 1 year, 11 months ago

**Selected Answer: D**

D is the correct answer  
Volume Gateway CACHED Vs STORED  
Cached = stores a data recently at local  
Stored = Retains the ENTIRE ("all file types") in on prem data centre

upvoted 2 times

  **rushi0611** 1 year, 12 months ago

**Selected Answer: D**

In the cached mode, your primary data is written to S3, while retaining your frequently accessed data locally in a cache for low-latency access. In the stored mode, your primary data is stored locally and your entire dataset is available for low-latency access while asynchronously backed up to AWS.  
Reference: <https://aws.amazon.com/storagegateway/faqs/>  
Good luck.

upvoted 3 times

  **kruasan** 2 years ago

**Selected Answer: D**

It is stated the company wants to keep the data locally and have DR plan in cloud. It points directly to the volume gateway

upvoted 2 times

  **UnluckyDucky** 2 years, 1 month ago

**Selected Answer: D**

"The company wants to ensure that end users retain immediate access to all file types from the on-premises systems "

D is the correct answer.  
upvoted 3 times

  **CapJackSparrow** 2 years, 1 month ago



**Selected Answer: C**

all file types, NOT all files. Volume mode can not cache 100TBs.  
upvoted 3 times

  **eddie5049** 1 year, 12 months ago

<https://docs.aws.amazon.com/storagegateway/latest/vgw/StorageGatewayConcepts.html>

Stored volumes can range from 1 GiB to 16 TiB in size and must be rounded to the nearest GiB. Each gateway configured for stored volumes can support up to 32 volumes and a total volume storage of 512 TiB (0.5 PiB).  
upvoted 2 times

  **MssP** 2 years, 1 month ago

all file types. Cached only save the most frecuently or lastest accesed. If you didn´t access any type for a long time, you will not cache it -> No immediate access  
upvoted 4 times

  **pentium75** 1 year, 4 months ago

Also the solution is meant for DR purposes, it's not like they need more storage or so.  
upvoted 2 times

  **WherecanIstart** 2 years, 1 month ago

**Selected Answer: D**

"The company wants to ensure that end users retain immediate access to all file types from the on-premises systems "

This points to stored volumes..  
upvoted 2 times

A company is hosting a web application from an Amazon S3 bucket. The application uses Amazon Cognito as an identity provider to authenticate users and return a JSON Web Token (JWT) that provides access to protected resources that are stored in another S3 bucket.

Upon deployment of the application, users report errors and are unable to access the protected content. A solutions architect must resolve this issue by providing proper permissions so that users can access the protected content.

Which solution meets these requirements?

- A. Update the Amazon Cognito identity pool to assume the proper IAM role for access to the protected content.
- B. Update the S3 ACL to allow the application to access the protected content.
- C. Redeploy the application to Amazon S3 to prevent eventually consistent reads in the S3 bucket from affecting the ability of users to access the protected content.
- D. Update the Amazon Cognito pool to use custom attribute mappings within the identity pool and grant users the proper permissions to access the protected content.

**Correct Answer: A**

*Community vote distribution*

A (95%)

5%

  alexandercamachop Highly Voted 11 months, 1 week ago

**Selected Answer: A**

To resolve the issue and provide proper permissions for users to access the protected content, the recommended solution is:

A. Update the Amazon Cognito identity pool to assume the proper IAM role for access to the protected content.

Explanation:

Amazon Cognito provides authentication and user management services for web and mobile applications. In this scenario, the application is using Amazon Cognito as an identity provider to authenticate users and obtain JSON Web Tokens (JWTs). The JWTs are used to access protected resources stored in another S3 bucket. To grant users access to the protected content, the proper IAM role needs to be assumed by the identity pool in Amazon Cognito. By updating the Amazon Cognito identity pool with the appropriate IAM role, users will be authorized to access the protected content in the S3 bucket.

upvoted 16 times

  alexandercamachop 11 months, 1 week ago

Option B is incorrect because updating the S3 ACL (Access Control List) will only affect the permissions of the application, not the users accessing the content.

Option C is incorrect because redeploying the application to Amazon S3 will not resolve the issue related to user access permissions.

Option D is incorrect because updating custom attribute mappings in Amazon Cognito will not directly grant users the proper permissions to access the protected content.

upvoted 12 times

  LuckyAro Highly Voted 1 year, 2 months ago

**Selected Answer: A**

A is the best solution as it directly addresses the issue of permissions and grants authenticated users the necessary IAM role to access the protected content.

A suggests updating the Amazon Cognito identity pool to assume the proper IAM role for access to the protected content. This is a valid solution, as it would grant authenticated users the necessary permissions to access the protected content.

upvoted 6 times

  Marco\_St Most Recent 4 months, 3 weeks ago

**Selected Answer: A**

IAM role is assinged to IAM users or groups or assumed by AWS service. So IAM role is given to AWS Cognito service which provides temporary AWS credentials to authenticated users. so technically When a user is authenticated by Cognito, they receive temporary credentials based on the IAM role tied to the Cognito identity pool. If this IAM role has permissions to access certain S3 buckets or objects, the authenticated user will be able to access those resources as allowed by the role. This service is used under the hood by Cognito to provide these temporary credentials. The credentials are limited in time and scope based on the permissions defined in the IAM role.

upvoted 2 times

  Guru4Cloud 8 months ago

**Selected Answer: A**

A. Update the Amazon Cognito identity pool to assume the proper IAM role for access to the protected content.  
upvoted 3 times

  **Abrar2022** 10 months, 4 weeks ago



**Selected Answer: A**

Services access other services via IAM Roles. Hence why updating AWS Cognito identity pool to assume proper IAM Role is the right solution.  
upvoted 2 times

  **shanwford** 1 year ago

**Selected Answer: A**



Amazon Cognito identity pools assign your authenticated users a set of temporary, limited-privilege credentials to access your AWS resources. The permissions for each user are controlled through IAM roles that you create. <https://docs.aws.amazon.com/cognito/latest/developerguide/role-based-access-control.html>  
upvoted 3 times

  **Brak** 1 year, 1 month ago



**Selected Answer: D**

A makes no sense - Cognito is not accessing the S3 resource. It just returns the JWT token that will be attached to the S3 request.



D is the right answer, using custom attributes that are added to the JWT and used to grant permissions in S3. See <https://docs.aws.amazon.com/cognito/latest/developerguide/using-attributes-for-access-control-policy-example.html> for an example.  
upvoted 2 times

  **Abhineet9148232** 1 year, 1 month ago

But even D requires setting up the permissions as bucket policy (as show in the shared example) which includes higher overhead than managing permissions attached to specific roles.  
upvoted 3 times

  **asoli** 1 year, 1 month ago

A says "Identity Pool"  
According to AWS: "With an identity pool, your users can obtain temporary AWS credentials to access AWS services, such as Amazon S3 and DynamoDB."  
So, answer is A  
upvoted 4 times



  **Steve\_4542636** 1 year, 2 months ago

**Selected Answer: A**

Services access other services via IAM Roles.  
upvoted 2 times

  **jennyka76** 1 year, 2 months ago

ANSWER - A  
<https://docs.aws.amazon.com/cognito/latest/developerguide/tutorial-create-identity-pool.html>  
You have to create an custom role such as read-only  
upvoted 5 times

  **zTopic** 1 year, 2 months ago

**Selected Answer: A**

Answer is A  
upvoted 3 times

An image hosting company uploads its large assets to Amazon S3 Standard buckets. The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent. The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.

Which combination of actions should a solutions architect recommend to meet these requirements? (Choose two.)

- A. Move assets to S3 Intelligent-Tiering after 30 days.
- B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
- C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
- D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

**Correct Answer: AB**

*Community vote distribution*

AB (65%)

BD (25%)

9%

- Neha999**

Highly Voted

2 years, 2 months ago

AB

A : Access Pattern for each object inconsistent, Infrequent Access

B : Deleting Incomplete Multipart Uploads to Lower Amazon S3 Costs

upvoted 25 times
- TungPham**

Highly Voted

2 years, 2 months ago

**Selected Answer: AB**

B because Abort Incomplete Multipart Uploads Using S3 Lifecycle => <https://aws.amazon.com/blogs/aws-cloud-financial-management/discovering-and-deleting-incomplete-multipart-uploads-to-lower-amazon-s3-costs/>

A because The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent => random access => S3 Intelligent-Tiering

upvoted 16 times
- LeonSauveterre**

Most Recent

5 months, 1 week ago

**Selected Answer: AB**

Inconsistent access patterns: A is correct, and D is out (could work but less ideal and costs more).

C: "object delete markers" only apply to versioned buckets.

upvoted 1 times
- ChymKuBoy**

9 months, 2 weeks ago

**Selected Answer: AB**

AB for sure

upvoted 3 times
- bujuman**

1 year, 1 month ago

**Selected Answer: BD**

If we consider these statements:

1. For the first 30 days after upload, the objects will be accessed frequently

2.The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent

3.The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.

4.The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again.

Statements 1 and 2 cloudl be completed with option D and not A because datas are infrequently accessed only after 30 days.

Due to usage of multipart upload, to meet requirement regarding cost optimization, option B will be used to clean up buckets uncompleted file parts(statements 3 & 4).

upvoted 4 times
- NayeraB**

1 year, 2 months ago

**Selected Answer: AD**

Because A & D address the main ask, there's no mention of cost optimization.

upvoted 2 times
- NayeraB**

1 year, 2 months ago

\*Facepalm\* It does ask for reducing the cost, A&B it is!

upvoted 3 times

  **NayeraB** 1 year, 2 months ago

**Selected Answer: AC**

Because A & C address the main ask, there's no mention of cost optimization.

upvoted 1 times

  **NayeraB** 1 year, 2 months ago

Not C ':D, I meant to say A&D. Added another vote for that one.

upvoted 2 times

  **awsgeek75** 1 year, 4 months ago

**Selected Answer: AB**

A as the access pattern for each object is inconsistent so let AWS AWS do the handling.

B deals with multi-part duplication issues and saves money by deleting incomplete uploads

C No mention of deleted object so this is a distractor

D The objects will be accessed in unpredictable pattern so can't use this

E Not HA compliant

upvoted 3 times

  **awsgeek75** 1 year, 4 months ago

Also, don't be confused by 30 days. The question has tricky wording: " The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent"

It does NOT say that objects will be accessed less frequently after 30 days. It says the access is unpredictable which means it could go up or down. Don't make assumptions.

upvoted 4 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: AB**



C is nonsense

E does not meet the "high availability and resiliency" requirement

B is obvious (incomplete multipart uploads consume space -> cost money)

The tricky part is A vs. D. However, 'inconsistent access patterns' are the primary use case for Intelligent-Tiering. There are probably objects that will never be accessed and that would be moved to Glacier Instant Retrieval by Intelligent-Tiering, thus the overall cost would be lower than with D.

upvoted 4 times

  **osmk** 1 year, 4 months ago

bd <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-infreq-data-access> => S3 Standard-IA objects are resilient to the loss of an Availability Zone. This storage class offers greater availability and resiliency than the S3 One Zone-IA class

upvoted 2 times

  **raymondfekry** 1 year, 4 months ago

**Selected Answer: AB**

I wouldnt go with D since " the access patterns for each object will be inconsistent.", so we cannot move all assets to IA

upvoted 2 times

  **Marco\_St** 1 year, 4 months ago

**Selected Answer: AB**

incosistent access pattern brings more sense to use Intelligent-Tiering after 30 days which also covers infrequent access.

upvoted 2 times



  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: AB**

A. Move assets to S3 Intelligent-Tiering after 30 days.

B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.

upvoted 2 times

  **vini15** 1 year, 9 months ago

should be A and B

upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: BD**

Option A has not been mentioned for resiliency in S3, check the page: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/disaster-recovery-resiliency.html>

Therefore, I am with B & D choices.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

Intelligent-Tiering just moves to Standard-IA or Glacier Instant Access based on access patterns. This does not affect resiliency.

upvoted 2 times

  **alexandercamachop** 1 year, 11 months ago

**Selected Answer: AB**

- A. Move assets to S3 Intelligent-Tiering after 30 days.
- B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.

Explanation:

A. Moving assets to S3 Intelligent-Tiering after 30 days: This storage class automatically analyzes the access patterns of objects and moves them between frequent access and infrequent access tiers. Since the objects will be accessed frequently for the first 30 days, storing them in the frequent access tier during that period optimizes performance. After 30 days, when the access patterns become inconsistent, S3 Intelligent-Tiering will automatically move the objects to the infrequent access tier, reducing storage costs.

B. Configuring an S3 Lifecycle policy to clean up incomplete multipart uploads: Multipart uploads are used for large objects, and incomplete multipart uploads can consume storage space if not cleaned up. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, unnecessary storage costs can be avoided.

upvoted 2 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: AD**

AD.

B makes no sense because multipart uploads overwrite objects that are already uploaded. The question never says this is a problem.

upvoted 2 times

  **VellaDevil** 1 year, 10 months ago

Questions says to optimize cost and if incomplete multiparts are not aborted it will still use capacity on S3 Bucket thus increase unnecessary cost.

upvoted 3 times

A solutions architect must secure a VPC network that hosts Amazon EC2 instances. The EC2 instances contain highly sensitive data and run in a private subnet. According to company policy, the EC2 instances that run in the VPC can access only approved third-party software repositories on the internet for software product updates that use the third party’s URL. Other internet traffic must be blocked.

Which solution meets these requirements?

- A. Update the route table for the private subnet to route the outbound traffic to an AWS Network Firewall firewall. Configure domain list rule groups.
- B. Set up an AWS WAF web ACL. Create a custom set of rules that filter traffic requests based on source and destination IP address range sets.
- C. Implement strict inbound security group rules. Configure an outbound rule that allows traffic only to the authorized software repositories on the internet by specifying the URLs.
- D. Configure an Application Load Balancer (ALB) in front of the EC2 instances. Direct all outbound traffic to the ALB. Use a URL-based rule listener in the ALB’s target group for outbound access to the internet.

**Correct Answer: A**

*Community vote distribution*

A (91%)

9%

- Bhawesh**

Highly Voted

1 year, 8 months ago

Selected Answer: A

Correct Answer A. Send the outbound connection from EC2 to Network Firewall. In Network Firewall, create stateful outbound rules to allow certain domains for software patch download and deny all other domains.

<https://docs.aws.amazon.com/network-firewall/latest/developerguide/suricata-examples.html#suricata-example-domain-filtering>

upvoted 15 times
- Guru4Cloud**

1 year, 2 months ago

Option A uses a network firewall which is overkill for instance-level rules.

upvoted 2 times
- UnluckyDucky**

Highly Voted

1 year, 7 months ago

Selected Answer: A

Can't use URLs in outbound rule of security groups. URL Filtering screams Firewall.

upvoted 11 times
- TheFivePips**

Most Recent

8 months, 1 week ago

Selected Answer: A

Security Groups operate at the transport layer (Layer 4) of the OSI model and are primarily concerned with controlling traffic based on IP addresses, ports, and protocols. They do not have the capability to inspect or filter traffic based on URLs.

The solution to restrict outbound internet traffic based on specific URLs typically involves using a proxy or firewall that can inspect the application layer (Layer 7) of the OSI model, where URL information is available.

AWS Network Firewall operates at the network and application layers, allowing for more granular control, including the ability to inspect and filter traffic based on domain names or URLs.

By configuring domain list rule groups in AWS Network Firewall, you can specify which URLs are allowed for outbound traffic.

This option is more aligned with the requirement of allowing access to approved third-party software repositories based on their URLs.

upvoted 5 times
- awsgeek75**

10 months ago

Selected Answer: A

<https://aws.amazon.com/network-firewall/features/>

"Web filtering:

AWS Network Firewall supports inbound and outbound web filtering for unencrypted web traffic. For encrypted web traffic, Server Name Indication (SNI) is used for blocking access to specific sites. SNI is an extension to Transport Layer Security (TLS) that remains unencrypted in the traffic flow and indicates the destination hostname a client is attempting to access over HTTPS. In addition, \*\*AWS Network Firewall can filter fully qualified domain names (FQDN).\*\*"

Always use an AWS product if the advertisement meets the use case.

upvoted 4 times
- farnamjam**



10 months, 1 week ago



Selected Answer: A



AWS Network Firewall



  - Protect your entire Amazon VPC



• From Layer 3 to Layer 7 protection  
• Any direction, you can inspect  
Traffic filtering: Allow, drop, or alert for the traffic that matches the rules, • Active flow inspection to intrusion prevention  
upvoted 2 times



  **Subhrangsu** 10 months, 3 weeks ago  
D not possible?  
upvoted 1 times



  **awsgeek75** 10 months ago  
ALB is for inbound traffic. D is not possible as it is suggesting to direct OUTBOUND traffic.  
upvoted 4 times



  **Cyberkayu** 10 months, 3 weeks ago  
**Selected Answer: A**  
AWS network firewall is stateful, providing control and visibility to Layer 3-7 network traffic, thus cover the application too  
upvoted 2 times



  **TariqKipkemei** 1 year ago  
**Selected Answer: A**  
Just tried on the console to set up an outbound rule, and URLs cannot be used as a destination. I will opt for A.  
upvoted 2 times



  **Guru4Cloud** 1 year, 2 months ago  
**Selected Answer: C**  
Implement strict inbound security group rules  
Configure an outbound security group rule to allow traffic only to the approved software repository URLs  
The key points:  
  
Highly sensitive EC2 instances in private subnet that can access only approved URLs  
Other internet access must be blocked  
Security groups act as a firewall at the instance level and can control both inbound and outbound traffic.  
upvoted 2 times



  **pentium75** 10 months, 1 week ago  
Security Groups work with CIDR ranges, not URLs.  
upvoted 3 times



  **kelvintoys93** 1 year, 4 months ago  
Isn't private subnet not connectible to internet at all, unless with a NAT gateway?  
upvoted 4 times



  **VeseljkoD** 1 year, 7 months ago  
**Selected Answer: A**  
We can't specify URL in outbound rule of security group. Create free tier AWS account and test it.  
upvoted 2 times

  **Leo301** 1 year, 8 months ago  
**Selected Answer: C**  
CCCCCCCCCCCC  
upvoted 1 times

  **pentium75** 10 months, 1 week ago  
Security Groups with IP ranges, not URLs  
upvoted 2 times

  **Brak** 1 year, 8 months ago  
It can't be C. You cannot use URLs in the outbound rules of a security group.  
upvoted 3 times

  **johnmcclane78** 1 year, 8 months ago  
Option C is the best solution to meet the requirements of this scenario. Implementing strict inbound security group rules that only allow traffic from approved sources can help secure the VPC network that hosts Amazon EC2 instances. Additionally, configuring an outbound rule that allows traffic only to the authorized software repositories on the internet by specifying the URLs will ensure that only approved third-party software repositories can be accessed from the EC2 instances. This solution does not require any additional AWS services and can be implemented using VPC security groups.  
  
Option A is not the best solution as it involves the use of AWS Network Firewall, which may introduce additional operational overhead. While domain list rule groups can be used to block all internet traffic except for the approved third-party software repositories, this solution is more complex than necessary for this scenario.  
upvoted 2 times

  **pentium75** 10 months, 1 week ago  
How do you use a Security Group to allow access to https://server.com/repoa while denying access to https://server.com/repoB ? Security Groups work with IP ranges.

upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

**Selected Answer: C**

In the security group, only allow inbound traffic originating from the VPC. Then only allow outbound traffic with a whitelisted IP address. The question asks about blocking EC2 instances, which is best for security groups since those are at the EC2 instance level. A network firewall is at the VPC level, which is not what the question is asking to protect.

upvoted 1 times

  **pentium75** 10 months, 1 week ago



Security Groups work with IP ranges, not URLs.

upvoted 1 times

  **Theodorz** 1 year, 8 months ago

Is Security Group able to allow a specific URL? According to [https://docs.aws.amazon.com/vpc/latest/userguide/VPC\\_SecurityGroups.html](https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html), I cannot find such description.

upvoted 2 times

  **KZM** 1 year, 8 months ago



I am confused that It seems both options A and C are valid solutions.

upvoted 3 times

  **Mia2009687** 1 year, 3 months ago

I think C is in private subnet. Even with security group, it could not go public to download the software.

upvoted 1 times

  **Zohx** 1 year, 8 months ago

Same here - why is C not a valid option?

upvoted 2 times

  **pentium75** 10 months, 1 week ago


Because you want to filter based on URLs, not IP ranges.

upvoted 2 times

  **Karlos99** 1 year, 8 months ago



And it is easier to do it at the VPC level

upvoted 2 times

  **Karlos99** 1 year, 8 months ago

And it is easier to do it at the level

upvoted 2 times

  **ruqui** 1 year, 5 months ago

C is not valid. Security groups can allow only traffic from specific ports and/or IPs, you can't use an URL. Correct answer is A



upvoted 3 times

  **jennyka76** 1 year, 8 months ago

Answer - A

<https://aws.amazon.com/premiumsupport/knowledge-center/ec2-al1-al2-update-yum-without-internet/>

upvoted 6 times

  **asoli** 1 year, 7 months ago

Although the answer is A, the link you provided here is not related to this question.

The information about "Network Firewall" and how it can help this issue is here:

<https://docs.aws.amazon.com/network-firewall/latest/developerguide/suricata-examples.html#suricata-example-domain-filtering>

(thanks to "@Bhawesh" to provide the link in their answer)

upvoted 4 times

A company is hosting a three-tier ecommerce application in the AWS Cloud. The company hosts the website on Amazon S3 and integrates the website with an API that handles sales requests. The company hosts the API on three Amazon EC2 instances behind an Application Load Balancer (ALB). The API consists of static and dynamic front-end content along with backend workers that process sales requests asynchronously.

The company is expecting a significant and sudden increase in the number of sales requests during events for the launch of new products.

What should a solutions architect recommend to ensure that all the requests are processed successfully?

- A. Add an Amazon CloudFront distribution for the dynamic content. Increase the number of EC2 instances to handle the increase in traffic.
- B. Add an Amazon CloudFront distribution for the static content. Place the EC2 instances in an Auto Scaling group to launch new instances based on network traffic.
- C. Add an Amazon CloudFront distribution for the dynamic content. Add an Amazon ElastiCache instance in front of the ALB to reduce traffic for the API to handle.
- D. Add an Amazon CloudFront distribution for the static content. Add an Amazon Simple Queue Service (Amazon SQS) queue to receive requests from the website for later processing by the EC2 instances.

**Correct Answer:** D

*Community vote distribution*

D (73%)B (27%)

- Steve\_4542636**

Highly Voted

 2 years, 2 months ago

Selected Answer: B

The auto-scaling would increase the rate at which sales requests are "processed", whereas a SQS will ensure messages don't get lost. If you were at a fast food restaurant with a long line with 3 cash registers, would you want more cash registers or longer ropes to handle longer lines? Same concept here.

upvoted 22 times
- AdamVigas** 2 months, 2 weeks ago

they want to process it sucessfully not faster (SQS) is the way...

upvoted 1 times
- Chef\_couincouin** 1 year, 5 months ago

ensure that all the requests are processed successfully? doesn't mean more quickly

upvoted 4 times
- lizzard812** 2 years, 1 month ago

Hell true: I'd rather combine the both options: a SQS + auto-scaled bound to the length of the queue.

upvoted 10 times
- joechen2023** 1 year, 10 months ago

As an architecture, it is not possible to add more backend workers (it is part of the HR and boss's job, not for architecture design the solution). So when the demand surge, the only correct choice is to buffer them using SQS so that workers can take their time to process it successfully

upvoted 2 times
- Abhineet9148232**

Highly Voted

 1 year, 12 months ago

Selected Answer: D

B doesn't fit because Auto Scaling alone does not guarantee that all requests will be processed successfully, which the question clearly asks for.

D ensures that all messages are processed.

upvoted 17 times
- ChhatwaniB**

Most Recent

 3 weeks, 3 days ago

Selected Answer: D

There is no scaling based on network traffic. since the backend worker works asynchronously SQS should be right here

upvoted 1 times
- samadal** 7 months ago

The problem states that the application consists of "static and dynamic front-end content." Static content typically includes cacheable resources such as HTML, CSS, and image files. Therefore, from this statement, one can infer that caching static content using CloudFront would improve performance. In other words, the mention of "static content" in the problem itself leads to the conclusion that CloudFront should be added for static content.

Additionally, the problem mentions "asynchronously processed backend workers." Asynchronous processing is well-suited for services like SQS, which can improve efficiency by handling dynamic requests that do not require immediate processing. The mention of "successfully processing all requests" also suggests that SQS is needed to ensure that all requests are handled properly.

Therefore, the correct answer is D.  
upvoted 3 times

  **Adinas\_** 1 year, 2 months ago

**Selected Answer: B**

Important question to answer D. Can you connect the website with SQS directly? How do you control access to who can put messages to SQS? I have never seen such a situation it has to be at least behind API gateway. So that conclusion brings me to answer B, application also can process async everything without SQS.  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: D**

I chose D because I love SQS! These questions are hammering SQS in every solution as a "protagonist" that saves the day. AC are clearly useless  
B can work but D is better because of SQS being better than EC2 scaling. The other part is that backend workers process the request asynchronously therefore a queue is better.  
upvoted 4 times

  **awsgeek75** 1 year, 4 months ago

**Selected Answer: D**

A and C don't solve anything so ignore them.  
Between B and D, D guarantees the scaling via SQS and order processing. B can also do that but it is not guaranteed that EC2 scaling will work to process the order.  
As usual, I suspect that this "brain dump" may be missing critical wording to differentiate between the options so read carefully in the exam.  
upvoted 4 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: D**

There are two components that we need  
\* Frontend: Hosted on S3, performance can be increased with CloudFront  
\* Backend: There's no reason to process all the orders instantly, so we should decouple the processing from the API which we do with SQS

Thus D, CloudFront + SQS  
upvoted 8 times

  **pentium75** 1 year, 4 months ago

And as others said, B might speed up the processing or reduce the number of lost orders, but we need to make sure that "ALL requests are processed successfully", NOT that "less requests are lost".  
upvoted 3 times

  **Marco\_St** 1 year, 4 months ago

**Selected Answer: D**

I picked B before I read D option. Read the question again, it concerns asynchronous processing of sales requests, Option D seems to align more closely with the requirements. So the requirement is ensuring all requests are processed successfully which means no request would be missed. So D is better option  
upvoted 4 times

  **wsdasdasdqwda** 1 year, 6 months ago

Amazon SQS will make sure that the requests are stored and didn't get lost. After that the workers asynchronously will process the requests. I would go for D  
upvoted 4 times

  **TariqKipkemei** 1 year, 6 months ago

Technically both option B and D would work. But, there's a need to process requests asynchronously, hence decoupling, hence Amazon SQS. I will settle with option D.  
upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

D is correct.  
upvoted 3 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: D**

D is correct.  
upvoted 2 times

  **kruasan** 2 years ago

**Selected Answer: D**

An SQS queue acts as a buffer between the frontend (website) and backend (API). Web requests can dump messages into the queue at a high throughput, then the queue handles delivering those messages to the API at a controlled rate that it can sustain. This prevents the API from being overwhelmed.

upvoted 3 times

  **kruasan** 2 years ago

Options A and B would help by scaling out more instances, however, this may not scale quickly enough and still risks overwhelming the API. Caching parts of the dynamic content (option C) may help but does not provide the buffering mechanism that a queue does.

upvoted 1 times

  **seifshendy99** 2 years ago

**Selected Answer: D**

D make sens

upvoted 2 times

  **kraken21** 2 years, 1 month ago

**Selected Answer: D**

D makes more sense

upvoted 2 times

  **kraken21** 2 years, 1 month ago

There is no clarity on what the asynchronous process is but D makes more sense if we want to process all requests successfully. The way the question is worded it looks like the msgs->SQS>ELB/Ec2. This ensures that the messages are processed but may be delayed as the load increases.

upvoted 1 times

A security audit reveals that Amazon EC2 instances are not being patched regularly. A solutions architect needs to provide a solution that will run regular security scans across a large fleet of EC2 instances. The solution should also patch the EC2 instances on a regular schedule and provide a report of each instance's patch status.

Which solution will meet these requirements?

- A. Set up Amazon Macie to scan the EC2 instances for software vulnerabilities. Set up a cron job on each EC2 instance to patch the instance on a regular schedule.
- B. Turn on Amazon GuardDuty in the account. Configure GuardDuty to scan the EC2 instances for software vulnerabilities. Set up AWS Systems Manager Session Manager to patch the EC2 instances on a regular schedule.
- C. Set up Amazon Detective to scan the EC2 instances for software vulnerabilities. Set up an Amazon EventBridge scheduled rule to patch the EC2 instances on a regular schedule.
- D. Turn on Amazon Inspector in the account. Configure Amazon Inspector to scan the EC2 instances for software vulnerabilities. Set up AWS Systems Manager Patch Manager to patch the EC2 instances on a regular schedule.

Correct Answer: D

Community vote distribution

D (100%)

  **elearningtakai** Highly Voted  1 year, 7 months ago

Selected Answer: D

Amazon Inspector is a security assessment service that automatically assesses applications for vulnerabilities or deviations from best practices. It can be used to scan the EC2 instances for software vulnerabilities. AWS Systems Manager Patch Manager can be used to patch the EC2 instances on a regular schedule. Together, these services can provide a solution that meets the requirements of running regular security scans and patching EC2 instances on a regular schedule. Additionally, Patch Manager can provide a report of each instance's patch status.  
upvoted 10 times

  **LuckyAro** Highly Voted  1 year, 8 months ago

Selected Answer: D

Amazon Inspector is a security assessment service that helps improve the security and compliance of applications deployed on Amazon Web Services (AWS). It automatically assesses applications for vulnerabilities or deviations from best practices. Amazon Inspector can be used to identify security issues and recommend fixes for them. It is an ideal solution for running regular security scans across a large fleet of EC2 instances.

AWS Systems Manager Patch Manager is a service that helps you automate the process of patching Windows and Linux instances. It provides a simple, automated way to patch your instances with the latest security patches and updates. Patch Manager helps you maintain compliance with security policies and regulations by providing detailed reports on the patch status of your instances.  
upvoted 5 times

  **awsgeek75** Most Recent  10 months ago


Selected Answer: D

A handy reference page for such questions is:  
<https://aws.amazon.com/products/security/>  
Amazon Inspector = vulnerability detection = patching  
<https://aws.amazon.com/inspector/>  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

Selected Answer: D

dddddddddd  
upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

Selected Answer: D

Inspector is for EC2 instances and network accessibility of those instances  
<https://portal.tutorialsdojo.com/forums/discussion/difference-between-security-hub-detective-and-inspector/>  
upvoted 2 times

  **TungPham** 1 year, 8 months ago

Selected Answer: D

Amazon Inspector for EC2  
[https://aws.amazon.com/vi/inspector/faqs/?nc1=f\\_ls](https://aws.amazon.com/vi/inspector/faqs/?nc1=f_ls)  
Amazon system manager Patch manager for automates the process of patching managed nodes with both security-related updates and other

types of updates.

<http://webcache.googleusercontent.com/search?q=cache:FbFTc6XKycwJ:https://medium.com/aws-architech/use-case-aws-inspector-vs-guarddduty-3662bf80767a&hl=vi&gl=kr&strip=1&vwsrc=0>

upvoted 3 times

  **jennyka76** 1 year, 8 months ago

answer - D

<https://aws.amazon.com/inspector/faqs/>

upvoted 3 times

  **Neha999** 1 year, 8 months ago

D as AWS Systems Manager Patch Manager can patch the EC2 instances.

upvoted 2 times

A company is planning to store data on Amazon RDS DB instances. The company must encrypt the data at rest.

What should a solutions architect do to meet this requirement?

- A. Create a key in AWS Key Management Service (AWS KMS). Enable encryption for the DB instances.
- B. Create an encryption key. Store the key in AWS Secrets Manager. Use the key to encrypt the DB instances.
- C. Generate a certificate in AWS Certificate Manager (ACM). Enable SSL/TLS on the DB instances by using the certificate.
- D. Generate a certificate in AWS Identity and Access Management (IAM). Enable SSL/TLS on the DB instances by using the certificate.

**Correct Answer: A**

*Community vote distribution*

A (100%)

- awsgeek75** 10 months ago

**Selected Answer: A**

A: Enable encryption  
B: KMS is for storage and doesn't directly integrate to DB without further work  
C and D are for data encryption in transit not at rest  
upvoted 4 times
- awsgeek75** 9 months, 2 weeks ago

Actually, D is total nonsense and no idea what it is saying  
upvoted 1 times
- robpalacios1** 11 months, 2 weeks ago

**Selected Answer: A**

KMS only generates and manages encryption keys. That's it. That's all it does. It's a fundamental service that you as well as other AWS Services (like Secrets Manager) use it to encrypt or decrypt.  
Key Management Service. Secrets Manager is for database connection strings.  
upvoted 3 times  
upvoted 4 times
- antropaws** 1 year, 5 months ago

OK, but why not B???  
upvoted 1 times
- aaroncelestin** 1 year, 2 months ago

KMS only generates and manages encryption keys. That's it. That's all it does. It's a fundamental service that you as well as other AWS Services (like Secrets Manager) use it to encrypt or decrypt.  
  
Secrets Manager stores actual secrets like passwords, pass phrases, and anything else you want encrypted. SM uses KMS to encrypt its secrets, it would be circular to get an encryption key from KMS to use SM to encrypt the encryption key.  
upvoted 5 times
- SkyZeroZx** 1 year, 6 months ago

**Selected Answer: A**

ANSWER - A  
upvoted 1 times
- datz** 1 year, 7 months ago

**Selected Answer: A**

A for sure  
upvoted 2 times
- PRASAD180** 1 year, 8 months ago

A is 100% Crt  
upvoted 2 times
- Steve\_4542636** 1 year, 8 months ago

**Selected Answer: A**

Key Management Service. Secrets Manager is for database connection strings.  
upvoted 4 times
- LuckyAro** 1 year, 8 months ago

**Selected Answer: A**

A is the correct solution to meet the requirement of encrypting the data at rest.

To encrypt data at rest in Amazon RDS, you can use the encryption feature of Amazon RDS, which uses AWS Key Management Service (AWS KMS). With this feature, Amazon RDS encrypts each database instance with a unique key. This key is stored securely by AWS KMS. You can manage your own keys or use the default AWS-managed keys. When you enable encryption for a DB instance, Amazon RDS encrypts the underlying storage, including the automated backups, read replicas, and snapshots.

upvoted 4 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: A**

AWS Key Management Service (KMS) is used to manage the keys used to encrypt and decrypt the data.

upvoted 2 times

  **pbpally** 1 year, 8 months ago

**Selected Answer: A**

Option A

upvoted 2 times

  **NolaH0la** 1 year, 8 months ago

A. Create a key in AWS Key Management Service (AWS KMS). Enable encryption for the DB instances is the correct answer to encrypt the data at rest in Amazon RDS DB instances.

Amazon RDS provides multiple options for encrypting data at rest. AWS Key Management Service (KMS) is used to manage the keys used to encrypt and decrypt the data. Therefore, a solution architect should create a key in AWS KMS and enable encryption for the DB instances to encrypt the data at rest.

upvoted 2 times

  **jennyka76** 1 year, 8 months ago

ANSWER - A

<https://docs.aws.amazon.com/whitepapers/latest/efs-encrypted-file-systems/managing-keys.html>

upvoted 2 times

  **Bhawesh** 1 year, 8 months ago

**Selected Answer: A**

A. Create a key in AWS Key Management Service (AWS KMS). Enable encryption for the DB instances.

<https://www.examttopics.com/discussions/amazon/view/80753-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 3 times

A company must migrate 20 TB of data from a data center to the AWS Cloud within 30 days. The company’s network bandwidth is limited to 15 Mbps and cannot exceed 70% utilization.

What should a solutions architect do to meet these requirements?

- A. Use AWS Snowball.
- B. Use AWS DataSync.
- C. Use a secure VPN connection.
- D. Use Amazon S3 Transfer Acceleration.

**Correct Answer: A**

*Community vote distribution*

A (93%)

7%

  **kruasan** Highly Voted 1 year, 6 months ago

**Selected Answer: A**

Don't mix up between Mbps and Mbs.  
The proper calculation is:

$10 \text{ MB/s} \times 86,400 \text{ seconds per day} \times 30 \text{ days} / 8 = 3,402,000 \text{ MB}$  or approximately 3.4 TB  
upvoted 14 times

  **awsgeek75** Highly Voted 10 months ago

**Selected Answer: A**

Honestly, the company has bigger problem with that slow connection :)  
30 days is the first clue so you can get snowball shipped and sent back (5 days each way)  
upvoted 7 times

  **cabta** Most Recent 10 months, 1 week ago

**Selected Answer: A**

aws snowball은 대용량 데이터 이전하기 위한 것 입니다.  
upvoted 1 times

  **wsdasdasdqwda** 1 year ago

$(15/8) = 1.875 \text{ MB/s}$   
 $1.875 \text{ MB/s} \times 0.7 = 1.3125 \text{ (70\% NW utilization) MB/s}$   
 $1.3125 \text{ MB/s} \times 3600 = 4725 \text{ MB (MB per 1 hour)}$   
 $4725 \times 24 = 113400 \text{ MB per 1 full day (24h)}$   
 $113400 \times 30 = 3402000 \text{ MB for 30 days}$   
 $3402000 / 1024 = 3322.265625 \text{ GB for 30 days}$   
 $3322.265625 / 1024 \sim 3.24 \text{ TB for 30 days} \Rightarrow \text{not enough for NW} \Rightarrow \text{Snowball which is A}$   
upvoted 4 times

  **TariqKipkemei** 1 year ago

**Selected Answer: A**

I wont try to think to much about it, AWS Snowball was designed for this  
upvoted 4 times

  **Guru4Cloud** 1 year, 1 month ago

**Selected Answer: A**

° 15 Mbps bandwidth with 70% max utilization limits the effective bandwidth to 10.5 Mbps or 1.31 MB/s.  
° 20 TB of data at 1.31 MB/s would take approximately 193 days to transfer over the network. ° This far exceeds the 30 day requirement.  
° AWS Snowball provides a physical storage device that can be shipped to the data center. Up to 80 TB can be loaded onto a Snowball device and shipped back to AWS.  
This allows the 20 TB of data to be transferred much faster by shipping rather than over the limited network bandwidth.  
° Snowball uses tamper-resistant enclosures and 256-bit encryption to keep the data secure during transit.  
° The data can be imported into Amazon S3 or Amazon Glacier once the Snowball is received by AWS.  
upvoted 5 times

  **UnluckyDucky** 1 year, 7 months ago

**Selected Answer: B**



$10 \text{ MB/s} \times 86,400 \text{ seconds per day} \times 30 \text{ days} = 25,920,000 \text{ MB}$  or approximately 25.2 TB

That's how much you can transfer with a 10 Mbps link (roughly 70% of the 15 Mbps connection).

With a consistent connection of 8~ Mbps, and 30 days, you can upload 20 TB of data.



My math says B, my brain wants to go with A. Take your pick.

upvoted 3 times

  **Zox42** 1 year, 7 months ago



15 Mbps \* 0.7 = 1.3125 MB/s and 1.3125 \* 86,400 \* 30 = 3.402.000 MB  
Answer A is correct.

upvoted 3 times

  **Zox42** 1 year, 7 months ago



3,402,000

upvoted 2 times

  **hozy\_** 1 year, 3 months ago

How can 15 \* 0.7 be 1.3125 LMAO

upvoted 1 times

  **hozy\_** 1 year, 3 months ago

OMG it was Mbps! Not MBps. You are right! awesome!!!

upvoted 3 times

  **Bilalazure** 1 year, 8 months ago

**Selected Answer: A**

Aws snowball

upvoted 3 times

  **PRASAD180** 1 year, 8 months ago

A is 100% Crt


upvoted 2 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: A**

AWS Snowball

upvoted 2 times

  **pbpally** 1 year, 8 months ago

**Selected Answer: A**

Option a

upvoted 2 times

  **jennyka76** 1 year, 8 months ago

ANSWER - A

<https://docs.aws.amazon.com/snowball/latest/ug/whatissnowball.html>

upvoted 2 times

  **AWSSHA1** 1 year, 8 months ago

**Selected Answer: A**

option A

upvoted 4 times

A company needs to provide its employees with secure access to confidential and sensitive files. The company wants to ensure that the files can be accessed only by authorized users. The files must be downloaded securely to the employees’ devices.

The files are stored in an on-premises Windows file server. However, due to an increase in remote usage, the file server is running out of capacity.

Which solution will meet these requirements?



- A. Migrate the file server to an Amazon EC2 instance in a public subnet. Configure the security group to limit inbound traffic to the employees’ IP addresses.
- B. Migrate the files to an Amazon FSx for Windows File Server file system. Integrate the Amazon FSx file system with the on-premises Active Directory. Configure AWS Client VPN.
- C. Migrate the files to Amazon S3, and create a private VPC endpoint. Create a signed URL to allow download.
- D. Migrate the files to Amazon S3, and create a public VPC endpoint. Allow employees to sign on with AWS IAM Identity Center (AWS Single Sign-On).


**Correct Answer: B**

*Community vote distribution*

B (93%)

7%



-   **elearningtakai**


Highly Voted 

 2 years, 1 month ago

Selected Answer: B

This solution addresses the need for secure access to confidential and sensitive files, as well as the increase in remote usage. Migrating the files to Amazon FSx for Windows File Server provides a scalable, fully managed file storage solution in the AWS Cloud that is accessible from on-premises and cloud environments. Integration with the on-premises Active Directory allows for a consistent user experience and centralized access control. AWS Client VPN provides a secure and managed VPN solution that can be used by employees to access the files securely.



upvoted 9 times
-   **pentium75**


Highly Voted 

 1 year, 4 months ago

Selected Answer: B

C has "signed URL", everyone who has the URL could download. Plus, only B ensure the "must be downloaded securely" part by using VPN.



upvoted 5 times
-   **NayeraB**

Most Recent 

 1 year, 2 months ago



Selected Answer: B

My money is on B, but it's still not mentioned that the customer used an on-prem Active Directory.



upvoted 3 times
-   **TariqKipkemei** 1 year, 6 months ago

Selected Answer: B

Windows file server = Amazon FSx for Windows File Server file system  
Files can be accessed only by authorized users = On-premises Active Directory



upvoted 4 times
-   **JA2018** 5 months, 2 weeks ago

"Data must be downloaded securely to users' devices" => AWS Client VPN



upvoted 1 times
-   **BrijMohan08** 1 year, 8 months ago

Selected Answer: C

Remember: The file server is running out of capacity.

upvoted 2 times
-   **awsgeek75** 1 year, 3 months ago


But then how do you download the files to user's machine in a secure way?

upvoted 2 times
-   **pentium75** 1 year, 4 months ago

That's why we're using FSX for Windows File Server in AWS.

"Signed URL to allow download" would allow everyone who has the URL to download the files, but we must "ensure that the files can be accessed only by authorized users". Plus, the "private VPC endpoint" is not really of use here, it's still S3 and the users are not in AWS.

upvoted 4 times

  **SkyZeroZx** 1 year, 12 months ago

**Selected Answer: B**

B is the correct answer

upvoted 2 times

  **LuckyAro** 2 years, 2 months ago

**Selected Answer: B**

B is the best solution for the given requirements. It provides a secure way for employees to access confidential and sensitive files from anywhere using AWS Client VPN. The Amazon FSx for Windows File Server file system is designed to provide native support for Windows file system features such as NTFS permissions, Active Directory integration, and Distributed File System (DFS). This means that the company can continue to use their on-premises Active Directory to manage user access to files.

upvoted 4 times

  **Bilalazure** 2 years, 2 months ago

B is the correct answer

upvoted 2 times

  **jennyka76** 2 years, 2 months ago

Answer - B

1- <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html>

2- <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/managing-storage-capacity.html>

upvoted 2 times

  **Neha999** 2 years, 2 months ago

B

Amazon FSx for Windows File Server file system

upvoted 3 times

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 runs. 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.

**Correct Answer: C**

*Community vote distribution*

C (100%)

🗲️ 👤 **TariqKipkemei** Highly Voted 👍 6 months, 3 weeks ago

**Selected Answer: C**

'On the first day of every month at midnight' = Scheduled scaling policy  
upvoted 5 times

🗲️ 👤 **elearningtakai** Most Recent ⌚ 1 year, 1 month ago

**Selected Answer: C**

By configuring a scheduled scaling policy, the EC2 Auto Scaling group can proactively launch additional EC2 instances before the CPU utilization peaks to 100%. This will ensure that the application can handle the workload during the month-end financial calculation batch, and avoid any disruption or downtime.

Configuring a simple scaling policy based on CPU utilization or adding Amazon CloudFront distribution or Amazon ElastiCache will not directly address the issue of handling the monthly peak workload.  
upvoted 4 times

🗲️ 👤 **Steve\_4542636** 1 year, 2 months ago

**Selected Answer: C**

If the scaling were based on CPU or memory, it requires a certain amount of time above that threshold, 5 minutes for example. That would mean the CPU would be at 100% for five minutes.  
upvoted 4 times

🗲️ 👤 **LuckyAro** 1 year, 2 months ago

**Selected Answer: C**

C: Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule is the best option because it allows for the proactive scaling of the EC2 instances before the monthly batch run begins. This will ensure that the application is able to handle the increased workload without experiencing downtime. The scheduled scaling policy can be configured to increase the number of instances in the Auto Scaling group a few hours before the batch run and then decrease the number of instances after the batch run is complete. This will ensure that the resources are available when needed and not wasted when not needed.

The most appropriate solution to handle the increased workload during the monthly batch run and avoid downtime would be to configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.  
upvoted 3 times

🗲️ 👤 **LuckyAro** 1 year, 2 months ago

Scheduled scaling policies allow you to schedule EC2 instance scaling events in advance based on a specified time and date. You can use this feature to plan for anticipated traffic spikes or seasonal changes in demand. By setting up scheduled scaling policies, you can ensure that you have the right number of instances running at the right time, thereby optimizing performance and reducing costs.

To set up a scheduled scaling policy in EC2 Auto Scaling, you need to specify the following:

Start time and date: The date and time when the scaling event should begin.

Desired capacity: The number of instances that you want to have running after the scaling event.

Recurrence: The frequency with which the scaling event should occur. This can be a one-time event or a recurring event, such as daily or weekly.  
upvoted 2 times

🗲️ 👤 **bdp123** 1 year, 2 months ago

**Selected Answer: C**

C is the correct answer as traffic spike is known  
upvoted 2 times

  **jennyka76** 1 year, 2 months ago

ANSWER - C

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-scheduled-scaling.html>

upvoted 3 times

  **Neha999** 1 year, 2 months ago

C as the schedule of traffic spike is known beforehand.

upvoted 2 times

A company wants to give a customer the ability to use on-premises Microsoft Active Directory to download files that are stored in Amazon S3. The customer’s application uses an SFTP client to download the files.



Which solution will meet these requirements with the LEAST operational overhead and no changes to the customer’s application?


- A. Set up AWS Transfer Family with SFTP for Amazon S3. Configure integrated Active Directory authentication.
- B. Set up AWS Database Migration Service (AWS DMS) to synchronize the on-premises client with Amazon S3. Configure integrated Active Directory authentication.
- C. Set up AWS DataSync to synchronize between the on-premises location and the S3 location by using AWS IAM Identity Center (AWS Single Sign-On).
- D. Set up a Windows Amazon EC2 instance with SFTP to connect the on-premises client with Amazon S3. Integrate AWS Identity and Access Management (IAM).

**Correct Answer: A**

*Community vote distribution*

A (100%)



-   **Steve\_4542636**


Highly Voted 

 2 years, 2 months ago

Selected Answer: A



SFTP, FTP - think "Transfer" during test time

upvoted 17 times
-   **wsdasdasdqwdaw**

Most Recent 

 1 year, 6 months ago

LEAST operational overhead => A, D is much more operational overhead



upvoted 3 times
-   **JA2018** 5 months, 2 weeks ago

1. AWS Transfer Family: Fully managed service that allows customers to transfer files over SFTP, FTPS, and FTP directly into and out of Amazon S3.

2. Eliminates the need to manage any infrastructure for file transfer, which reduces operational overhead.



3. You can also configure the service to use an existing Active Directory for authentication, =>>> no changes need to be made to the customer's application.

4. Rule of thumbs for AWS exams: Opt for AWS fully managed services. :p :p :p



upvoted 1 times
-   **TariqKipkemei** 1 year, 6 months ago

Selected Answer: A



SFTP, No changes to the customer’s application? = AWS Transfer Family

upvoted 2 times
-   **Guru4Cloud** 1 year, 8 months ago

Transfer family is used for SFTP



upvoted 2 times
-   **live\_reply\_developers** 1 year, 9 months ago

SFTP -> transfer family

upvoted 2 times
-   **antropaws** 1 year, 11 months ago



Selected Answer: A

A no doubt. Why the system gives B as the correct answer?

upvoted 2 times
-   **lht** 2 years ago

Selected Answer: A

just A

upvoted 2 times
-   **LuckyAro** 2 years, 2 months ago

**Selected Answer: A**

AWS Transfer Family  
upvoted 3 times

🗉 👤 **LuckyAro** 2 years, 2 months ago

AWS Transfer Family is a fully managed service that allows customers to transfer files over SFTP, FTPS, and FTP directly into and out of Amazon S3. It eliminates the need to manage any infrastructure for file transfer, which reduces operational overhead. Additionally, the service can be configured to use an existing Active Directory for authentication, which means that no changes need to be made to the customer's application.  
upvoted 3 times

🗉 👤 **bdp123** 2 years, 2 months ago

**Selected Answer: A**

Transfer family is used for SFTP  
upvoted 2 times

🗉 👤 **TungPham** 2 years, 2 months ago

**Selected Answer: A**

using AWS Batch to LEAST operational overhead  
and have SFTP to no changes to the customer's application

<https://aws.amazon.com/vi/blogs/architecture/managed-file-transfer-using-aws-transfer-family-and-amazon-s3/>  
upvoted 3 times

🗉 👤 **Bhawesh** 2 years, 2 months ago

**Selected Answer: A**

A. Set up AWS Transfer Family with SFTP for Amazon S3. Configure integrated Active Directory authentication.

<https://docs.aws.amazon.com/transfer/latest/userguide/directory-services-users.html>  
upvoted 4 times

A company is experiencing sudden increases in demand. The company needs to provision large Amazon EC2 instances from an Amazon Machine Image (AMI). The instances will run in an Auto Scaling group. The company needs a solution that provides minimum initialization latency to meet the demand.

Which solution meets these requirements?

- A. Use the `aws ec2 register-image` command to create an AMI from a snapshot. Use AWS Step Functions to replace the AMI in the Auto Scaling group.
- B. Enable Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot. Provision an AMI by using the snapshot. Replace the AMI in the Auto Scaling group with the new AMI.
- C. Enable AMI creation and define lifecycle rules in Amazon Data Lifecycle Manager (Amazon DLM). Create an AWS Lambda function that modifies the AMI in the Auto Scaling group.
- D. Use Amazon EventBridge to invoke AWS Backup lifecycle policies that provision AMIs. Configure Auto Scaling group capacity limits as an event source in EventBridge.

**Correct Answer:** B

*Community vote distribution*

B (94%)

4%

- danielklein09** Highly Voted

1 year, 11 months ago

readed the question 5 times, didn't understood a thing :(

upvoted 63 times
- elmyth**

7 months, 3 weeks ago

Me too((( terrible question

upvoted 2 times
- Guru4Cloud**

1 year, 8 months ago

Me too

upvoted 4 times
- lostmagnet001**

1 year, 2 months ago

the same here!

upvoted 1 times
- bdp123** Highly Voted

2 years, 2 months ago

**Selected Answer: B**

Enabling Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot allows you to quickly create a new Amazon Machine Image (AMI) from a snapshot, which can help reduce the initialization latency when provisioning new instances. Once the AMI is provisioned, you can replace the AMI in the Auto Scaling group with the new AMI. This will ensure that new instances are launched from the updated AMI and are able to meet the increased demand quickly.

upvoted 14 times

**LeonSauveterre** Most Recent

5 months, 1 week ago

**Selected Answer: B**

The question focuses on reducing initialization latency when launching large EC2 instances in an Auto Scaling group during sudden demand spikes.

So B with fast snapshot is obviously the choice. But without the choices listed there, I would simply integrate a Warm Pool for the Auto Scaling group with pre-initialized instances based on an optimized AMI :)

upvoted 3 times

**awsgeek75**

1 year, 4 months ago

**Selected Answer: B**

The question wording is pretty weird but the only thing of value is latency during initialisation which makes B the correct option.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-fast-snapshot-restore.html>

A only helps with creating the AMI  
C and D will probably work (ambiguous language) but won't handle initialising latency issues.

upvoted 6 times

**farnamjam**

1 year, 4 months ago

**Selected Answer: B**

Fast Snapshot Restore (FSR)

- Force full initialization of snapshot to have no latency on the first use

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: B**

"Fast snapshot restore" = pre-warmed snapshot  
AMI from such a snapshot is pre-warmed AMI

upvoted 4 times

  **master9** 1 year, 4 months ago

**Selected Answer: D**

Amazon Data Lifecycle Manager (DLM) is a feature of Amazon EBS that automates the creation, retention, and deletion of snapshots, which are used to back up your Amazon EBS volumes. With DLM, you can protect your data by implementing a backup strategy that aligns with your business requirements.

You can create lifecycle policies to automate snapshot management. Each policy includes a schedule of when to create snapshots, a retention rule with a defined period to retain each snapshot, and a set of Amazon EBS volumes to assign to the policy.

This service helps simplify the management of your backups, ensure compliance, and reduce costs.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

We're not asked to "simplify the management of our backups, ensure compliance, and reduce costs", we're asked to "provide minimum initialization latency" for an auto-scaling group.

upvoted 2 times

  **master9** 1 year, 4 months ago

Sorry, its "C" and not "D"

upvoted 1 times

  **Nisarg2121** 1 year, 4 months ago

**Selected Answer: B**

b is correct

upvoted 1 times

  **meowruki** 1 year, 5 months ago

B. Enable Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot. Provision an AMI by using the snapshot. Replace the AMI in the Auto Scaling group with the new AMI.

Here's the reasoning:

Amazon EBS Fast Snapshot Restore: This feature allows you to quickly create new EBS volumes (and subsequently AMIs) from snapshots. Fast Snapshot Restore optimizes the initialization process by pre-warming the snapshots, reducing the time it takes to create volumes from those snapshots.

Provision an AMI using the snapshot: By using fast snapshot restore, you can efficiently provision an AMI from the pre-warmed snapshot, minimizing the initialization latency.

Replace the AMI in the Auto Scaling group: This allows you to update the instances in the Auto Scaling group with the new AMI efficiently, ensuring that the new instances are launched with minimal delay.

upvoted 2 times

  **meowruki** 1 year, 5 months ago

Option A (Use aws ec2 register-image command and AWS Step Functions): While this approach can be used to automate the creation of an AMI and update the Auto Scaling group, it may not offer the same level of optimization for initialization latency as Amazon EBS fast snapshot restore.

Option C (Enable AMI creation and define lifecycle rules in Amazon Data Lifecycle Manager, create a Lambda function): While Amazon DLM can help manage the lifecycle of your AMIs, it might not provide the same level of speed and responsiveness needed for sudden increases in demand.

Option D (Use Amazon EventBridge and AWS Backup): AWS Backup is primarily designed for backup and recovery, and it might not be as optimized for quickly provisioning instances in response to sudden demand spikes. EventBridge can be used for event-driven architectures, but in this context, it might introduce unnecessary complexity.

upvoted 2 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: B**

Enable Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot. Provision an AMI by using the snapshot. Replace the AMI in the Auto Scaling group with the new AMI

upvoted 2 times

  **kambarami** 1 year, 7 months ago

Pleaw3 reword 5he question. Can not understand a thing!  
upvoted 1 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Enable EBS fast snapshot restore on a snapshot  
Create an AMI from the snapshot  
Replace the AMI used by the Auto Scaling group with this new AMI

The key points:

- ° Need to launch large EC2 instances quickly from an AMI in an Auto Scaling group
- ° Looking to minimize instance initialization latency

upvoted 3 times

  **antropaws** 1 year, 11 months ago



**Selected Answer: B**

B most def  
upvoted 2 times

  **elearningtakai** 2 years, 1 month ago

**Selected Answer: B**

B: "EBS fast snapshot restore": minimizes initialization latency. This is a good choice.  
upvoted 3 times

  **Zox42** 2 years, 1 month ago

**Selected Answer: B**

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-fast-snapshot-restore.html>  
upvoted 3 times

  **geekgirl22** 2 years, 2 months ago

Keyword, minimize initilization latency == snapshot. A and B have snapshots in them, but B is the one that makes sense.  
C has DLP that can create machines from AMI, but that does not talk about latency and snapshots.  
upvoted 4 times

  **LuckyAro** 2 years, 2 months ago

**Selected Answer: B**

Enabling Amazon Elastic Block Store (Amazon EBS) fast snapshot restore on a snapshot allows for rapid restoration of EBS volumes from snapshots. This reduces the time required to create an AMI from a snapshot, which is useful for quickly provisioning large Amazon EC2 instances.

Provisioning an AMI by using the fast snapshot restore feature is a fast and efficient way to create an AMI. Once the AMI is created, it can be replaced in the Auto Scaling group without any downtime or disruption to running instances.  
upvoted 2 times

A company hosts a multi-tier web application that uses an Amazon Aurora MySQL DB cluster for storage. The application tier is hosted on Amazon EC2 instances. The company's IT security guidelines mandate that the database credentials be encrypted and rotated every 14 days.

What should a solutions architect do to meet this requirement with the LEAST operational effort?

- A. Create a new AWS Key Management Service (AWS KMS) encryption key. Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Associate the secret with the Aurora DB cluster. Configure a custom rotation period of 14 days.
- B. Create two parameters in AWS Systems Manager Parameter Store: one for the user name as a string parameter and one that uses the SecureString type for the password. Select AWS Key Management Service (AWS KMS) encryption for the password parameter, and load these parameters in the application tier. Implement an AWS Lambda function that rotates the password every 14 days.
- C. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system in all EC2 instances of the application tier. Restrict the access to the file on the file system so that the application can read the file and that only super users can modify the file. Implement an AWS Lambda function that rotates the key in Aurora every 14 days and writes new credentials into the file.
- D. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon S3 bucket that the application uses to load the credentials. Download the file to the application regularly to ensure that the correct credentials are used. Implement an AWS Lambda function that rotates the Aurora credentials every 14 days and uploads these credentials to the file in the S3 bucket.

Correct Answer: A

Community vote distribution

A (100%)

  **elearningtakai** Highly Voted 2 years, 1 month ago

Selected Answer: A

AWS Secrets Manager allows you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. With this service, you can automate the rotation of secrets, such as database credentials, on a schedule that you choose. The solution allows you to create a new secret with the appropriate credentials and associate it with the Aurora DB cluster. You can then configure a custom rotation period of 14 days to ensure that the credentials are automatically rotated every two weeks, as required by the IT security guidelines. This approach requires the least amount of operational effort as it allows you to manage secrets centrally without modifying your application code or infrastructure.

upvoted 7 times

  **jennyka76** Highly Voted 2 years, 2 months ago

Answer is A

To implement password rotation lifecycles, use AWS Secrets Manager. You can rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle using Secrets Manager.

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types-oracle/>

upvoted 5 times

  **MJ45** Most Recent 4 months, 1 week ago

Selected Answer: A

Most questions related to rotating credentials - AWS Secrets Manager  
AWS System Parameter Store CANNOT rotate credentials.

upvoted 1 times

  **LeonSauveterre** 5 months, 1 week ago

Selected Answer: A

Well... I would simply rule out all the options with "Implement an AWS Lambda function" given that we need a solution with least operational effort. Why manually compose codes to rotate when you can achieve the same goal automatically?

upvoted 1 times

  **TariqKipkemei** 1 year, 6 months ago

Selected Answer: A

Create a new AWS Key Management Service (AWS KMS) encryption key. Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Associate the secret with the Aurora DB cluster. Configure a custom rotation period of 14 days

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

Selected Answer: A

Use AWS Secrets Manager to store the Aurora credentials as a secret  
Encrypt the secret with a KMS key  
Configure 14 day automatic rotation for the secret

Associate the secret with the Aurora DB cluster  
The key points:

Aurora MySQL credentials must be encrypted and rotated every 14 days  
Want to minimize operational effort  
upvoted 3 times

🗲️ 👤 **elearningtakai** 2 years, 1 month ago

**Selected Answer: A**

A: AWS Secrets Manager. Simply this supported rotate feature, and secure to store credentials instead of EFS or S3.  
upvoted 2 times

🗲️ 👤 **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: A**

Voting A  
upvoted 2 times

🗲️ 👤 **LuckyAro** 2 years, 2 months ago

**Selected Answer: A**

A proposes to create a new AWS KMS encryption key and use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Then, the secret will be associated with the Aurora DB cluster, and a custom rotation period of 14 days will be configured. AWS Secrets Manager will automate the process of rotating the database credentials, which will reduce the operational effort required to meet the IT security guidelines.  
upvoted 2 times

🗲️ 👤 **Neha999** 2 years, 2 months ago

A  
<https://www.examttopics.com/discussions/amazon/view/59985-exam-aws-certified-solutions-architect-associate-saa-c02/>  
upvoted 2 times

A company has deployed a web application on AWS. The company hosts the backend database on Amazon RDS for MySQL with a primary DB instance and five read replicas to support scaling needs. The read replicas must lag no more than 1 second behind the primary DB instance. The database routinely runs scheduled stored procedures.

As traffic on the website increases, the replicas experience additional lag during periods of peak load. A solutions architect must reduce the replication lag as much as possible. The solutions architect must minimize changes to the application code and must minimize ongoing operational overhead.

Which solution will meet these requirements?



- A. Migrate the database to Amazon Aurora MySQL. Replace the read replicas with Aurora Replicas, and configure Aurora Auto Scaling. Replace the stored procedures with Aurora MySQL native functions.
- B. Deploy an Amazon ElastiCache for Redis cluster in front of the database. Modify the application to check the cache before the application queries the database. Replace the stored procedures with AWS Lambda functions.
- C. Migrate the database to a MySQL database that runs on Amazon EC2 instances. Choose large, compute optimized EC2 instances for all replica nodes. Maintain the stored procedures on the EC2 instances.
- D. Migrate the database to Amazon DynamoDB. Provision a large number of read capacity units (RCUs) to support the required throughput, and configure on-demand capacity scaling. Replace the stored procedures with DynamoDB streams.



Correct Answer: A



Community vote distribution



A (90%)



10%

  **fkie4** Highly Voted 2 years, 1 month ago  
i hate this kind of question  
upvoted 62 times



  **asoli** Highly Voted 2 years, 1 month ago  
**Selected Answer: A**  
Using Cache required huge changes in the application. Several things need to change to use cache in front of the DB in the application. So, option B is not correct.  
Aurora will help to reduce replication lag for read replica  
upvoted 13 times

  **sheilawu** Most Recent 10 months, 4 weeks ago  
**Selected Answer: A**  
You need to read the question carefully.  
The solutions architect must minimize changes to the application code = therefore A  
If this question without this statement, B will be a better choice.  
upvoted 4 times

  **JackyCCK** 1 year, 1 month ago  
minimize ongoing operational overhead = Not B  
Using ElastiCache require app change  
upvoted 4 times

  **awsgeek75** 1 year, 4 months ago  
**Selected Answer: A**  
AWS Aurora and Native Functions are least application changes while providing better performance and minimum latency.  
<https://aws.amazon.com/rds/aurora/faqs/>

B, C, D require lots of changes to the application so relatively speaking A is least code change and least maintenance/operational overhead.  
upvoted 9 times

  **JA2018** 5 months, 2 weeks ago  
from the same URL:

Is Amazon Aurora MySQL compatible?  
Amazon Aurora is drop-in compatible with existing MySQL open-source databases and adds support for new releases regularly. This means you can easily migrate MySQL databases to and from Aurora using standard import/export tools or snapshots. It also means that most of the code, applications, drivers, and tools you already use with MySQL databases today can be used with Aurora with little or no change. This makes it easy to move applications between the two engines.

You can see the current Amazon Aurora MySQL release compatibility information in the documentation.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: A**

A: Minimal changes to the application code, < 1 second lag  
B: Does not address the replication lag issue at all, requires code changes and adds overhead  
C: Moving from managed RDS to self-managed database on EC2 is ADDING, not minimizing, overhead, PLUS it does not address the replication lag issue  
D: DynamoDB is a NoSQL DB, would require MASSIVE changes to application code and probably even application logic

upvoted 5 times

  **Murtadhaceit** 1 year, 4 months ago

**Selected Answer: A**

imho, B is not valid because it involves extra coding and the question specifically mentions no more coding. Therefore, replacing the current db with another one is not considered as more coding.

upvoted 3 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: A**

Migrate the database to Amazon Aurora MySQL. Replace the read replicas with Aurora Replicas, and configure Aurora Auto Scaling. Replace the stored procedures with Aurora MySQL native functions

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

Migrate the RDS MySQL database to Amazon Aurora MySQL  
Use Aurora Replicas for read scaling instead of RDS read replicas  
Configure Aurora Auto Scaling to handle load spikes  
Replace stored procedures with Aurora MySQL native functions

upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: A**

First, ElastiCache involves heavy change on application code. The question mentioned that "he solutions architect must minimize changes to the application code". Therefore B is not suitable and A is more appropriate for the question requirement.

upvoted 3 times

  **aaroncelestin** 1 year, 8 months ago

... but migrating their ENTIRE prod database and its replicas to a new platform is not a heavy change?



upvoted 3 times

  **KMohsoe** 1 year, 11 months ago

**Selected Answer: B**

Why not B? Please explain to me.

upvoted 2 times

  **Terion** 1 year, 7 months ago

It wouldn't have the most up to date info since it must no lag in relation to the main DB

upvoted 2 times

  **pentium75** 1 year, 4 months ago

How would adding a cache "reduce the replication lag" between the primary instance and the read replicas? Plus, it would require "changes to the application code" that we want to avoid. The "AWS Lambda functions" would create "ongoing operational overhead" that we're also asked to avoid.



upvoted 2 times

  **kaushald** 2 years, 1 month ago

Option A is the most appropriate solution for reducing replication lag without significant changes to the application code and minimizing ongoing operational overhead. Migrating the database to Amazon Aurora MySQL allows for improved replication performance and higher scalability compared to Amazon RDS for MySQL. Aurora Replicas provide faster replication, reducing the replication lag, and Aurora Auto Scaling ensures that there are enough Aurora Replicas to handle the incoming traffic. Additionally, Aurora MySQL native functions can replace the stored procedures, reducing the load on the database and improving performance.

Option B is not the best solution since adding an ElastiCache for Redis cluster does not address the replication lag issue, and the cache may not have the most up-to-date information. Additionally, replacing the stored procedures with AWS Lambda functions adds additional complexity and may not improve performance.

upvoted 5 times

  **njufi** 1 year, 1 month ago

I agree with your explanation. Additionally, considering the requirement that "the read replicas must lag no more than 1 second behind the primary DB instance," it's crucial to ensure that ElastiCache for Redis also maintains this tight synchronization window. This implies that the main RDS instance would need to synchronize an additional database, potentially exacerbating lag during peak times rather than alleviating it.

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 1 month ago

**Selected Answer: B**

a,b are confusing me..  
i would like to go with b..  
upvoted 1 times

🗨️ 👤 **bangfire** 2 years, 1 month ago

Option B is incorrect because it suggests using ElastiCache for Redis as a caching layer in front of the database, but this would not necessarily reduce the replication lag on the read replicas. Additionally, it suggests replacing the stored procedures with AWS Lambda functions, which may require significant changes to the application code.  
upvoted 6 times

🗨️ 👤 **lizzard812** 2 years, 1 month ago

Yes and moreover Redis requires app refactoring which is a solid operational overhead  
upvoted 2 times

🗨️ 👤 **Nel8** 2 years, 2 months ago

**Selected Answer: B**

By using ElastiCache you avoid a lot of common issues you might encounter. ElastiCache is a database caching solution. ElastiCache Redis per se, supports failover and Multi-AZ. And Most of all, ElastiCache is well suited to place in front of RDS.

Migrating a database such as option A, requires operational overhead.  
upvoted 2 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

Database migration is one-time work, NOT "operational overhead". Plus, RDS for MySQL to Aurora with MySQL compatibility is not a big deal, and "minimizes changes to the application code" as requested.  
upvoted 2 times

🗨️ 👤 **bdp123** 2 years, 2 months ago

**Selected Answer: A**

Aurora can have up to 15 read replicas - much faster than RDS  
<https://aws.amazon.com/rds/aurora/>  
upvoted 5 times

🗨️ 👤 **ChrisG1454** 2 years, 2 months ago

" As a result, all Aurora Replicas return the same data for query results with minimal replica lag. This lag is usually much less than 100 milliseconds after the primary instance has written an update "

Reference:  
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>  
upvoted 3 times

🗨️ 👤 **ChrisG1454** 2 years, 1 month ago

You can invoke an Amazon Lambda function from an Amazon Aurora MySQL-Compatible Edition DB cluster with the "native function"....  
  
[https://docs.amazonaws.cn/en\\_us/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html](https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html)  
upvoted 1 times

🗨️ 👤 **jennyka76** 2 years, 2 months ago

Answer - A  
[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_PostgreSQL.Replication.ReadReplicas.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PostgreSQL.Replication.ReadReplicas.html)

-----  
You can scale reads for your Amazon RDS for PostgreSQL DB instance by adding read replicas to the instance. As with other Amazon RDS database engines, RDS for PostgreSQL uses the native replication mechanisms of PostgreSQL to keep read replicas up to date with changes on the source DB. For general information about read replicas and Amazon RDS, see Working with read replicas.  
upvoted 3 times

A solutions architect must create a disaster recovery (DR) plan for a high-volume software as a service (SaaS) platform. All data for the platform is stored in an Amazon Aurora MySQL DB cluster.

The DR plan must replicate data to a secondary AWS Region.

Which solution will meet these requirements MOST cost-effectively?

- A. Use MySQL binary log replication to an Aurora cluster in the secondary Region. Provision one DB instance for the Aurora cluster in the secondary Region.
- B. Set up an Aurora global database for the DB cluster. When setup is complete, remove the DB instance from the secondary Region.
- C. Use AWS Database Migration Service (AWS DMS) to continuously replicate data to an Aurora cluster in the secondary Region. Remove the DB instance from the secondary Region.
- D. Set up an Aurora global database for the DB cluster. Specify a minimum of one DB instance in the secondary Region.

Correct Answer: B

Community vote distribution

B (57%)

D (32%)

6%

  **awsgeek75** Highly Voted 1 year, 4 months ago

Selected Answer: B

I originally went for D but now I think B is correct. D is active-active cluster so whereas B is active-passive (headless cluster) so it is cheaper than D.

<https://aws.amazon.com/blogs/database/achieve-cost-effective-multi-region-resiliency-with-amazon-aurora-global-database-headless-clusters/>  
upvoted 19 times

  **jennyka76** Highly Voted 2 years, 2 months ago



Answer - A

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Replication.CrossRegion.html>

Before you begin

Before you can create an Aurora MySQL DB cluster that is a cross-Region read replica, you must turn on binary logging on your source Aurora MySQL DB cluster. Cross-region replication for Aurora MySQL uses MySQL binary replication to replay changes on the cross-Region read replica DB cluster.

upvoted 9 times

  **ChrisG1454** 2 years, 2 months ago

The question states " The DR plan must replicate data to a "secondary" AWS Region."

In addition to Aurora Replicas, you have the following options for replication with Aurora MySQL:

Aurora MySQL DB clusters in different AWS Regions.

You can replicate data across multiple Regions by using an Aurora global database. For details, see High availability across AWS Regions with Aurora global databases



You can create an Aurora read replica of an Aurora MySQL DB cluster in a different AWS Region, by using MySQL binary log (binlog) replication. Each cluster can have up to five read replicas created this way, each in a different Region.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>  
upvoted 1 times

  **ChrisG1454** 2 years, 2 months ago

The question is asking for the most cost-effective solution.  
Aurora global databases are more expensive.

<https://aws.amazon.com/rds/aurora/pricing/>  
upvoted 1 times



  **leoattf** 2 years, 2 months ago

On this same URL you provided, there is a note highlighted, stating the following:

"Replication from the primary DB cluster to all secondaries is handled by the Aurora storage layer rather than by the database engine, so lag time for replicating changes is minimal—typically, less than 1 second. Keeping the database engine out of the replication process means that the database engine is dedicated to processing workloads. It also means that you don't need to configure or manage the Aurora MySQL binlog (binary logging) replication."

So, answer should be A

upvoted 2 times

  **leoattf** 2 years, 2 months ago

Correction: So, answer should be D

upvoted 3 times

  **Yak\_Yeti** Most Recent 1 week ago

**Selected Answer: B**

When you set up an Aurora Global Database, the storage layer (Aurora Storage) is continuously and automatically replicated across Regions — independent of whether you have DB instances (readers/writers) running in the secondary Region.

If you remove the DB instance from the secondary Region:

The replicated storage remains.

No compute (no DB instances) = No cost for compute, only storage costs.

BUT:

In a real failover event, you must promote the secondary cluster.

When you promote, you must first create a new DB instance in the secondary Region from the existing storage. This takes some time (typically a few minutes), because Aurora has to spin up the DB instance.

upvoted 1 times

  **tch** 1 month, 3 weeks ago

**Selected Answer: D**

Answer could be D, as this is DR plan, answer B is ask to remove the DB instance which might not suitable for DR plan..... if you remove it... how to replicate data again? you setup again for next replication ?

upvoted 1 times

  **FlyingHawk** 3 months, 1 week ago

**Selected Answer: D**

Giving the requirement is the DR plan for a high-volume software as a service (SaaS) platform, the RTO and RPO should be small, needs a quick recovery, D(warm standby) is a better choice than A ( cold standby, pilot light)

upvoted 2 times

  **LeonSauveterre** 5 months, 1 week ago

**Selected Answer: D**

A - No. Higher cost & higher latency.

B - No. A DB instance is mandatory in each Region for the global database to function.

C - Removing would make DR incomplete, because the secondary Region cannot process failovers or support read replicas.

D - Cost slightly higher, yes, but feasible & cost less than A.

Overall, I don't like this question ...

upvoted 2 times

  **themachine** 10 months, 3 weeks ago

**Selected Answer: B**

Aurora Global Databases offer a cost-effective way to replicate data to a secondary region for disaster recovery. By removing the secondary DB instance after setup, you only pay for storage and minimal compute resources.

upvoted 2 times

  **thewalker** 1 year, 3 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database.html#aurora-global-database.advantages>

upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

Wrong, while D will work, B is cheaper. This question is about DR, not cross region scaling

upvoted 2 times

  **upliftinghut** 1 year, 3 months ago

**Selected Answer: D**

B is more cost-effective however because this is DR so when the region fails => still need a DB to fail over and if setting up a DB from snapshot at the time of failure will be risky => D is the answer

upvoted 3 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: B**

"Achieve cost-effective multi-Region resiliency with Amazon Aurora Global Database headless clusters" is exactly the topic here. "A headless secondary Amazon Aurora database cluster is one without a database instance. This type of configuration can lower expenses for an Aurora global database."

<https://aws.amazon.com/blogs/database/achieve-cost-effective-multi-region-resiliency-with-amazon-aurora-global-database-headless-clusters/>

upvoted 8 times

  **minagaboya** 1 year, 5 months ago

shd be D i guess .. Migrating the database to Amazon Aurora MySQL allows for improved replication performance and higher scalability compared to Amazon RDS for MySQL. Aurora Replicas provide faster replication, reducing the replication lag, and Aurora Auto Scaling ensures that there are enough Aurora Replicas to handle the incoming traffic. Additionally, Aurora MySQL native functions can replace the stored procedures, reducing the load on the database and improving performance.

Option B is not the best solution since adding an ElastiCache for Redis cluster does not address the replication lag issue, and the cache may not have the most up-to-date information. Additionally, replacing the stored procedures with AWS Lambda functions adds additional complexity and may not improve performance.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

This is about a different question



upvoted 4 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: D**

Set up an Aurora global database for the DB cluster. Specify a minimum of one DB instance in the secondary Region

upvoted 1 times


  **vini15** 1 year, 9 months ago

should be B for most cost effective solution.

see the link - Achieve cost-effective multi-Region resiliency with Amazon Aurora Global Database headless clusters

<https://aws.amazon.com/blogs/database/achieve-cost-effective-multi-region-resiliency-with-amazon-aurora-global-database-headless-clusters/>

upvoted 2 times

  **luisgu** 1 year, 11 months ago

**Selected Answer: B**

MOST cost-effective --> B

See section "Creating a headless Aurora DB cluster in a secondary Region" on the link

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>



"Although an Aurora global database requires at least one secondary Aurora DB cluster in a different AWS Region than the primary, you can use a headless configuration for the secondary cluster. A headless secondary Aurora DB cluster is one without a DB instance. This type of configuration can lower expenses for an Aurora global database. In an Aurora DB cluster, compute and storage are decoupled. Without the DB instance, you're not charged for compute, only for storage. If it's set up correctly, a headless secondary's storage volume is kept in-sync with the primary Aurora DB cluster."

upvoted 7 times

  **bsbs1234** 1 year, 7 months ago

upvoted your message, but still think D is correct. Because the question is to design a DR plan.In case of DR, B need to create an instance in DR region manually.

upvoted 2 times

  **Abhineet9148232** 2 years, 1 month ago

**Selected Answer: D**

D: With Amazon Aurora Global Database, you pay for replicated write I/Os between the primary Region and each secondary Region (in this case 1).

Not A because it achieves the same, would be equally costly and adds overhead.

upvoted 3 times

  **[Removed]** 2 years, 2 months ago

**Selected Answer: C**

CCCCC



upvoted 3 times

  **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: D**

I think Amazon is looking for D here. I don' think A is intended because that would require knowledge of MySQL, which isn't what they are testing us on. Not option C because the question states large volume. If the volume were low, then DMS would be better. This question is not a good question.

upvoted 3 times

  **fkie4** 2 years, 1 month ago

very true. Amazon wanna everyone to use AWS, why do they sell for MySQL?

upvoted 1 times

A company has a custom application with embedded credentials that retrieves information from an Amazon RDS MySQL DB instance. Management says the application must be made more secure with the least amount of programming effort.

What should a solutions architect do to meet these requirements?

- A. Use AWS Key Management Service (AWS KMS) to create keys. Configure the application to load the database credentials from AWS KMS. Enable automatic key rotation.
- B. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manager. Configure the application to load the database credentials from Secrets Manager. Create an AWS Lambda function that rotates the credentials in Secret Manager.
- C. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manager. Configure the application to load the database credentials from Secrets Manager. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Secrets Manager.
- D. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Systems Manager Parameter Store. Configure the application to load the database credentials from Parameter Store. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Parameter Store.

**Correct Answer:** C

*Community vote distribution*

C (100%)

  **cloudbusting** Highly Voted 2 years, 2 months ago

Parameter Store does not provide automatic credential rotation.  
upvoted 16 times

  **Bhawesh** Highly Voted 2 years, 2 months ago

**Selected Answer: C**

C. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manager. Configure the application to load the database credentials from Secrets Manager. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Secrets Manager.


<https://www.examttopics.com/discussions/amazon/view/46483-exam-aws-certified-solutions-architect-associate-saa-c02/>  
upvoted 12 times

  **JA2018** Most Recent 5 months, 2 weeks ago

**Selected Answer: C**

Management says the application must be made more secure with the least amount of programming effort.

IMHO, this will rule out Lambda for this STEM.  
upvoted 1 times

  **Gape4** 10 months, 2 weeks ago

**Selected Answer: C**

credentials from Secrets Manager...  
upvoted 2 times

  **d401c0d** 1 year, 1 month ago

question is asking for "more secure with the least amount of programming effort." = Secrets Manager + Secretes Manager's built in rotation schedule instead of Lambda.  
upvoted 2 times

  **awsgeek75** 1 year, 4 months ago

**Selected Answer: C**

A KMS is for encryption keys specifically so this is a long way of doing the credentials storage  
B is too much work for rotation  
C exactly what secrets manager is designed for  
D You can do that if C wasn't an option  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

Store the RDS credentials in Secrets Manager  
Configure the application to retrieve the credentials from Secrets Manager  
Use Secrets Manager's built-in rotation to rotate the RDS credentials automatically  
upvoted 2 times

🗲️ 👤 **Hades2231** 1 year, 8 months ago

**Selected Answer: C**

Secrets Manager can handle the rotation, so no need for Lambda to rotate the keys.  
upvoted 2 times

🗲️ 👤 **chen0305\_099** 1 year, 8 months ago

WHY NOT B ?  
upvoted 1 times

🗲️ 👤 **StacyY** 1 year, 8 months ago

B, we need lambda for password rotation, confirmed!  
upvoted 2 times

🗲️ 👤 **Nikki013** 1 year, 8 months ago

It is not needed for certain types RDS, including MySQL as Secrets Manager has built-in rotation capabilities for it:  
<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>  
upvoted 3 times

🗲️ 👤 **Abrar2022** 1 year, 10 months ago

**Selected Answer: C**

If you need your DB to store credentials then use AWS Secret Manager. System Manager Paramater Store is for CloudFormation (no rotation)  
upvoted 2 times

🗲️ 👤 **AlessandraSAA** 2 years, 1 month ago

why it's not A?  
upvoted 4 times

🗲️ 👤 **MssP** 2 years, 1 month ago

It is asking for credentials, not for encryption keys.  
upvoted 7 times

🗲️ 👤 **PoisonBlack** 2 years ago

So credentials rotation is secrets manager and key rotation is KMS?  
upvoted 3 times

🗲️ 👤 **bdp123** 2 years, 2 months ago

**Selected Answer: C**

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>  
upvoted 2 times

🗲️ 👤 **LuckyAro** 2 years, 2 months ago

**Selected Answer: C**

C is a valid solution for securing the custom application with the least amount of programming effort. It involves creating credentials on the RDS for MySQL database for the application user and storing them in AWS Secrets Manager. The application can then be configured to load the database credentials from Secrets Manager. Additionally, the solution includes setting up a credentials rotation schedule for the application user in the RDS for MySQL database using Secrets Manager, which will automatically rotate the credentials at a specified interval without requiring any programming effort.  
upvoted 4 times

🗲️ 👤 **bdp123** 2 years, 2 months ago

**Selected Answer: C**

[https://docs.aws.amazon.com/secretsmanager/latest/userguide/create\\_database\\_secret.html](https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_database_secret.html)  
upvoted 3 times

🗲️ 👤 **jennyka76** 2 years, 2 months ago

Answer - C  
<https://ws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>  
upvoted 4 times

A media company hosts its website on AWS. The website application’s architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora. The company’s cybersecurity team reports that the application is vulnerable to SQL injection.

How should the company resolve this issue?

- A. Use AWS WAF in front of the ALB. Associate the appropriate web ACLs with AWS WAF.
- B. Create an ALB listener rule to reply to SQL injections with a fixed response.
- C. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.
- D. Set up Amazon Inspector to block all SQL injection attempts automatically.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **Bhawesh** Highly Voted 1 year, 2 months ago

**Selected Answer: A**

A. Use AWS WAF in front of the ALB. Associate the appropriate web ACLs with AWS WAF.

SQL Injection - AWS WAF  
DDoS - AWS Shield  
upvoted 24 times

  **jennyka76** Highly Voted 1 year, 2 months ago

Answer - A  
<https://aws.amazon.com/premiumsupport/knowledge-center/waf-block-common-attacks/#:~:text=To%20protect%20your%20applications%20against,%2C%20query%20string%2C%20or%20URI.>  
-----  
Protect against SQL injection and cross-site scripting  
To protect your applications against SQL injection and cross-site scripting (XSS) attacks, use the built-in SQL injection and cross-site scripting engines. Remember that attacks can be performed on different parts of the HTTP request, such as the HTTP header, query string, or URI. Configure the AWS WAF rules to inspect different parts of the HTTP request against the built-in mitigation engines.  
upvoted 8 times

  **wsdasdasdqwdaw** Most Recent 6 months, 2 weeks ago

AWS WAF - for SQL Injection ---> A  
  
AWS Shield - for DDOS  
Amazon Inspector - for automated security assessment, like known vulnerability  
upvoted 4 times

  **Guru4Cloud** 8 months ago

**Selected Answer: A**  
° Use AWS WAF in front of the Application Load Balancer  
° Configure appropriate WAF web ACLs to detect and block SQL injection patterns  
The key points:  
° Website hosted on EC2 behind an ALB with Aurora database  
° Application is vulnerable to SQL injection attacks  
° AWS WAF is designed to detect and block SQL injection and other common web exploits. It can be placed in front of the ALB to inspect all incoming requests. WAF rules can identify malicious SQL patterns and block them.  
upvoted 2 times

  **KMohsoe** 11 months, 2 weeks ago

**Selected Answer: A**  
SQL injection -> WAF  
upvoted 2 times

  **lexotan** 1 year ago

**Selected Answer: A**  
WAF is the right one  
upvoted 2 times

  **akram\_akram** 1 year ago

**Selected Answer: A**



SQL Injection - AWS WAF  
DDoS - AWS Shield

upvoted 2 times

  **movva12** 1 year, 1 month ago

Answer C - Shield Advanced (WAF + Firewall Manager)

upvoted 1 times

  **fkie4** 1 year, 1 month ago

**Selected Answer: A**

It is A. I am happy to see Amazon gives out score like this...

upvoted 2 times

  **LuckyAro** 1 year, 2 months ago

**Selected Answer: A**

AWS WAF is a managed service that protects web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF enables customers to create custom rules that block common attack patterns, such as SQL injection attacks.



By using AWS WAF in front of the ALB and associating the appropriate web ACLs with AWS WAF, the company can protect its website application from SQL injection attacks. AWS WAF will inspect incoming traffic to the website application and block requests that match the defined SQL injection patterns in the web ACLs. This will help to prevent SQL injection attacks from reaching the application, thereby improving the overall security posture of the application.

upvoted 3 times

  **LuckyAro** 1 year, 2 months ago

B, C, and D are not the best solutions for this issue. Replying to SQL injections with a fixed response  
(B) is not a recommended approach as it does not actually fix the vulnerability, but only masks the issue. Subscribing to AWS Shield Advanced  
(C) is useful to protect against DDoS attacks but does not protect against SQL injection vulnerabilities. Amazon Inspector  
(D) is a vulnerability assessment tool and can identify vulnerabilities but cannot block attacks in real-time.

upvoted 3 times

  **pbpally** 1 year, 2 months ago

**Selected Answer: A**

Bhawesh answers it perfect so I'm avoiding redundancy but agree on it being A.

upvoted 2 times

A company has an Amazon S3 data lake that is governed by AWS Lake Formation. The company wants to create a visualization in Amazon QuickSight by joining the data in the data lake with operational data that is stored in an Amazon Aurora MySQL database. The company wants to enforce column-level authorization so that the company’s marketing team can access only a subset of columns in the database.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon EMR to ingest the data directly from the database to the QuickSight SPICE engine. Include only the required columns.
- B. Use AWS Glue Studio to ingest the data from the database to the S3 data lake. Attach an IAM policy to the QuickSight users to enforce column-level access control. Use Amazon S3 as the data source in QuickSight.
- C. Use AWS Glue Elastic Views to create a materialized view for the database in Amazon S3. Create an S3 bucket policy to enforce column-level access control for the QuickSight users. Use Amazon S3 as the data source in QuickSight.
- D. Use a Lake Formation blueprint to ingest the data from the database to the S3 data lake. Use Lake Formation to enforce column-level access control for the QuickSight users. Use Amazon Athena as the data source in QuickSight.

Correct Answer: D

Community vote distribution

D (100%)

  **K0nAn** Highly Voted 1 year, 8 months ago

Selected Answer: D

This solution leverages AWS Lake Formation to ingest data from the Aurora MySQL database into the S3 data lake, while enforcing column-level access control for QuickSight users. Lake Formation can be used to create and manage the data lake's metadata and enforce security and governance policies, including column-level access control. This solution then uses Amazon Athena as the data source in QuickSight to query the data in the S3 data lake. This solution minimizes operational overhead by leveraging AWS services to manage and secure the data, and by using a standard query service (Amazon Athena) to provide a SQL interface to the data.

upvoted 14 times

  **jennyka76** Highly Voted 1 year, 8 months ago

Answer - D

<https://aws.amazon.com/blogs/big-data/enforce-column-level-authorization-with-amazon-quicksight-and-aws-lake-formation/>

upvoted 10 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

Selected Answer: D

<https://docs.aws.amazon.com/lake-formation/latest/dg/workflows-about.html>

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

Selected Answer: D

Use a Lake Formation blueprint to ingest data from the Aurora database into the S3 data lake  
Leverage Lake Formation to enforce column-level access control for the marketing team  
Use Amazon Athena as the data source in QuickSight  
The key points:

Need to join S3 data lake data with Aurora MySQL data  
Require column-level access controls for marketing team in QuickSight  
Minimize operational overhead  
upvoted 4 times

  **LuckyAro** 1 year, 8 months ago

Selected Answer: D

Using a Lake Formation blueprint to ingest the data from the database to the S3 data lake, using Lake Formation to enforce column-level access control for the QuickSight users, and using Amazon Athena as the data source in QuickSight. This solution requires the least operational overhead as it utilizes the features provided by AWS Lake Formation to enforce column-level authorization, which simplifies the process and reduces the need for additional configuration and maintenance.

upvoted 5 times

  **Bhawesh** 1 year, 8 months ago

Selected Answer: D

D. Use a Lake Formation blueprint to ingest the data from the database to the S3 data lake. Use Lake Formation to enforce column-level access control for the QuickSight users. Use Amazon Athena as the data source in QuickSight.

<https://www.examttopics.com/discussions/amazon/view/80865-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 3 times

A transaction processing company has weekly scripted batch jobs that run on Amazon EC2 instances. The EC2 instances are in an Auto Scaling group. The number of transactions can vary, but the baseline CPU utilization that is noted on each run is at least 60%. The company needs to provision the capacity 30 minutes before the jobs run.

Currently, engineers complete this task by manually modifying the Auto Scaling group parameters. The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts. The company needs an automated way to modify the Auto Scaling group's desired capacity.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a dynamic scaling policy for the Auto Scaling group. Configure the policy to scale based on the CPU utilization metric. Set the target value for the metric to 60%.
- B. Create a scheduled scaling policy for the Auto Scaling group. Set the appropriate desired capacity, minimum capacity, and maximum capacity. Set the recurrence to weekly. Set the start time to 30 minutes before the batch jobs run.
- C. Create a predictive scaling policy for the Auto Scaling group. Configure the policy to scale based on forecast. Set the scaling metric to CPU utilization. Set the target value for the metric to 60%. In the policy, set the instances to pre-launch 30 minutes before the jobs run.
- D. Create an Amazon EventBridge event to invoke an AWS Lambda function when the CPU utilization metric value for the Auto Scaling group reaches 60%. Configure the Lambda function to increase the Auto Scaling group's desired capacity and maximum capacity by 20%.

**Correct Answer: C**

*Community vote distribution*

C (69%)B (28%)

**fkie4**

Highly Voted

2 years, 1 month ago

Selected Answer: C

B is NOT correct. the question said "The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts."

answer B said "Set the appropriate desired capacity, minimum capacity, and maximum capacity".

how can someone set desired capacity if he has no resources to analyze the required capacity.

Read carefully Amigo

upvoted 26 times

**omoakin**

1 year, 11 months ago

scheduled scaling....

upvoted 4 times

**jjcode**

1 year, 2 months ago

works loads can vary, how can you predict something that is random?

upvoted 2 times

**ealpuche**

1 year, 11 months ago

But you can make a vague estimation according to the resources used; you don't need to make machine learning models to do that. You only need common sense.

upvoted 1 times

**Murtadhaceit**

1 year, 4 months ago

Your explanation is contradicting your answer. Since "the company does not have the resources to analyze the required capacity trend for the ASG", how come they can create and ASG based on a historic trend?

C doesn't make sense for me.

upvoted 4 times

**neverdie**

Highly Voted



2 years, 1 month ago



Selected Answer: B



A scheduled scaling policy allows you to set up specific times for your Auto Scaling group to scale out or scale in. By creating a scheduled scaling policy for the Auto Scaling group, you can set the appropriate desired capacity, minimum capacity, and maximum capacity, and set the recurrence to weekly. You can then set the start time to 30 minutes before the batch jobs run, ensuring that the required capacity is provisioned before the jobs run.



Option C, creating a predictive scaling policy for the Auto Scaling group, is not necessary in this scenario since the company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts. This would require analyzing the required capacity trends for the Auto Scaling group counts to determine the appropriate scaling policy.



upvoted 6 times



  **[Removed]** 2 years, 1 month ago  
(typo above) C is correct..  
upvoted 1 times



  **MssP** 2 years, 1 month ago  
Look at fkie4 comment... no way to know desired capacity!!! -> B not correct  
upvoted 1 times



  **Lalo** 1 year, 11 months ago  
the text says  
1.-"A transaction processing company has weekly scripted batch jobs", there is a Schedule  
2.-" The company does not have the resources to analyze the required capacity trends for the Auto Scaling " Do not use  
the answer is B  
upvoted 4 times



  **[Removed]** 2 years, 1 month ago  
B is correct. "Predictive scaling uses machine learning to predict capacity requirements based on historical data from CloudWatch.", meaning the company does not have to analyze the capacity trends themselves. <https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-predictive-scaling.html>  
upvoted 3 times



  **bora4motion** Most Recent 4 days, 16 hours ago  
**Selected Answer: B**  
B is correct  
You know the schedule: weekly jobs → exact match.  
You don't need to analyze utilization trends → simple config.  
1-time setup → minimal overhead.  
C is wrong, at first glance, looks great.  
✗ BUT it requires historical data and analysis — which the company cannot support.  
Setup and tuning introduce overhead.  
✗ Breaks the “company lacks resources to analyze trends” condition.  
✗ Incorrect.  
upvoted 1 times



  **9820ad3** 1 month ago  
**Selected Answer: B**  
B is correct answer .....  
upvoted 1 times



  **Denise123** 4 months, 2 weeks ago  
**Selected Answer: C**  
Option A - Dynamic scaling policy is reactive not proactive.  
Option B - Scheduled policy is not suitable as it says the work loads can vary  
Option C - Correct  
Option D - Irrelevant  
upvoted 1 times

  **EllenLiu** 4 months, 2 weeks ago  
**Selected Answer: C**  
"The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts. " ==> this means this company can take advantage of cloudwatch metric and auto scaling on AWS to set up a predictive scaling group, so that no need to have any resources for this kind of task any more.??  
upvoted 1 times

  **studydue** 6 months ago  
Answer should be B  
Pref: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-predictive-scaling.html>  
Pref: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/predictive-scaling-policy-overview.html>  
keyword for this question: "s. The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts"  
upvoted 2 times

  **Hkayne** 1 year ago  
**Selected Answer: C**  
B or C.  
I think C because the company needs an automated way to modify the autoscaling desired capacity  
upvoted 1 times

  **jjcode** 1 year, 2 months ago  
How does C works with : transactions can vary, clearly C is designed for workloads that are predictable, if the transactions can vary then predictive scaling will not work. The only one that will work is scheduled since its based on time not workload intensity.  
upvoted 3 times

  **pentium75** 1 year, 4 months ago  
**Selected Answer: C**

C per <https://docs.aws.amazon.com/autoscaling/ec2/userguide/predictive-scaling-create-policy.html>.

B is out because it wants the company to 'set the desired/minimum/maximum capacity' but "the company does not have the resources to analyze the required capacity".

upvoted 6 times

  **Cyberkayu** 1 year, 4 months ago

Lambda did not appear to take over scripting/batch job, what a surprise

upvoted 4 times

  **daniel1** 1 year, 6 months ago

**Selected Answer: B**

From GPT4:  
mong the provided options, creating a scheduled scaling policy (Option B) is the most direct and efficient way to ensure that the necessary capacity is provisioned 30 minutes before the weekly batch jobs run, with the least operational overhead. Here's a breakdown of Option B:

B. Create a scheduled scaling policy for the Auto Scaling group. Set the appropriate desired capacity, minimum capacity, and maximum capacity. Set the recurrence to weekly. Set the start time to 30 minutes before the batch jobs run.

Scheduled scaling allows you to change the desired capacity of your Auto Scaling group based on a schedule. In this case, setting the recurrence to weekly and adjusting the start time to 30 minutes before the batch jobs run will ensure that the necessary capacity is available when needed, without requiring manual intervention.

upvoted 5 times

  **TheFivePips** 1 year, 2 months ago

yeah chatgpt told me this, so maybe dont take its word as gospel:

Upon reviewing the question again, it appears that the requirements emphasize the need to provision capacity 30 minutes before the batch jobs run and the company's constraint of not having resources to analyze capacity trends. In this context, the most suitable solution is C.

Predictive Scaling can use historical data to forecast future capacity needs.  
Configuring the policy to scale based on CPU utilization with a target value of 60% aligns with the baseline CPU utilization mentioned in the scenario.  
Setting instances to pre-launch 30 minutes before the jobs run provides the desired capacity just in time.

upvoted 2 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: C**

Predictive scaling: increases the number of EC2 instances in your Auto Scaling group in advance of daily and weekly patterns in traffic flows. If you have regular patterns of traffic increases use predictive scaling, to help you scale faster by launching capacity in advance of forecasted load. You don't have to spend time reviewing your application's load patterns and trying to schedule the right amount of capacity using scheduled scaling. Predictive scaling uses machine learning to predict capacity requirements based on historical data from CloudWatch. The machine learning algorithm consumes the available historical data and calculates capacity that best fits the historical load pattern, and then continuously learns based on new data to make future forecasts more accurate.

upvoted 2 times

  **bsbs1234** 1 year, 7 months ago

should be C. Question does not say how long the job will run. don't know when to set the end time in the schedule policy.

upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: C**

C is correct!



upvoted 2 times

  **Abrar2022** 1 year, 10 months ago

**Selected Answer: C**

if the baseline CPU utilization is 60%, then that's enough information needed to determaine you to predict some aspect of the usage in the future. So key word "predictive" judging by past usage.

upvoted 2 times

  **omoakin** 1 year, 11 months ago

BBBBBBBBBBBBBB

upvoted 1 times

A solutions architect is designing a company’s disaster recovery (DR) architecture. The company has a MySQL database that runs on an Amazon EC2 instance in a private subnet with scheduled backup. The DR design needs to include multiple AWS Regions.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Migrate the MySQL database to multiple EC2 instances. Configure a standby EC2 instance in the DR Region. Turn on replication.
- B. Migrate the MySQL database to Amazon RDS. Use a Multi-AZ deployment. Turn on read replication for the primary DB instance in the different Availability Zones.
- C. Migrate the MySQL database to an Amazon Aurora global database. Host the primary DB cluster in the primary Region. Host the secondary DB cluster in the DR Region.
- D. Store the scheduled backup of the MySQL database in an Amazon S3 bucket that is configured for S3 Cross-Region Replication (CRR). Use the data backup to restore the database in the DR Region.

**Correct Answer: C**

*Community vote distribution*

C (100%)

- LuckyAro**

Highly Voted

1 year, 8 months ago

C: Migrate MySQL database to an Amazon Aurora global database is the best solution because it requires minimal operational overhead. Aurora is a managed service that provides automatic failover, so standby instances do not need to be manually configured. The primary DB cluster can be hosted in the primary Region, and the secondary DB cluster can be hosted in the DR Region. This approach ensures that the data is always available and up-to-date in multiple Regions, without requiring significant manual intervention.

upvoted 10 times
- AlessandraSAA**

Highly Voted

1 year, 7 months ago

**Selected Answer: C**

A. Multiple EC2 instances to be configured and updated manually in case of DR.  
B. Amazon RDS=Multi-AZ while it asks to be multi-region  
C. correct, see comment from LuckyAro  
D. Manual process to start the DR, therefore same limitation as answer A

upvoted 10 times
- gulmichamagaun5**

Most Recent

10 months, 1 week ago

hello friends, question required: The DR design needs to include multiple AWS Regions, but the correct answer is B, how it comes, because the DR here is on AZ not Different Region so the i would go with D

upvoted 1 times
- TariqKipkemei**

1 year ago

**Selected Answer: C**

LEAST operational overhead = Serverless = Amazon Aurora global database

upvoted 3 times
- Guru4Cloud**

1 year, 2 months ago

**Selected Answer: C**

Amazon Aurora global database can span and replicate DB Servers between multiple AWS Regions. And also compatible with MySQL.

upvoted 2 times
- GalileoEC2**

1 year, 7 months ago

C, Why B? B is multi zone in one region, C is multi region as it was requested

upvoted 2 times
- lucdt4**

1 year, 5 months ago

" The DR design needs to include multiple AWS Regions."  
with the requirement "DR SITE multiple AWS region" -> B is wrong, because it deploy multy AZ (this is not multi region)

upvoted 2 times
- KZM**

1 year, 8 months ago

Amazon Aurora global database can span and replicate DB Servers between multiple AWS Regions. And also compatible with MySQL.

upvoted 4 times
- LuckyAro**



1 year, 8 months ago

With dynamic scaling, the Auto Scaling group will automatically adjust the number of instances based on the actual workload. The target value for the CPU utilization metric is set to 60%, which is the baseline CPU utilization that is noted on each run, indicating that this is a reasonable level of



utilization for the workload. This solution does not require any scheduling or forecasting, reducing the operational overhead.  
upvoted 1 times

  **LuckyAro** 1 year, 8 months ago

Sorry, Posted right answer to the wrong question, mistakenly clicked the next question, sorry.  
upvoted 6 times

  **geekgirl22** 1 year, 8 months ago

C is the answer as RDS is only multi-zone not multi region.  
upvoted 2 times


  **bdp123** 1 year, 8 months ago

**Selected Answer: C**

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>  
upvoted 2 times



  **SMAZ** 1 year, 8 months ago

C  
option A has operation overhead whereas option C not.  
upvoted 2 times

  **alexman** 1 year, 8 months ago

**Selected Answer: C**

C mentions multiple regions. Option B is within the same region  
upvoted 4 times

  **jennyka76** 1 year, 8 months ago

ANSWER - B ?? NOT SURE  
upvoted 1 times

A company has a Java application that uses Amazon Simple Queue Service (Amazon SQS) to parse messages. The application cannot parse messages that are larger than 256 KB in size. The company wants to implement a solution to give the application the ability to parse messages as large as 50 MB.

Which solution will meet these requirements with the FEWEST changes to the code?

- A. Use the Amazon SQS Extended Client Library for Java to host messages that are larger than 256 KB in Amazon S3.
- B. Use Amazon EventBridge to post large messages from the application instead of Amazon SQS.
- C. Change the limit in Amazon SQS to handle messages that are larger than 256 KB.
- D. Store messages that are larger than 256 KB in Amazon Elastic File System (Amazon EFS). Configure Amazon SQS to reference this location in the messages.

**Correct Answer: A**

*Community vote distribution*

A (100%)

- LuckyAro**

Highly Voted

2 years, 2 months ago

**Selected Answer: A**

A. Use the Amazon SQS Extended Client Library for Java to host messages that are larger than 256 KB in Amazon S3.

Amazon SQS has a limit of 256 KB for the size of messages. To handle messages larger than 256 KB, the Amazon SQS Extended Client Library for Java can be used. This library allows messages larger than 256 KB to be stored in Amazon S3 and provides a way to retrieve and process them. Using this solution, the application code can remain largely unchanged while still being able to process messages up to 50 MB in size.

upvoted 17 times
- Neha999**

Highly Voted

2 years, 2 months ago

A

For messages > 256 KB, use Amazon SQS Extended Client Library for Java

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/quotas-messages.html>

upvoted 7 times
- tonywang0418**

Most Recent

6 months, 2 weeks ago

who would know this...

upvoted 3 times
- Gape4**

10 months, 2 weeks ago

**Selected Answer: A**

To send messages larger than 256 KiB, you can use the Amazon SQS Extended Client Library for Java...

upvoted 2 times
- TariqKipkemei**

1 year, 6 months ago

**Selected Answer: A**

The Amazon SQS Extended Client Library for Java enables you to manage Amazon SQS message payloads with Amazon S3. This is especially useful for storing and retrieving messages with a message payload size greater than the current SQS limit of 256 KB, up to a maximum of 2 GB.

upvoted 4 times
- Guru4Cloud**

1 year, 8 months ago

**Selected Answer: A**

The SQS Extended Client Library enables storing large payloads in S3 while referenced via SQS. The application code can stay almost entirely unchanged - it sends/receives SQS messages normally. The library handles transparently routing the large payloads to S3 behind the scenes

upvoted 2 times
- james2033**

1 year, 9 months ago

**Selected Answer: A**

Quote "The Amazon SQS Extended Client Library for Java enables you to manage Amazon SQS message payloads with Amazon S3." and "An extension to the Amazon SQS client that enables sending and receiving messages up to 2GB via Amazon S3." at

<https://github.com/awslabs/amazon-sqs-java-extended-client-lib>

upvoted 2 times
- Abrar2022**



1 year, 10 months ago

**Selected Answer: A**

Amazon SQS has a limit of 256 KB for the size of messages.

To handle messages larger than 256 KB, the Amazon SQS Extended Client Library for Java can be used.

upvoted 2 times

  **gold4otas** 2 years, 1 month ago

The Amazon SQS Extended Client Library for Java enables you to publish messages that are greater than the current SQS limit of 256 KB, up to a maximum of 2 GB.

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-s3-messages.html>

upvoted 2 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: A**

<https://github.com/aws-labs/amazon-sqs-java-extended-client-lib>

upvoted 4 times

  **Arathore** 2 years, 2 months ago

**Selected Answer: A**

To send messages larger than 256 KiB, you can use the Amazon SQS Extended Client Library for Java. This library allows you to send an Amazon SQS message that contains a reference to a message payload in Amazon S3. The maximum payload size is 2 GB.

upvoted 5 times

A company wants to restrict access to the content of one of its main web applications and to protect the content by using authorization techniques available on AWS. The company wants to implement a serverless architecture and an authentication solution for fewer than 100 users. The solution needs to integrate with the main web application and serve web content globally. The solution must also scale as the company's user base grows while providing the lowest login latency possible.

Which solution will meet these requirements MOST cost-effectively?

- A. Use Amazon Cognito for authentication. Use Lambda@Edge for authorization. Use Amazon CloudFront to serve the web application globally.
- B. Use AWS Directory Service for Microsoft Active Directory for authentication. Use AWS Lambda for authorization. Use an Application Load Balancer to serve the web application globally.
- C. Use Amazon Cognito for authentication. Use AWS Lambda for authorization. Use Amazon S3 Transfer Acceleration to serve the web application globally.
- D. Use AWS Directory Service for Microsoft Active Directory for authentication. Use Lambda@Edge for authorization. Use AWS Elastic Beanstalk to serve the web application globally.

**Correct Answer: A**

*Community vote distribution*

A (100%)

- Lonojack**

Highly Voted

1 year, 8 months ago

Selected Answer: A

CloudFront=globally  
Lambda@edge = Authorization/ Latency  
Cognito=Authentication for Web apps  
upvoted 15 times
- Lin878**

Most Recent

4 months, 4 weeks ago

Selected Answer: A

<https://aws.amazon.com/blogs/networking-and-content-delivery/external-server-authorization-with-lambdaedge/>  
upvoted 2 times
- Cyberkayu**

10 months, 3 weeks ago

fewer than 100 users but scattered around the globe, lowest latency.

Should have do nothing, most cost effective.  
upvoted 3 times
- TariqKipkemei**

1 year ago

Selected Answer: A

Use Amazon Cognito for authentication. Use Lambda@Edge for authorization. Use Amazon CloudFront to serve the web application globally  
upvoted 3 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: A

Amazon Cognito is a serverless authentication service that can be used to easily add user sign-up and authentication to web and mobile apps. It is a good choice for this scenario because it is scalable and can handle a small number of users without any additional costs.  
  
Lambda@Edge is a serverless compute service that can be used to run code at the edge of the AWS network. It is a good choice for this scenario because it can be used to perform authorization checks at the edge, which can improve the login latency.  
  
Amazon CloudFront is a content delivery network (CDN) that can be used to serve web content globally. It is a good choice for this scenario because it can cache web content closer to users, which can improve the performance of the web application.  
upvoted 4 times
- antropaws**

1 year, 5 months ago

Selected Answer: A

A is perfect.  
upvoted 2 times
- kraken21**


1 year, 7 months ago

**Selected Answer: A**

Lambda@Edge for authorization

<https://aws.amazon.com/blogs/networking-and-content-delivery/adding-http-security-headers-using-lambdaedge-and-amazon-cloudfront/>

upvoted 3 times

  **LuckyAro** 1 year, 8 months ago

**Selected Answer: A**

Amazon CloudFront is a global content delivery network (CDN) service that can securely deliver web content, videos, and APIs at scale. It integrates with Cognito for authentication and with Lambda@Edge for authorization, making it an ideal choice for serving web content globally.

Lambda@Edge is a service that lets you run AWS Lambda functions globally closer to users, providing lower latency and faster response times. It can also handle authorization logic at the edge to secure content in CloudFront. For this scenario, Lambda@Edge can provide authorization for the web application while leveraging the low-latency benefit of running at the edge.



upvoted 3 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: A**

CloudFront to serve globally

upvoted 2 times

  **SMAZ** 1 year, 8 months ago

A

Amazon Cognito for authentication and Lambda@Edge for authorization, Amazon CloudFront to serve the web application globally provides low-latency content delivery

upvoted 4 times

A company has an aging network-attached storage (NAS) array in its data center. The NAS array presents SMB shares and NFS shares to client workstations. The company does not want to purchase a new NAS array. The company also does not want to incur the cost of renewing the NAS array’s support contract. Some of the data is accessed frequently, but much of the data is inactive.

A solutions architect needs to implement a solution that migrates the data to Amazon S3, uses S3 Lifecycle policies, and maintains the same look and feel for the client workstations. The solutions architect has identified AWS Storage Gateway as part of the solution.

Which type of storage gateway should the solutions architect provision to meet these requirements?

- A. Volume Gateway
- B. Tape Gateway
- C. Amazon FSx File Gateway
- D. Amazon S3 File Gateway

**Correct Answer:** D

Community vote distribution

D (100%)

- LuckyAro**

Highly Voted

1 year, 8 months ago

Selected Answer: D

Amazon S3 File Gateway provides on-premises applications with access to virtually unlimited cloud storage using NFS and SMB file interfaces. It seamlessly moves frequently accessed data to a low-latency cache while storing colder data in Amazon S3, using S3 Lifecycle policies to transition data between storage classes over time.

In this case, the company's aging NAS array can be replaced with an Amazon S3 File Gateway that presents the same NFS and SMB shares to the client workstations. The data can then be migrated to Amazon S3 and managed using S3 Lifecycle policies

upvoted 17 times
- everfly**

Highly Voted

1 year, 8 months ago

Selected Answer: D

Amazon S3 File Gateway provides a file interface to objects stored in S3. It can be used for a file-based interface with S3, which allows the company to migrate their NAS array data to S3 while maintaining the same look and feel for client workstations. Amazon S3 File Gateway supports SMB and NFS protocols, which will allow clients to continue to access the data using these protocols. Additionally, Amazon S3 Lifecycle policies can be used to automate the movement of data to lower-cost storage tiers, reducing the storage cost of inactive data.

upvoted 7 times
- pentium75**

Most Recent

10 months, 1 week ago

Selected Answer: D

A - provides virtual disk via iSCSI  
B - provides virtual tape via iSCSI  
C - provides access to FSx via SMB

upvoted 4 times
- TariqKipkemei**

1 year ago

Selected Answer: D

The Amazon S3 File Gateway enables you to store and retrieve objects in Amazon Simple Storage Service (S3) using file protocols such as Network File System (NFS) and Server Message Block (SMB).

upvoted 4 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: D

It provides an easy way to lift-and-shift file data from the existing NAS to Amazon S3. The S3 File Gateway presents SMB and NFS file shares that client workstations can access just like the NAS shares.

Behind the scenes, it moves the file data to S3 storage, storing it durably and cost-effectively.

S3 Lifecycle policies can be used to transition less frequently accessed data to lower-cost S3 storage tiers like S3 Glacier.

From the client workstation perspective, access to files feels seamless and unchanged after migration to S3. The S3 File Gateway handles the underlying data transfers.


It is a simple, low-cost gateway option tailored for basic file share migration use cases.

upvoted 4 times
- james2033**

1 year, 3 months ago

Selected Answer: D

- Volume Gateway: <https://aws.amazon.com/storagegateway/volume/> (Remove A, related iSCSI)
- Tape Gateway <https://aws.amazon.com/storagegateway/vtl/> (Remove B)
- Amazon FSx File Gateway <https://aws.amazon.com/storagegateway/file/fsx/> (C)
- Why not choose C? Because need working with Amazon S3. (Answer D, and it is correct answer) <https://aws.amazon.com/storagegateway/file/s3/>  
upvoted 4 times

  **siyam008** 1 year, 8 months ago

**Selected Answer: D**

<https://aws.amazon.com/blogs/storage/how-to-create-smb-file-shares-with-aws-storage-gateway-using-hyper-v/>  
upvoted 3 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: D**

<https://aws.amazon.com/about-aws/whats-new/2018/06/aws-storage-gateway-adds-smb-support-to-store-objects-in-amazon-s3/>  
upvoted 4 times

A company has an application that is running on Amazon EC2 instances. A solutions architect has standardized the company on a particular instance family and various instance sizes based on the current needs of the company.

The company wants to maximize cost savings for the application over the next 3 years. The company needs to be able to change the instance family and sizes in the next 6 months based on application popularity and usage.



Which solution will meet these requirements MOST cost-effectively?


- A. Compute Savings Plan
- B. EC2 Instance Savings Plan
- C. Zonal Reserved Instances
- D. Standard Reserved Instances

**Correct Answer: A**

*Community vote distribution*

A (83%)	Other
---------	-------

-   **AlmeroSenior**



Highly Voted 

 1 year, 8 months ago

Selected Answer: A



Read Carefully guys , They need to be able to change FAMILY , and although EC2 Savings has a higher discount , its clearly documented as not allowed >

EC2 Instance Savings Plans provide savings up to 72 percent off On-Demand, in exchange for a commitment to a specific instance family in a chosen AWS Region (for example, M5 in Virginia). These plans automatically apply to usage regardless of size (for example, m5.xlarge, m5.2xlarge, etc.), OS (for example, Windows, Linux, etc.), and tenancy (Host, Dedicated, Default) within the specified family in a Region.

upvoted 21 times
-   **FFO** 1 year, 6 months ago

Savings Plans are a flexible pricing model that offer low prices on Amazon EC2, AWS Lambda, and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. When you sign up for a Savings Plan, you will be charged the discounted Savings Plans price for your usage up to your commitment.



The company wants savings over the next 3 years but wants to change the instance type in 6 months. This invalidates A


upvoted 5 times
-   **FFO** 1 year, 6 months ago

Disregard! found more information:

We recommend Savings Plans (over Reserved Instances). Like Reserved Instances, Savings Plans offer lower prices (up to 72% savings compared to On-Demand Instance pricing). In addition, Savings Plans offer you the flexibility to change your usage as your needs evolve. For example, with Compute Savings Plans, lower prices will automatically apply when you change from C4 to C6g instances, shift a workload from EU (Ireland) to EU (London), or move a workload from Amazon EC2 to AWS Fargate or AWS Lambda.

<https://aws.amazon.com/ec2/pricing/reserved-instances/pricing/>

upvoted 3 times
-   **awsgeek75**

Highly Voted 

 9 months, 2 weeks ago

Selected Answer: A



<https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/savings-plans.html>

Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66% (just like Convertible RIs). These plans automatically apply to EC2 instance usage regardless of instance family...

EC2 Instance Savings Plans provide the lowest prices, offering savings up to 72% (just like Standard RIs) in exchange for commitment to usage of individual instance families



Instance Savings "locks" you in that instance family which is not desired by the company hence A is the best plan as they can change the instance family anytime


upvoted 10 times

  **awsgeek75** 9 months, 2 weeks ago

Also, don't forget, the minimum commitment for both of these plans is 1 year and the company wants the ability to change in 6 months so it has to be a plan which allows changing of instance within the commitment window (no refunds!)

upvoted 4 times

  **xBUGx**

Most Recent 



 7 months ago

Selected Answer: A

EC2 Instance Savings Plans provide the lowest prices, offering savings up to 72% in exchange for commitment to usage of individual instance families in a region (e.g. M5 usage in N. Virginia). This automatically reduces your cost on the selected instance family in that region regardless of AZ, size, OS or tenancy. \*\*\*EC2 Instance Savings Plans give you the flexibility to change your usage between instances within a family in that region.\*\*\* For example, you can move from c5.xlarge running Windows to c5.2xlarge running Linux and automatically benefit from the Savings Plans prices.

<https://aws.amazon.com/savingsplans/faq/#:~:text=EC2%20Instance%20Savings%20Plans%20give,from%20the%20Savings%20Plans%20prices>.

upvoted 4 times

  **Mican07** 9 months, 3 weeks ago

B is the definite answer


upvoted 1 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: A**

B does not allow changing the instance family, despite all the ChatGPT-based answers claiming the opposite

upvoted 3 times

  **meowruki** 11 months, 1 week ago

**Selected Answer: A**

While EC2 Instance Savings Plans also provide cost savings over On-Demand pricing, they offer less flexibility in terms of changing instance families. They provide a discount in excha

upvoted 2 times

  **hungta** 11 months, 3 weeks ago

**Selected Answer: B**

EC2 Instance Savings Plans is most saving. And it is enough for required flexibility


EC2 Instance Savings Plans provide the lowest prices, offering savings up to 72% (just like Standard RIs) in exchange for commitment to usage of individual instance families in a Region (for example, M5 usage in N. Virginia). This automatically reduces your cost on the selected instance family in that region regardless of AZ, size, operating system, or tenancy. EC2 Instance Savings Plans give you the flexibility to change your usage between instances within a family in that Region. For example, you can move from c5.xlarge running Windows to c5.2xlarge running Linux and automatically benefit from the Savings Plans prices.

upvoted 1 times

  **pentium75** 10 months, 1 week ago

But it does not allow changing the instance family, which is a requirement here.

upvoted 3 times

  **Jazz888** 5 months ago

You voted against yourself. Did you mean to vote A?

upvoted 1 times

  **dilaaziz** 11 months, 3 weeks ago

**Selected Answer: A**

<https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/savings-plans.html>

upvoted 2 times

  **EdenWang** 12 months ago

**Selected Answer: B**

The most cost-effective solution that meets the company's requirements would be B. EC2 Instance Savings Plan.

EC2 Instance Savings Plans provide significant cost savings, allowing the company to commit to a consistent amount of usage (measured in \$/hour) for a 1- or 3-year term, and in return, receive a discount on the hourly rate for the instances that match the attributes of the plan.

With EC2 Instance Savings Plans, the company can benefit from the flexibility to change the instance family and sizes over the next 3 years, which aligns with their requirement to adjust based on application popularity and usage.

This option provides the best balance of cost savings and flexibility, making it the most suitable choice for the company's needs.

upvoted 2 times

  **pentium75** 10 months, 1 week ago

"With EC2 Instance Savings Plans, the company can benefit from the flexibility to change the instance family" NO, this is simply wrong. Is this from ChatGPT?

"EC2 Instance Savings Plans provide savings up to 72 percent off On-Demand, in exchange for a commitment to a specific instance family (!) in a chosen AWS Region ... With an EC2 Instance Savings Plan, you can change your instance size within the instance family (!)".

<https://docs.aws.amazon.com/savingsplans/latest/userguide/what-is-savings-plans.html>

upvoted 4 times

  **TariqKipkemei** 1 year ago

**Selected Answer: A**

Change instance family = Compute Savings Plans

upvoted 4 times

  **Wayne23Fang** 1 year, 1 month ago

**Selected Answer: A**

D is not right. D. Standard Reserved Instances. should be Convertible Reserved Instances if you need additional flexibility, such as the ability to use different instance families, operating systems.

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**



The key factors are:

Need to maximize cost savings over 3 years

Ability to change instance family and sizes in 6 months

Standardized on a particular instance family for now

upvoted 2 times

  **awsgeek75** 9 months, 2 weeks ago

"Ability to change instance family and sizes in 6 months" is not allowed in Instance Savings plan so B is wrong

upvoted 2 times

  **Kiki\_Pass** 1 year, 3 months ago

Why not C? Can do with Convertible Reserved Instance

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/reserved-instances-types.html>

upvoted 1 times

  **ITV2021** 1 year, 3 months ago

**Selected Answer: A**

<https://aws.amazon.com/savingsplans/compute-pricing/>

upvoted 2 times



  **Mia2009687** 1 year, 3 months ago

**Selected Answer: A**

EC2 Instance Savings Plan cannot change the family.

<https://docs.aws.amazon.com/savingsplans/latest/userguide/what-is-savings-plans.html>

upvoted 2 times

  **mattcl** 1 year, 4 months ago

Anser D: You can use Standard Reserved Instances when you know that you need a specific instance type.

upvoted 1 times

  **kruasan** 1 year, 6 months ago

**Selected Answer: A**

Savings Plans offer a flexible pricing model that provides savings on AWS usage. You can save up to 72 percent on your AWS compute workloads. Compute Savings Plans provide lower prices on Amazon EC2 instance usage regardless of instance family, size, OS, tenancy, or AWS Region. This also applies to AWS Fargate and AWS Lambda usage. SageMaker Savings Plans provide you with lower prices for your Amazon SageMaker instance usage, regardless of your instance family, size, component, or AWS Region.

<https://docs.aws.amazon.com/savingsplans/latest/userguide/what-is-savings-plans.html>

upvoted 3 times

  **kruasan** 1 year, 6 months ago

With an EC2 Instance Savings Plan, you can change your instance size within the instance family (for example, from c5.xlarge to c5.2xlarge) or the operating system (for example, from Windows to Linux), or move from Dedicated tenancy to Default and continue to receive the discounted rate provided by your EC2 Instance Savings Plan.

<https://docs.aws.amazon.com/savingsplans/latest/userguide/what-is-savings-plans.html>

upvoted 1 times

  **kruasan** 1 year, 6 months ago

The company needs to be able to change the instance family and sizes in the next 6 months based on application popularity and usage.

Therefore EC2 Instance Savings Plan prerequisites are not fulfilled

upvoted 3 times

A company collects data from a large number of participants who use wearable devices. The company stores the data in an Amazon DynamoDB table and uses applications to analyze the data. The data workload is constant and predictable. The company wants to stay at or below its forecasted budget for DynamoDB.

Which solution will meet these requirements MOST cost-effectively?



- A. Use provisioned mode and DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA). Reserve capacity for the forecasted workload.
- B. Use provisioned mode. Specify the read capacity units (RCUs) and write capacity units (WCUs).
- C. Use on-demand mode. Set the read capacity units (RCUs) and write capacity units (WCUs) high enough to accommodate changes in the workload.
- D. Use on-demand mode. Specify the read capacity units (RCUs) and write capacity units (WCUs) with reserved capacity.


**Correct Answer: B**

*Community vote distribution*

B (90%)

10%



-   **everfly**


Highly Voted 

 1 year, 8 months ago

Selected Answer: B

The data workload is constant and predictable.



upvoted 11 times
-   **hovnival**


Highly Voted 

 1 year ago

Selected Answer: B

I think it is not possible to set Read Capacity Units(RCU)/Write Capacity Units(WCU) in on-demand mode.

upvoted 6 times
-   **pentium75**



Most Recent 

 10 months ago

Selected Answer: B

C and D are impossible because you don't set or specify RCUs and WCUs in on-demand mode.



A is wrong because there is no indication of "infrequent access", and "the data workload is constant", there is no different between the current and the "forecasted" workload.

upvoted 4 times
-   **wsdasdasdqwda** 1 year ago

predictable/constant => provisioned mode. On-demand mode is more suitable for workloads that are unpredictable and can vary widely from minute to minute.



The use case is not for Standard-IA which is described here: <https://aws.amazon.com/dynamodb/standard-ia/>

=> Option B

upvoted 3 times
-   **TariqKipkemei** 1 year ago

Selected Answer: B

I rule out A because of this 'Standard-Infrequent Access ', clearly the company uses applications to analyze the data. The data workload is constant and predictable making provisioned mode the best option.

upvoted 2 times
-   **Guru4Cloud** 1 year, 2 months ago



Selected Answer: A

Option B lacks the cost benefits of Standard-IA.

Option C uses more expensive on-demand pricing.

Option D does not actually allow reserving capacity with on-demand mode.

So option A leverages provisioned mode, Standard-IA, and reserved capacity to meet the requirements in a cost-optimal way.

upvoted 1 times
-   **MrAWSAssociate** 1 year, 4 months ago

Selected Answer: A

A is correct!

upvoted 1 times

  **MrAWSAssociate** 1 year, 4 months ago

Sorry, A will not work, since Reserved Capacity can only be used with DynamoDB Standard table class. So, B is right for this case.  
upvoted 4 times



  **UNGMAN** 1 year, 7 months ago

**Selected Answer: B**



예측가능..  
upvoted 4 times

  **kayodea25** 1 year, 7 months ago

Option C is the most cost-effective solution for this scenario. In on-demand mode, DynamoDB automatically scales up or down based on the current workload, so the company only pays for the capacity it uses. By setting the RCUs and WCUs high enough to accommodate changes in the workload, the company can ensure that it always has the necessary capacity without overprovisioning and incurring unnecessary costs. Since the workload is constant and predictable, using provisioned mode with reserved capacity (Options A and D) may result in paying for unused capacity during periods of low demand. Option B, using provisioned mode without reserved capacity, may result in throttling during periods of high demand if the provisioned capacity is not sufficient to handle the workload.  
upvoted 3 times

  **pentium75** 10 months ago

You can't 'set RCUs and WCUs' in on-demand mode.  
upvoted 2 times

  **boxu03** 1 year, 7 months ago

you forgot "The data workload is constant and predictable", should be B  
upvoted 3 times

  **Bofi** 1 year, 7 months ago

Kayode olode..lol  
upvoted 1 times

  **Steve\_4542636** 1 year, 8 months ago

"The data workload is constant and predictable."  
<https://docs.aws.amazon.com/wellarchitected/latest/serverless-applications-lens/capacity.html>

"With provisioned capacity you pay for the provision of read and write capacity units for your DynamoDB tables. Whereas with DynamoDB on-demand you pay per request for the data reads and writes that your application performs on your tables."  
upvoted 2 times

  **Charly0710** 1 year, 8 months ago

**Selected Answer: B**

The data workload is constant and predictable, then, isn't on-demand mode.  
DynamoDB Standard-IA is not necessary in this context  
upvoted 2 times

  **Lonojack** 1 year, 8 months ago

**Selected Answer: B**

The problem with (A) is: "Standard-Infrequent Access". In the question, they say the company has to analyze the Data.  
That's why the Correct answer is (B)  
upvoted 4 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: A**

workload is constant  
upvoted 2 times

  **Lonojack** 1 year, 8 months ago

The problem with (A) is: "Standard-Infrequent Access".  
In the question, they say the company has to analyze the Data.  
Correct answer is (B)  
upvoted 3 times

  **Samuel03** 1 year, 8 months ago

**Selected Answer: B**

As the numbers are already known  
upvoted 4 times

A company stores confidential data in an Amazon Aurora PostgreSQL database in the ap-southeast-3 Region. The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key. The company was recently acquired and must securely share a backup of the database with the acquiring company’s AWS account in ap-southeast-3.

What should a solutions architect do to meet these requirements?

- A. Create a database snapshot. Copy the snapshot to a new unencrypted snapshot. Share the new snapshot with the acquiring company’s AWS account.
- B. Create a database snapshot. Add the acquiring company’s AWS account to the KMS key policy. Share the snapshot with the acquiring company’s AWS account.
- C. Create a database snapshot that uses a different AWS managed KMS key. Add the acquiring company’s AWS account to the KMS key alias. Share the snapshot with the acquiring company's AWS account.
- D. Create a database snapshot. Download the database snapshot. Upload the database snapshot to an Amazon S3 bucket. Update the S3 bucket policy to allow access from the acquiring company’s AWS account.

**Correct Answer: B**

*Community vote distribution*

B (100%)

- Abrar2022**

Highly Voted

1 year, 10 months ago

**Selected Answer: B**

A. - "So let me get this straight, with the current company the data is protected and encrypted. However, for the acquiring company the data is unencrypted? How is that fair?"

C - Wouldn't recommended this option because using a different AWS managed KMS key will not allow the acquiring company's AWS account to access the encrypted data.

D. - Don't risk it for a biscuit and get fired!!!! - by downloading the database snapshot and uploading it to an Amazon S3 bucket. This will increase the risk of data leakage or loss of confidentiality during the transfer process.

B - CORRECT

upvoted 15 times
- njufi**

Most Recent

1 year, 1 month ago

I believe the reason why option C is not the correct answer is that adding the acquiring company's AWS account to the KMS key alias doesn't directly control access to the encrypted data. KMS key aliases are simply alternative names for KMS keys and do not affect access control. Access to encrypted data is goverened by KMS key policies, which define who can use the key for encryption and decryption.

upvoted 2 times
- TariqKipkemei**

1 year, 6 months ago

**Selected Answer: B**

Create a database snapshot. Add the acquiring company’s AWS account to the KMS key policy. Share the snapshot with the acquiring company’s AWS account.

upvoted 2 times
- Vuuu**

1 year, 9 months ago

**Selected Answer: B**

B. Create a database snapshot. Add the acquiring company’s AWS account to the KMS key policy. Share the snapshot with the acquiring company’s AWS account. Most Voted

upvoted 2 times
- Abrar2022**

1 year, 10 months ago

Create a database snapshot of the encrypted. Add the acquiring company’s AWS account to the KMS key policy. Share the snapshot with the acquiring company’s AWS account.

upvoted 2 times
- SkyZeroZx**

2 years ago

**Selected Answer: B**

To securely share a backup of the database with the acquiring company's AWS account in the same Region, a solutions architect should create a database snapshot, add the acquiring company's AWS account to the AWS KMS key policy, and share the snapshot with the acquiring company's AWS account.

Option A, creating an unencrypted snapshot, is not recommended as it will compromise the confidentiality of the data. Option C, creating a

snapshot that uses a different AWS managed KMS key, does not provide any additional security and will unnecessarily complicate the solution. Option D, downloading the database snapshot and uploading it to an S3 bucket, is not secure as it can expose the data during transit.

Therefore, the correct option is B: Create a database snapshot. Add the acquiring company's AWS account to the KMS key policy. Share the snapshot with the acquiring company's AWS account.

upvoted 2 times

🗨️ 👤 **JA2018** 5 months, 2 weeks ago

KMS key aliases are simply alternative names for KMS keys and do not affect access control.

Access to encrypted data is governed by KMS key policies, that define who can use the key for encryption and decryption.

upvoted 1 times

🗨️ 👤 **elearningtakai** 2 years, 1 month ago

**Selected Answer: B**

Option B is the correct answer.

Option A is not recommended because copying the snapshot to a new unencrypted snapshot will compromise the confidentiality of the data.

Option C is not recommended because using a different AWS managed KMS key will not allow the acquiring company's AWS account to access the encrypted data.

Option D is not recommended because downloading the database snapshot and uploading it to an Amazon S3 bucket will increase the risk of data leakage or loss of confidentiality during the transfer process.

upvoted 3 times

🗨️ 👤 **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html>

upvoted 3 times

🗨️ 👤 **geekgirl22** 2 years, 2 months ago

It is C, you have to create a new key. Read below

You can't share a snapshot that's encrypted with the default AWS KMS key. You must create a custom AWS KMS key instead. To share an encrypted Aurora DB cluster snapshot:

Create a custom AWS KMS key.

Add the target account to the custom AWS KMS key.

Create a copy of the DB cluster snapshot using the custom AWS KMS key. Then, share the newly copied snapshot with the target account.

Copy the shared DB cluster snapshot from the target account

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted-snapshot/>

upvoted 1 times

🗨️ 👤 **leoattf** 2 years, 2 months ago

I also thought straight away that it could be C, however, the questions mentions that the database is encrypted with an AWS KMS custom key already. So maybe the letter B could be right, since it already has a custom key, not the default KMS Key.

What do you think?

upvoted 3 times

🗨️ 👤 **enzomv** 2 years, 1 month ago

It is B.

There's no need to create another custom AWS KMS key.

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted-snapshot/>

Give target account access to the custom AWS KMS key within the source account

1. Log in to the source account, and go to the AWS KMS console in the same Region as the DB cluster snapshot.

2. Select Customer-managed keys from the navigation pane.

3. Select your custom AWS KMS key (ALREADY CREATED)

4. From the Other AWS accounts section, select Add another AWS account, and then enter the AWS account number of your target account.

Then:

Copy and share the DB cluster snapshot

upvoted 3 times

🗨️ 👤 **KZM** 2 years, 2 months ago

Yes, as per the given information "The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key", it may not be the default AWS KMS key.

upvoted 1 times

🗨️ 👤 **KZM** 2 years, 2 months ago

Yes, can't share a snapshot that's encrypted with the default AWS KMS key.

But as per the given information "The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key", it may not be the default AWS KMS key.

upvoted 3 times

🗨️ 👤 **enzomv** 2 years, 1 month ago

I agree with KZM.

It is B.

There's no need to create another custom AWS KMS key.

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted-snapshot/>

Give target account access to the custom AWS KMS key within the source account

1. Log in to the source account, and go to the AWS KMS console in the same Region as the DB cluster snapshot.

- 2. Select Customer-managed keys from the navigation pane.
- 3. Select your custom AWS KMS key (ALREADY CREATED)
- 4. From the Other AWS accounts section, select Add another AWS account, and then enter the AWS account number of your target account.

Then:  
Copy and share the DB cluster snapshot  
upvoted 3 times

  **nyx12345** 2 years, 2 months ago



Is it bad that in answer B the acquiring company is using the same KMS key? Should a new KMS key not be used?  
upvoted 2 times

  **geekgirl22** 2 years, 2 months ago

Yes, you are right, read my comment above.  
upvoted 1 times

  **bsbs1234** 1 year, 7 months ago

I think I would agree with you if option C say using a new "customer managed key" instead of AWS managed key  
upvoted 1 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: B**

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted-snapshot/>  
upvoted 3 times

  **jennyka76** 2 years, 2 months ago

ANSWER - B  
upvoted 2 times

A company uses a 100 GB Amazon RDS for Microsoft SQL Server Single-AZ DB instance in the us-east-1 Region to store customer transactions. The company needs high availability and automatic recovery for the DB instance.

The company must also run reports on the RDS database several times a year. The report process causes transactions to take longer than usual to post to the customers’ accounts. The company needs a solution that will improve the performance of the report process.



Which combination of steps will meet these requirements? (Choose two.)


- A. Modify the DB instance from a Single-AZ DB instance to a Multi-AZ deployment.
- B. Take a snapshot of the current DB instance. Restore the snapshot to a new RDS deployment in another Availability Zone.
- C. Create a read replica of the DB instance in a different Availability Zone. Point all requests for reports to the read replica.
- D. Migrate the database to RDS Custom.
- E. Use RDS Proxy to limit reporting requests to the maintenance window.

**Correct Answer: AC**

*Community vote distribution*



AC (100%)


-   **elearningtakai**

Highly Voted 

 2 years, 1 month ago



A and C are the correct choices.  
B. It will not help improve the performance of the report process.  
D. Migrating to RDS Custom does not address the issue of high availability and automatic recovery.  
E. RDS Proxy can help with scalability and high availability but it does not address the issue of performance for the report process. Limiting the reporting requests to the maintenance window will not provide the required availability and recovery for the DB instance.

upvoted 8 times
-   **rockyykrish**

Most Recent 



 5 months, 3 weeks ago

A and C

upvoted 2 times
-   **TariqKipkemei** 1 year, 6 months ago

Selected Answer: AC

Create a Multi-AZ deployment, create a read replica of the DB instance in the second Availability Zone, point all requests for reports to the read replica



upvoted 4 times
-   **Guru4Cloud** 1 year, 8 months ago

Selected Answer: AC

The correct answers are A and C.



A. Modify the DB instance from a Single-AZ DB instance to a Multi-AZ deployment. This will provide high availability and automatic recovery for the DB instance. If the primary DB instance fails, the standby DB instance will automatically become the primary DB instance. This will ensure that the database is always available.

C. Create a read replica of the DB instance in a different Availability Zone. Point all requests for reports to the read replica. This will improve the performance of the report process by offloading the read traffic from the primary DB instance to the read replica. The read replica is a fully synchronized copy of the primary DB instance, so the reports will be accurate.

upvoted 4 times
-   **elearningtakai** 2 years, 1 month ago



Selected Answer: AC

A and C.

upvoted 3 times
-   **WherecanIstart** 2 years, 1 month ago

Selected Answer: AC

Options A & C...

upvoted 4 times
-   **KZM** 2 years, 2 months ago

Options A+C

upvoted 3 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: AC**

<https://medium.com/awesome-cloud/aws-difference-between-multi-az-and-read-replicas-in-amazon-rds-60fe848ef53a>  
upvoted 4 times

  **jennyka76** 2 years, 2 months ago

ANSWER - A & C  
upvoted 4 times

A company is moving its data management application to AWS. The company wants to transition to an event-driven architecture. The architecture needs to be more distributed and to use serverless concepts while performing the different aspects of the workflow. The company also wants to minimize operational overhead.

Which solution will meet these requirements?

- A. Build out the workflow in AWS Glue. Use AWS Glue to invoke AWS Lambda functions to process the workflow steps.
- B. Build out the workflow in AWS Step Functions. Deploy the application on Amazon EC2 instances. Use Step Functions to invoke the workflow steps on the EC2 instances.
- C. Build out the workflow in Amazon EventBridge. Use EventBridge to invoke AWS Lambda functions on a schedule to process the workflow steps.
- D. Build out the workflow in AWS Step Functions. Use Step Functions to create a state machine. Use the state machine to invoke AWS Lambda functions to process the workflow steps.

**Correct Answer: D**

*Community vote distribution*

D (91%)

9%

- Lonojack**

Highly Voted

1 year, 8 months ago

Selected Answer: D

This is why I'm voting D.....QUESTION ASKED FOR IT TO: use serverless concepts while performing the different aspects of the workflow. Is option D utilizing Serverless concepts?

upvoted 13 times
- geekgirl22**

Highly Voted

1 year, 8 months ago

It is D. Cannot be C because C is "scheduled"

upvoted 8 times
- bujuman**

Most Recent

7 months, 2 weeks ago

Selected Answer: D

While considering this requirement: The architecture needs to be more distributed and to use serverless concepts while performing the different aspects of the workflow

And checking the following link : [https://aws.amazon.com/step-functions/?nc1=h\\_ls](https://aws.amazon.com/step-functions/?nc1=h_ls), Answer D is the best for this use case

upvoted 3 times
- TariqKipkemei**

1 year ago

Selected Answer: D

One of the use cases for step functions is to Automate extract, transform, and load (ETL) processes.

<https://aws.amazon.com/step-functions/#:~:text=for%20modern%20applications.-,Use%20cases,-Automate%20extract%2C%20transform>

upvoted 2 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: D

AWS Step functions is serverless Visual workflows for distributed applications

<https://aws.amazon.com/step-functions/>

upvoted 2 times
- TariqKipkemei**

1 year, 5 months ago

Selected Answer: D

Step Functions is based on state machines and tasks. A state machine is a workflow. A task is a state in a workflow that represents a single unit of work that another AWS service performs. Each step in a workflow is a state.

Depending on your use case, you can have Step Functions call AWS services, such as Lambda, to perform tasks.

<https://docs.aws.amazon.com/step-functions/latest/dg/welcome.html>

upvoted 3 times
- TariqKipkemei**

1 year, 5 months ago

Answer is D.

Step Functions is based on state machines and tasks. A state machine is a workflow. A task is a state in a workflow that represents a single unit of work that another AWS service performs. Each step in a workflow is a state.

Depending on your use case, you can have Step Functions call AWS services, such as Lambda, to perform tasks.

<https://docs.aws.amazon.com/step-functions/latest/dg/welcome.html>

upvoted 3 times

  **Karlos99** 1 year, 8 months ago

**Selected Answer: C**

There are two main types of routers used in event-driven architectures: event buses and event topics. At AWS, we offer Amazon EventBridge to build event buses and Amazon Simple Notification Service (SNS) to build event topics. <https://aws.amazon.com/event-driven-architecture/>  
upvoted 1 times

  **pentium75** 10 months, 1 week ago

How do you 'build out a workflow' in EventBridge?  
upvoted 4 times

  **TungPham** 1 year, 8 months ago

**Selected Answer: D**

Step 3: Create a State Machine  
Use the Step Functions console to create a state machine that invokes the Lambda function that you created earlier in Step 1.  
<https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html>  
In Step Functions, a workflow is called a state machine, which is a series of event-driven steps. Each step in a workflow is called a state.  
upvoted 3 times

  **Bilalazure** 1 year, 8 months ago



**Selected Answer: D**

Distrubuted\*\*\*\*  
upvoted 2 times

  **Americo32** 1 year, 8 months ago

**Selected Answer: C**


Vou de C, orientada a eventos  
upvoted 2 times

  **MssP** 1 year, 7 months ago

It is true that an Event-driven is made with EventBridge but with a Lambda on schedule??? It is a mismatch, isn't it?  
upvoted 2 times


  **kraken21** 1 year, 7 months ago

Tricky question huh!  
upvoted 3 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: D**



AWS Step functions is serverless Visual workflows for distributed applications  
<https://aws.amazon.com/step-functions/>  
upvoted 2 times

  **leoattf** 1 year, 8 months ago

Besides, "Visualize and develop resilient workflows for EVENT-DRIVEN architectures."  
upvoted 2 times

  **tellmenowwww** 1 year, 8 months ago

Could it be a C because it's event-driven architecture?  
upvoted 3 times

  **SMAZ** 1 year, 8 months ago

Option D..  
AWS Step functions are used for distributed applications  
upvoted 3 times

A company is designing the network for an online multi-player game. The game uses the UDP networking protocol and will be deployed in eight AWS Regions. The network architecture needs to minimize latency and packet loss to give end users a high-quality gaming experience.

Which solution will meet these requirements?

- A. Setup a transit gateway in each Region. Create inter-Region peering attachments between each transit gateway.
- B. Set up AWS Global Accelerator with UDP listeners and endpoint groups in each Region.
- C. Set up Amazon CloudFront with UDP turned on. Configure an origin in each Region.
- D. Set up a VPC peering mesh between each Region. Turn on UDP for each VPC.

Correct Answer: B

Community vote distribution

B (100%)

lucdt4

Highly Voted

1 year, 11 months ago

Selected Answer: B

AWS Global Accelerator = TCP/UDP minimize latency

upvoted 12 times

OAdekunle

Highly Voted

2 years ago

General

Q: What is AWS Global Accelerator?

A: AWS Global Accelerator is a networking service that helps you improve the availability and performance of the applications that you offer to your global users. AWS Global Accelerator is easy to set up, configure, and manage. It provides static IP addresses that provide a fixed entry point to your applications and eliminate the complexity of managing specific IP addresses for different AWS Regions and Availability Zones. AWS Global Accelerator always routes user traffic to the optimal endpoint based on performance, reacting instantly to changes in application health, your user’s location, and policies that you configure. You can test the performance benefits from your location with a speed comparison tool. Like other AWS services, AWS Global Accelerator is a self-service, pay-per-use offering, requiring no long term commitments or minimum fees.

<https://aws.amazon.com/global-accelerator/faqs/>

upvoted 5 times

Danilus

Most Recent

5 months, 3 weeks ago

Selected Answer: B

key-UDP networking protocol

key-minimize latency and packet loss

Transit gateway is to conect multiple vpc

The solution is B because its accept udp protocol and its s global .

The solution is not C because cloudfront just accept http and https protocols

upvoted 2 times

mwwt2022

1 year, 4 months ago

online game -> Global Accelerator

cloudfront is for static/dynamic content caching

upvoted 4 times

Guru4Cloud

1 year, 8 months ago

Selected Answer: B

Set up AWS Global Accelerator with UDP listeners and endpoint groups in each Region.

upvoted 3 times

TariqKipkemei

1 year, 11 months ago

Selected Answer: B

Connect to up to 10 regions within the AWS global network using the AWS Global Accelerator.

upvoted 2 times

TariqKipkemei

1 year, 6 months ago

UDP = Global Accelerator

upvoted 2 times

elearningtakai

2 years, 1 month ago

Selected Answer: B

Global Accelerator supports the User Datagram Protocol (UDP) and Transmission Control Protocol (TCP), making it an excellent choice for an online multi-player game using UDP networking protocol. By setting up Global Accelerator with UDP listeners and endpoint groups in each Region, the network architecture can minimize latency and packet loss, giving end users a high-quality gaming experience.



upvoted 5 times

  **Bofi** 2 years, 2 months ago

**Selected Answer: B**

AWS Global Accelerator is a service that improves the availability and performance of applications with local or global users. Global Accelerator improves performance for a wide range of applications over TCP or UDP by proxying packets at the edge to applications running in one or more AWS Regions. Global Accelerator is a good fit for non-HTTP use cases, such as gaming (UDP), IoT (MQTT), or Voice over IP, as well as for HTTP use cases that specifically require static IP addresses or deterministic, fast regional failover. Both services integrate with AWS Shield for DDoS protection.


upvoted 2 times

  **K0nAn** 2 years, 2 months ago

**Selected Answer: B**

Global Accelerator for UDP and TCP traffic

upvoted 2 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: B**

Global Accelerator

upvoted 2 times

  **Neha999** 2 years, 2 months ago

B

Global Accelerator for UDP traffic

upvoted 2 times

A company hosts a three-tier web application on Amazon EC2 instances in a single Availability Zone. The web application uses a self-managed MySQL database that is hosted on an EC2 instance to store data in an Amazon Elastic Block Store (Amazon EBS) volume. The MySQL database currently uses a 1 TB Provisioned IOPS SSD (io2) EBS volume. The company expects traffic of 1,000 IOPS for both reads and writes at peak traffic.

The company wants to minimize any disruptions, stabilize performance, and reduce costs while retaining the capacity for double the IOPS. The company wants to move the database tier to a fully managed solution that is highly available and fault tolerant.

Which solution will meet these requirements MOST cost-effectively?

- A. Use a Multi-AZ deployment of an Amazon RDS for MySQL DB instance with an io2 Block Express EBS volume.
- B. Use a Multi-AZ deployment of an Amazon RDS for MySQL DB instance with a General Purpose SSD (gp2) EBS volume.
- C. Use Amazon S3 Intelligent-Tiering access tiers.
- D. Use two large EC2 instances to host the database in active-passive mode.

**Correct Answer: B**

*Community vote distribution*

B (92%)

8%

- AlmeroSenior**

Highly Voted

2 years, 2 months ago

Selected Answer: B

RDS does not support IO2 or IO2express . GP2 can do the required IOPS

RDS supported Storage >  
[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)  
GP2 max IOPS >  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/general-purpose.html#gp2-performance>  
upvoted 16 times

**sophieb**

Highly Voted

1 year, 1 month ago

Selected Answer: B

RDS now supports io2 but it might still be an overkill given Gp2 is enough and we are looking for the most cost effective solution.

upvoted 8 times

**EllenLiu**

Most Recent

4 months, 2 weeks ago

Selected Answer: B

gp2: 1,000 GiB => 3000 Baseline performance (IOPS) match the requirement which is 2000 only  
Baseline IOPS performance : 100 ~16,000 at a rate of 3 IOPS per GiB of volume size.

upvoted 2 times

**Guru4Cloud**

1 year, 8 months ago

Selected Answer: B

RDS does not support IO2 or IO2express . GP2 can do the required IOPS

upvoted 4 times

**Gooniegoogoo**

1 year, 10 months ago

The Options is A only because it is sufficient.. Provisioned IOPS are available but overkill.. just want to make sure we understand why its A for the right reason

upvoted 1 times

**dkw2342**

1 year, 2 months ago

Provisioned IOPS are available, but not io2, just io1.

upvoted 1 times

**Abrar2022**

1 year, 10 months ago

Simplified by Almero - thanks.

RDS does not support IO2 or IO2express . GP2 can do the required IOPS

upvoted 2 times



**TariqKipkemei**

1 year, 11 months ago

Selected Answer: B

I tried on the portal and only gp3 and io1 are supported.  
This is 11 May 2023.

upvoted 4 times

  **ruqui** 1 year, 11 months ago

it doesn't matter whether or no io\* is supported, using io2 is overkill, you only need 1K IOPS, B is the correct answer  
upvoted 2 times

  **SimiTik** 2 years ago

A  
Amazon RDS supports the use of Amazon EBS Provisioned IOPS (io2) volumes. When creating a new DB instance or modifying an existing one, you can select the io2 volume type and specify the amount of IOPS and storage capacity required. RDS also supports the newer io2 Block Express volumes, which can deliver even higher performance for mission-critical database workloads.  
upvoted 2 times

  **TariqKipkemei** 1 year, 11 months ago

Impossible. I just tried on the portal and only io1 and gp3 are supported.  
upvoted 2 times

  **klayytech** 2 years, 1 month ago

**Selected Answer: B**

he most cost-effective solution that meets the requirements is to use a Multi-AZ deployment of an Amazon RDS for MySQL DB instance with a General Purpose SSD (gp2) EBS volume. This solution will provide high availability and fault tolerance while minimizing disruptions and stabilizing performance. The gp2 EBS volume can handle up to 16,000 IOPS. You can also scale up to 64 TiB of storage.

Amazon RDS for MySQL provides automated backups, software patching, and automatic host replacement. It also provides Multi-AZ deployments that automatically replicate data to a standby instance in another Availability Zone. This ensures that data is always available even in the event of a failure.  
upvoted 2 times

  **test\_devops\_aws** 2 years, 1 month ago

**Selected Answer: B**

RDS does not support io2 !!!  
upvoted 1 times

  **Maximus007** 2 years, 1 month ago

B:gp3 would be the better option, but considering we have only gp2 option and such storage volume - gp2 will be the right choice  
upvoted 4 times



  **Nel8** 2 years, 1 month ago

**Selected Answer: B**

I thought the answer here is A. But when I found the link from Amazon website; as per AWS:

Amazon RDS provides three storage types: General Purpose SSD (also known as gp2 and gp3), Provisioned IOPS SSD (also known as io1), and magnetic (also known as standard). They differ in performance characteristics and price, which means that you can tailor your storage performance and cost to the needs of your database workload. You can create MySQL, MariaDB, Oracle, and PostgreSQL RDS DB instances with up to 64 terabytes (TiB) of storage. You can create SQL Server RDS DB instances with up to 16 TiB of storage. For this amount of storage, use the Provisioned IOPS SSD and General Purpose SSD storage types.

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)  
upvoted 1 times

  **JA2018** 5 months, 2 weeks ago

[updated - 21 Nove 2024]

Amazon RDS provides three storage types: Provisioned IOPS SSD (also known as io1 and io2 Block Express), General Purpose SSD (also known as gp2 and gp3), and magnetic (also known as standard). They differ in performance characteristics and price, which means that you can tailor your storage performance and cost to the needs of your database workload. You can create Db2, MySQL, MariaDB, Oracle, SQL Server, and PostgreSQL RDS DB instances with up to 64 terabytes (TiB) of storage. RDS for Db2 doesn't support the gp2 and magnetic storage types.  
upvoted 1 times

  **Steve\_4542636** 2 years, 2 months ago

**Selected Answer: B**


for DB instances between 1 TiB and 4 TiB, storage is striped across four Amazon EBS volumes providing burst performance of up to 12,000 IOPS.

from "[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)"  
upvoted 2 times

  **TungPham** 2 years, 2 months ago

**Selected Answer: B**



[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)  
Amazon RDS provides three storage types: General Purpose SSD (also known as gp2 and gp3), Provisioned IOPS SSD (also known as io1), and magnetic (also known as standard)  
B - MOST cost-effectively  
upvoted 4 times

  **KZM** 2 years, 2 months ago

The baseline IOPS performance of gp2 volumes is 3 IOPS per GB, which means that a 1 TB gp2 volume will have a baseline performance of 3,000 IOPS. However, the volume can also burst up to 16,000 IOPS for short periods, but this burst performance is limited and may not be sustained for long durations.

So, I am more prefer option A.

upvoted 1 times

  **KZM** 2 years, 2 months ago

If a 1 TB gp3 EBS volume is used, the maximum available IOPS according to calculations is 3000. This means that the storage can support a requirement of 1000 IOPS, and even 2000 IOPS if the requirement is doubled.

I am confusing between choosing A or B.

upvoted 2 times

  **mark16dc** 2 years, 2 months ago

**Selected Answer: A**

Option A is the correct answer. A Multi-AZ deployment provides high availability and fault tolerance by automatically replicating data to a standby instance in a different Availability Zone. This allows for seamless failover in the event of a primary instance failure. Using an io2 Block Express EBS volume provides the needed IOPS performance and capacity for the database. It is also designed for low latency and high durability, which makes it a good choice for a database tier.

upvoted 1 times

  **CapJackSparrow** 2 years, 1 month ago

How will you select io2 when RDS only offers io1.....magic?

upvoted 1 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: B**

Correction - hit wrong answer button - meant 'B'

Amazon RDS provides three storage types: General Purpose SSD (also known as gp2 and gp3), Provisioned IOPS SSD (also known as io1)

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)

upvoted 2 times

A company hosts a serverless application on AWS. The application uses Amazon API Gateway, AWS Lambda, and an Amazon RDS for PostgreSQL database. The company notices an increase in application errors that result from database connection timeouts during times of peak traffic or unpredictable traffic. The company needs a solution that reduces the application failures with the least amount of change to the code.

What should a solutions architect do to meet these requirements?

- A. Reduce the Lambda concurrency rate.
- B. Enable RDS Proxy on the RDS DB instance.
- C. Resize the RDS DB instance class to accept more connections.
- D. Migrate the database to Amazon DynamoDB with on-demand scaling.

**Correct Answer: B**

*Community vote distribution*

B (100%)

- TariqKipkemei**

Highly Voted

1 year, 11 months ago

Selected Answer: B

Many applications, including those built on modern serverless architectures, can have a large number of open connections to the database server and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%.

<https://aws.amazon.com/rds/proxy/>

upvoted 11 times
- Murtadhaceit**

Highly Voted

1 year, 4 months ago

Selected Answer: B

A. Reduce the Lambda concurrency rate? Has nothing to do with decreasing connections timeout.

B. Enable RDS Proxy on the RDS DB instance. Correct answer

C. Resize the RDS DB instance class to accept more connections? More connections means worse performance. Therefore, not correct.

D. Migrate the database to Amazon DynamoDB with on-demand scaling? DynamoDB is a noSQL database. Not correct.

upvoted 6 times
- JA2018**

Most Recent

5 months, 2 weeks ago

Selected Answer: B

Keys found in STEM:

  1. The company notices an increase in application errors that result from database connection timeouts during times of peak traffic or unpredictable traffic.
  2. The company needs a solution that reduces the application failures with the least amount of change to the code.

upvoted 1 times
- Guru4Cloud**

1 year, 8 months ago

Selected Answer: B

RDS Proxy is a fully managed, highly available, and scalable proxy for Amazon Relational Database Service (RDS) that makes it easy to connect to your RDS instances from applications running on AWS Lambda. RDS Proxy offloads the management of connections to the database, which can help to improve performance and reliability.

upvoted 4 times
- elearningtakai**

2 years, 1 month ago

Selected Answer: B

To reduce application failures resulting from database connection timeouts, the best solution is to enable RDS Proxy on the RDS DB instance

upvoted 2 times
- WherecanIstart**

2 years, 1 month ago

Selected Answer: B

RDS Proxy

upvoted 4 times
- nder**

2 years, 2 months ago

Selected Answer: B

RDS Proxy will pool connections, no code changes need to be made

upvoted 2 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: B**

RDS proxy

upvoted 2 times

  **Neha999** 2 years, 2 months ago

B RDS Proxy

<https://aws.amazon.com/rds/proxy/>

upvoted 3 times

A company is migrating an old application to AWS. The application runs a batch job every hour and is CPU intensive. The batch job takes 15 minutes on average with an on-premises server. The server has 64 virtual CPU (vCPU) and 512 GiB of memory.

Which solution will run the batch job within 15 minutes with the LEAST operational overhead?

- A. Use AWS Lambda with functional scaling.
- B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
- C. Use Amazon Lightsail with AWS Auto Scaling.
- D. Use AWS Batch on Amazon EC2.

Correct Answer: D

Community vote distribution

D (98%)

  **NolaH0la** Highly Voted 2 years, 2 months ago

The amount of CPU and memory resources required by the batch job exceeds the capabilities of AWS Lambda and Amazon Lightsail with AWS Auto Scaling, which offer limited compute resources. AWS Fargate offers containerized application orchestration and scalable infrastructure, but may require additional operational overhead to configure and manage the environment. AWS Batch is a fully managed service that automatically provisions the required infrastructure for batch jobs, with options to use different instance types and launch modes.

Therefore, the solution that will run the batch job within 15 minutes with the LEAST operational overhead is D. Use AWS Batch on Amazon EC2. AWS Batch can handle all the operational aspects of job scheduling, instance management, and scaling while using Amazon EC2

upvoted 22 times

  **everfly** Highly Voted 2 years, 2 months ago

Selected Answer: D

AWS Batch is a fully-managed service that can launch and manage the compute resources needed to execute batch jobs. It can scale the compute environment based on the size and timing of the batch jobs.


upvoted 12 times

  **FlyingHawk** Most Recent 3 months, 1 week ago

Selected Answer: D

Lambda has a hard limit of 10 GB of memory and 6 vCPUs per function execution, Amazon Light sail has the limit of 8vCPUs and 32GB memory, so A and C out, B and D can meet vCPU and memory requirements, but B is more expensive than D, so correct answer is D.

upvoted 1 times

  **Danilus** 5 months, 3 weeks ago

Selected Answer: D

key-The batch job takes 15 minutes on average



key-LEAST operational overhead

key-the batch job

is not A because lambda execution takes 15 max and lambda dont support 512 GIB of memory

the answer is D batch it is designed to run large-scale batch jobs an automatically manages the scaling of resources also batch allow jobs to be distributed accross multiple instances because supports parallel execution

upvoted 3 times

  **Ramdi1** 1 year, 7 months ago

Selected Answer: D

The question needs to be phrased differently. I assume at first it was Lambda, because it says 15 minutes in the question which can be done. Yes it also does say CPU intensive however they go on with a full stop and then give you the server specs. It does not say it uses that much of the specs so they need to really rephrase the questions.

upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

Selected Answer: D

The main reasons are:

AWS Batch can easily schedule and run batch jobs on EC2 instances. It can scale up to the required vCPUs and memory to match the on-premises server.

Using EC2 provides full control over the instance type to meet the resource needs.

No servers or clusters to manage like with ECS/Fargate or Lightsail. AWS Batch handles this automatically.

More cost effective and operationally simple compared to Lambda which is not ideal for long running batch jobs.

upvoted 7 times

🗳️ 👤 **BrijMohan08** 1 year, 8 months ago

**Selected Answer: A**

On-Prem was avg 15 min, but target state architecture is expected to finish within 15 min  
upvoted 1 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

How? The on-prem server has 64 CPUs and 512 GB RAM, Lambda offers much less. And even on-prem it takes 15 minutes ON AVERAGE, sometimes more.  
upvoted 5 times

🗳️ 👤 **jayce5** 1 year, 9 months ago

**Selected Answer: D**

Not Lambda, "average 15 minutes" means there are jobs with running more and less than 15 minutes. Lambda max is 15 minutes.  
upvoted 3 times

🗳️ 👤 **Gooniegoogoo** 1 year, 10 months ago

This is for certain a tough one. I do see that they have thrown a curve ball in making it Lambda Functional scaling, however what we dont know is if this application has many request or one large one.. looks like Lambda can scale and use the same lambda env.. seems intensive tho so will go with D  
upvoted 5 times

🗳️ 👤 **TariqKipkemei** 1 year, 11 months ago

**Selected Answer: D**

AWS Batch  
upvoted 3 times

🗳️ 👤 **JLII** 2 years, 1 month ago

**Selected Answer: D**

Not A because: "AWS Lambda now supports up to 10 GB of memory and 6 vCPU cores for Lambda Functions." <https://aws.amazon.com/about-aws/whats-new/2020/12/aws-lambda-supports-10gb-memory-6-vcpu-cores-lambda-functions/> vs. "The server has 64 virtual CPU (vCPU) and 512 GiB of memory" in the question.  
upvoted 7 times

🗳️ 👤 **geekgirl22** 2 years, 2 months ago

A is the answer. Lambda is known that has a limit of 15 minutes. So for as long as it says "within 15 minutes" that should be a clear indication it is Lambda  
upvoted 2 times

🗳️ 👤 **nder** 2 years, 2 months ago

Wrong, the job takes "On average 15 minutes" and requires more cpu and ram than lambda can deal with. AWS Batch is correct in this scenario  
upvoted 4 times

🗳️ 👤 **geekgirl22** 2 years, 2 months ago

read the rest of the question which gives the answer:  
"Which solution will run the batch job within 15 minutes with the LEAST operational overhead?"  
Keyword "Within 15 minutes"  
upvoted 3 times

🗳️ 👤 **Lonojack** 2 years, 2 months ago

What happens if it EXCEEDS the 15 min AVERAGE?  
Average = possibly can be more than 15min.  
The safer bet would be option D: AWS Batch on EC2  
upvoted 8 times

🗳️ 👤 **Terion** 1 year, 7 months ago

I think what he means is that it takes on average 15 min on prem only  
upvoted 2 times

🗳️ 👤 **awsgeek75** 1 year, 3 months ago

How are you going to get 64 vCPUS to a Lambda function?  
upvoted 2 times

🗳️ 👤 **bdp123** 2 years, 2 months ago

**Selected Answer: D**

AWS batch on EC2  
upvoted 2 times

A company stores its data objects in Amazon S3 Standard storage. A solutions architect has found that 75% of the data is rarely accessed after 30 days. The company needs all the data to remain immediately accessible with the same high availability and resiliency, but the company wants to minimize storage costs.

Which storage solution will meet these requirements?

- A. Move the data objects to S3 Glacier Deep Archive after 30 days.
- B. Move the data objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- C. Move the data objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.
- D. Move the data objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately.

**Correct Answer: B**

*Community vote distribution*

B (100%)

- Lonojack**

Highly Voted

2 years, 2 months ago

Selected Answer: B

Needs immediate accessibility after 30days, IF they need to be accessed.  
upvoted 6 times
- Piccalo**

Highly Voted

2 years, 1 month ago

Highly available so One Zone IA is out the question  
Glacier Deep archive isn't immediately accessible 12-48 hours  
B is the answer.  
upvoted 5 times
- Danilus**

Most Recent

5 months, 3 weeks ago

Selected Answer: B

Key - high availability and resiliency  
Key-the data to remain immediately accessible  
  
Its not C or D because its One zone  
its no B even if glacier is cheaper the problem is that data must be immediately accesible so the answer is B  
upvoted 2 times
- Aru\_1994**

9 months, 3 weeks ago

Option B  
upvoted 2 times
- Apexakil1996**

1 year, 4 months ago

One -zone -infrequent access cannot be the answer because it requires high availability so standard infrequent access should be the answer  
upvoted 5 times
- TariqKipkemei**

1 year, 6 months ago

Selected Answer: B

high availability, resiliency = multi AZ  
75% of the data is rarely accessed but remain immediately accessible = Standard-Infrequent Access  
upvoted 4 times
- Guru4Cloud**

1 year, 8 months ago

Selected Answer: B

The correct answer is B.  
  
S3 Standard-IA is a storage class that is designed for infrequently accessed data. It offers lower storage costs than S3 Standard, but it has a retrieval latency of 1-5 minutes.  
upvoted 4 times
- elearningtakai**

2 years, 1 month ago

Selected Answer: B

S3 Glacier Deep Archive is intended for data that is rarely accessed and can tolerate retrieval times measured in hours. Moving data to S3 One Zone-IA immediately would not meet the requirement of immediate accessibility with the same high availability and resiliency.  
upvoted 2 times

  **KS2020** 2 years, 1 month ago

The answer should be C.

S3 One Zone-IA is for data that is accessed less frequently but requires rapid access when needed. Unlike other S3 Storage Classes which store data in a minimum of three Availability Zones (AZs), S3 One Zone-IA stores data in a single AZ and costs 20% less than S3 Standard-IA.

<https://aws.amazon.com/s3/storage-classes/#:~:text=S3%20One%20Zone%2DIA%20is,less%20than%20S3%20Standard%2DIA>.

upvoted 1 times

  **shanwford** 2 years, 1 month ago

The Question emphasises to kepp same high availability class - S3 One Zone-IA doesnt support multiple Availability Zone data resilience model like S3 Standard-Infrequent Access.

upvoted 3 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: B**

S3 Standard-Infrequent Access after 30 days

upvoted 3 times

  **NolaH0la** 2 years, 2 months ago

B

Option B - Move the data objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days - will meet the requirements of keeping the data immediately accessible with high availability and resiliency, while minimizing storage costs. S3 Standard-IA is designed for infrequently accessed data, and it provides a lower storage cost than S3 Standard, while still offering the same low latency, high throughput, and high durability as S3 Standard.

upvoted 5 times

A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.



Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)


- A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge.
- B. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
- C. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
- D. Store the server-side code on Amazon FSx for Windows File Server. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
- E. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volume. Mount the EBS volume on each EC2 instance to share the files.

**Correct Answer:** AD

*Community vote distribution*

AD (100%)

-   **Steve\_4542636**

Highly Voted 



 2 years, 2 months ago

Selected Answer: AD



A because ElastiCache, despite being ideal for leaderboards per Amazon, doesn't cache at edge locations. D because FSx has higher performance for low latency needs.

<https://www.techtarget.com/searchaws/tip/Amazon-FSx-vs-EFS-Compare-the-AWS-file-services>



"FSx is built for high performance and submillisecond latency using solid-state drive storage volumes. This design enables users to select storage capacity and latency independently. Thus, even a subterabyte file system can have 256 Mbps or higher throughput and support volumes up to 64 TB."


upvoted 7 times
-   **Nel8** 2 years, 1 month ago

Just to add, ElastiCache is use in front of AWS database.

upvoted 2 times
-   **baba365** 1 year, 7 months ago

Why not EFS?

upvoted 1 times
-   **Guru4Cloud**



Highly Voted 


 1 year, 8 months ago

Selected Answer: AD

The reasons are:

Storing static files in S3 with CloudFront provides durability, high availability, and low latency by caching at edge locations. FSx for Windows File Server provides a fully managed Windows native file system that can be accessed from the Windows EC2 instances to share server-side code. It is designed for high availability and scales up to 10s of GBPS throughput. EFS and EBS volumes can be attached to a single AZ. FSx and S3 are replicated across AZs for high availability.

upvoted 6 times
-   **Danilus**

Most Recent 

 5 months, 3 weeks ago



Selected Answer: AD

Key: Static files and dynamic server-side code.  
Key: Windows Server.

The answer is A because S3 is for static files, and CloudFront is a CDN.

The second answer is D because FSx works for Windows File Server.

It's not C because EFS is just for Linux.  
It's not E because EBS doesn't share files between EC2 instances.

upvoted 2 times
-   **rodrigoleoncio** 11 months, 2 weeks ago

Selected Answer: AD

A because Elasticache doesn't cache at edge locations. D because FSx has higher performance for low latency needs.  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

The question and options are badly worded. How does (D) storing server side code on a file server makes it executable?  
upvoted 2 times

  **4fad2f8** 1 year, 3 months ago

you can't mount efs on windows  
upvoted 3 times



  **WherecanIstart** 2 years, 1 month ago

**Selected Answer: AD**

A & D for sure  
upvoted 5 times

  **KZM** 2 years, 2 months ago

It is obvious that A and D.  
upvoted 2 times

  **bdp123** 2 years, 2 months ago

**Selected Answer: AD**

both A and D seem correct  
upvoted 2 times

  **NolaH0la** 2 years, 2 months ago

A and D seems correct  
upvoted 2 times

A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Install an external image management library on an EC2 instance. Use the image management library to process the images.
- B. Create a CloudFront origin request policy. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
- C. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
- D. Create a CloudFront response headers policy. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

Correct Answer: C

Community vote distribution

C (88%)

12%

  **NolaH0la** Highly Voted 1 year, 8 months ago

Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.

Using a Lambda@Edge function with an external image management library is the best solution to resize the images dynamically and serve appropriate formats to clients. Lambda@Edge is a serverless computing service that allows running custom code in response to CloudFront events, such as viewer requests and origin requests. By using a Lambda@Edge function, it's possible to process images on the fly and modify the CloudFront response before it's sent back to the client. Additionally, Lambda@Edge has built-in support for external libraries that can be used to process images. This approach will reduce operational overhead and scale automatically with traffic.

upvoted 21 times

  **TariqKipkemei** Highly Voted 1 year ago

**Selected Answer: C**

The moment there is a need to implement some logic at the CDN think Lambda@Edge.

upvoted 10 times

  **Guru4Cloud** Most Recent 1 year, 2 months ago

**Selected Answer: C**

The correct answer is C.

A Lambda@Edge function is a serverless function that runs at the edge of the CloudFront network. This means that the function is executed close to the user, which can improve performance.

An external image management library can be used to resize images and to serve the appropriate format. Associating the Lambda@Edge function with the CloudFront behaviors that serve the images ensures that the function is executed for all requests that are served by those behaviors.

upvoted 4 times

  **BrijMohan08** 1 year, 2 months ago

**Selected Answer: B**

If the user asks for the most optimized image format (JPEG, WebP, or AVIF) using the directive format=auto, CloudFront Function will select the best format based on the Accept header present in the request.

Latest documentation: <https://aws.amazon.com/blogs/networking-and-content-delivery/image-optimization-using-amazon-cloudfront-and-aws-lambda/>

upvoted 3 times

  **pentium75** 10 months, 1 week ago

But a policy alone cannot resize images.


upvoted 1 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: C**

<https://aws.amazon.com/cn/blogs/networking-and-content-delivery/resizing-images-with-amazon-cloudfront-lambdaedge-aws-cdn-blog/>

upvoted 5 times

  **everfly** 1 year, 8 months ago

**Selected Answer: C**

<https://aws.amazon.com/cn/blogs/networking-and-content-delivery/resizing-images-with-amazon-cloudfront-lambdaedge-aws-cdn-blog/>  
upvoted 3 times

A hospital needs to store patient records in an Amazon S3 bucket. The hospital’s compliance team must ensure that all protected health information (PHI) is encrypted in transit and at rest. The compliance team must administer the encryption key for data at rest.

Which solution will meet these requirements?



- A. Create a public SSL/TLS certificate in AWS Certificate Manager (ACM). Associate the certificate with Amazon S3. Configure default encryption for each S3 bucket to use server-side encryption with AWS KMS keys (SSE-KMS). Assign the compliance team to manage the KMS keys.
- B. Use the aws:SecureTransport condition on S3 bucket policies to allow only encrypted connections over HTTPS (TLS). Configure default encryption for each S3 bucket to use server-side encryption with S3 managed encryption keys (SSE-S3). Assign the compliance team to manage the SSE-S3 keys.
- C. Use the aws:SecureTransport condition on S3 bucket policies to allow only encrypted connections over HTTPS (TLS). Configure default encryption for each S3 bucket to use server-side encryption with AWS KMS keys (SSE-KMS). Assign the compliance team to manage the KMS keys.
- D. Use the aws:SecureTransport condition on S3 bucket policies to allow only encrypted connections over HTTPS (TLS). Use Amazon Macie to protect the sensitive data that is stored in Amazon S3. Assign the compliance team to manage Macie.

**Correct Answer: C**

*Community vote distribution*



C (89%)

9%



-   **NolaH0la** Highly Voted 1 year, 8 months ago

Option C is correct because it allows the compliance team to manage the KMS keys used for server-side encryption, thereby providing the necessary control over the encryption keys. Additionally, the use of the "aws:SecureTransport" condition on the bucket policy ensures that all connections to the S3 bucket are encrypted in transit.

option B might be misleading but using SSE-S3, the encryption keys are managed by AWS and not by the compliance team

upvoted 28 times
-   **Lonojack** 1 year, 8 months ago

Perfect explanation. I Agree



upvoted 5 times
-   **pentium75** Highly Voted 10 months, 1 week ago

**Selected Answer: C**

Not A, Certificate Manager has nothing to do with S3

Not B, SSE-S3 does not allow compliance team to manage the key

Not D, Macie is for identifying sensitive data, not protecting it



upvoted 13 times
-   **Guru4Cloud** Most Recent 1 year, 2 months ago

**Selected Answer: C**

Macie does not encrypt the data like the question is asking

<https://docs.aws.amazon.com/macie/latest/user/what-is-macie.html>



Also, SSE-S3 encryption is fully managed by AWS so the Compliance Team can't administer this.

upvoted 3 times
-   **Yadav\_Sanjay** 1 year, 5 months ago

**Selected Answer: C**



D - Can't be because - Amazon Macie is a data security service that uses machine learning (ML) and pattern matching to discover and help protect your sensitive data.

Macie discovers sensitive information, can help in protection but can't protect

upvoted 3 times
-   **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: C**

B can work if they do not want control over encryption keys.

upvoted 2 times
-   **Russs99** 1 year, 7 months ago

**Selected Answer: A**

Option A proposes creating a public SSL/TLS certificate in AWS Certificate Manager and associating it with Amazon S3. This step ensures that data is encrypted in transit. Then, the default encryption for each S3 bucket will be configured to use server-side encryption with AWS KMS keys (SSE-KMS), which will provide encryption at rest for the data stored in S3. In this solution, the compliance team will manage the KMS keys, ensuring that they control the encryption keys for data at rest.

upvoted 1 times

  **pentium75** 10 months, 1 week ago



ACM is for website certificates, has nothing to do with S3.

upvoted 3 times

  **Shrestwt** 1 year, 6 months ago

ACM cannot be integrated with Amazon S3 bucket directly.

upvoted 3 times

  **Bofi** 1 year, 7 months ago

**Selected Answer: C**

Option C seems to be the correct answer, option A is also close but ACM cannot be integrated with Amazon S3 bucket directly, hence, u can not attached TLS to S3. You can only attached TLS certificate to ALB, API Gateway and CloudFront and maybe Global Accelerator but definitely NOT EC2 instance and S3 bucket



upvoted 2 times

  **CapJackSparrow** 1 year, 7 months ago

**Selected Answer: C**

D makes no sense.

upvoted 3 times

  **Dody** 1 year, 7 months ago

**Selected Answer: C**

Correct Answer is "C"

"D" is not correct because Amazon Macie securely stores your data at rest using AWS encryption solutions. Macie encrypts data, such as findings, using an AWS managed key from AWS Key Management Service (AWS KMS). However, in the question there is a requirement that the compliance team must administer the encryption key for data at rest.

<https://docs.aws.amazon.com/macie/latest/user/data-protection.html>

upvoted 3 times

  **cegama543** 1 year, 8 months ago

**Selected Answer: C**

Option C will meet the requirements.

Explanation:

The compliance team needs to administer the encryption key for data at rest in order to ensure that protected health information (PHI) is encrypted in transit and at rest. Therefore, we need to use server-side encryption with AWS KMS keys (SSE-KMS). The default encryption for each S3 bucket can be configured to use SSE-KMS to ensure that all new objects in the bucket are encrypted with KMS keys.

Additionally, we can configure the S3 bucket policies to allow only encrypted connections over HTTPS (TLS) using the aws:SecureTransport condition. This ensures that the data is encrypted in transit.



upvoted 2 times

  **Karlos99** 1 year, 8 months ago

**Selected Answer: C**

We must provide encrypted in transit and at rest. Macie is needed to discover and recognize any PII or Protected Health Information. We already know that the hospital is working with the sensitive data ) so protect them with KMS and SSL. Answer D is unnecessary

upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago



**Selected Answer: C**

Macie does not encrypt the data like the question is asking

<https://docs.aws.amazon.com/macie/latest/user/what-is-macie.html>

Also, SSE-S3 encryption is fully managed by AWS so the Compliance Team can't administer this.

upvoted 3 times



  **Abhineet9148232** 1 year, 8 months ago

**Selected Answer: C**

C [Correct]: Ensures Https only traffic (encrypted transit), Enables compliance team to govern encryption key.

D [Incorrect]: Misleading; PHI is required to be encrypted not discovered. Maice is a discovery service. (<https://aws.amazon.com/macie/>)

upvoted 5 times

  **Nel8** 1 year, 8 months ago

**Selected Answer: D**

Correct answer should be D. "Use Amazon Macie to protect the sensitive data..."

As requirement says "The hospitals's compliance team must ensure that all protected health information (PHI) is encrypted in transit and at rest."

Macie protects personal record such as PHI. Macie provides you with an inventory of your S3 buckets, and automatically evaluates and monitors

the buckets for security and access control. If Macie detects a potential issue with the security or privacy of your data, such as a bucket that becomes publicly accessible, Macie generates a finding for you to review and remediate as necessary.

upvoted 4 times

  **Drayen25** 1 year, 8 months ago

Option C should be  
upvoted 3 times

A company uses Amazon API Gateway to run a private gateway with two REST APIs in the same VPC. The BuyStock RESTful web service calls the CheckFunds RESTful web service to ensure that enough funds are available before a stock can be purchased. The company has noticed in the VPC flow logs that the BuyStock RESTful web service calls the CheckFunds RESTful web service over the internet instead of through the VPC. A solutions architect must implement a solution so that the APIs communicate through the VPC.

Which solution will meet these requirements with the FEWEST changes to the code?

- A. Add an X-API-Key header in the HTTP header for authorization.
- B. Use an interface endpoint.
- C. Use a gateway endpoint.
- D. Add an Amazon Simple Queue Service (Amazon SQS) queue between the two REST APIs.

**Correct Answer: B**

*Community vote distribution*

B (94%)

6%

- everfly**

Highly Voted

1 year, 8 months ago

Selected Answer: B

an interface endpoint is a horizontally scaled, redundant VPC endpoint that provides private connectivity to a service. It is an elastic network interface with a private IP address that serves as an entry point for traffic destined to the AWS service. Interface endpoints are used to connect VPCs with AWS services

upvoted 23 times
- lucdt4**

Highly Voted

1 year, 5 months ago

Selected Answer: B

C. Use a gateway endpoint is wrong because gateway endpoints only support for S3 and dynamoDB, so B is correct

upvoted 12 times
- zdi561**

Most Recent

3 months ago

Selected Answer: B

AWS PrivateLink connects resources within one VPC (the consumer) to services within another VPC (the provider) using private IP addresses. But since it uses API gateway, you have to use VPC endpoint to keep private communication though there are other options to avoid using VPC endpoint

upvoted 1 times
- Rcosmos**

3 months, 1 week ago

Selected Answer: B

B. Usar um endpoint de interface:

Embora os endpoints de interface também possam facilitar a comunicação privada, eles são projetados para integrar serviços AWS ou serviços baseados em VPC endpoints. Para o caso de APIs REST privadas do API Gateway, o endpoint de gateway é mais apropriado.

upvoted 1 times
- Rcosmos**

3 months, 1 week ago

Selected Answer: B

Por que o endpoint de gateway é a solução ideal?  
Comunicação por meio da VPC:Um endpoint de gateway no Amazon API Gateway permite que os serviços na mesma VPC se comuniquem diretamente, sem sair da rede interna ou passar pela Internet pública.  
Menor impacto no código:O uso de um endpoint de gateway no API Gateway requer apenas a reconfiguração do endpoint na VPC, sem mudanças significativas no código das APIs.  
Segurança:Com um endpoint de gateway, o tráfego é mantido dentro da rede privada da VPC, reduzindo os riscos de segurança associados ao tráfego pela Internet.  
Custo-benefício:Um endpoint de gateway é econômico, pois utiliza a infraestrutura existente da VPC e evita custos adicionais relacionados ao tráfego pela Internet.

upvoted 1 times
- meowruki**

11 months, 1 week ago

Selected Answer: B

B. Use an interface endpoint.

Here's the reasoning:

Interface Endpoint (Option B): An interface endpoint (also known as VPC endpoint) allows communication between resources in your VPC and

services without traversing the public internet. In this case, you can create an interface endpoint for API Gateway in your VPC. This enables the communication between the BuyStock and CheckFunds RESTful web services within the VPC, and it doesn't require significant changes to the code.

X-API-Key header (Option A): Adding an X-API-Key header for authorization doesn't address the issue of ensuring that the APIs communicate through the VPC. It's more related to authentication and authorization mechanisms.

upvoted 5 times

  **liux99** 12 months ago

The question here is that the BuyStock RESTful web service calls the CheckFunds RESTful web service through API gateway (internet), not directly. How does API gateway connect the services BuyStock and CheckFunds? It connects the Interface Endpoint of the services through Privatelink. The interface endpoints provide direct connection between services within the same private subnet. Answer B is correct.

upvoted 3 times

  **yodelin** 1 year ago

how is it even possible, I mean if it's private and both are in the same VPC then we shouldn't even have such an issue right?

upvoted 4 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

B. Use an interface endpoint.

upvoted 2 times

  **envest** 1 year, 5 months ago

Answer B (from abylead)

With API GW, you can create multiple prv REST APIs, only accessible with an interface VPC endpt. To allow/ deny simple or cross acc access to your API from selected VPCs & its endpts, you use resource plcys. In addition, you can also use DX for a connection between onprem network to VPC or your prv API.

API GW to VPC: <https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-private-apis.html>


Less correct & incorrect (infeasible & inadequate) answers:

A)X-API-Key in HTTP header for authorization needs auto-process fcts & changes: inadequate.

C)VPC GW endpts for S3 or DynamDB aren't for RESTful svcs: infeasible.

D)SQS que between 2 REST APIs needs endpts & some changes: inadequate.

upvoted 2 times

  **aqmdla2002** 1 year, 5 months ago

**Selected Answer: C**


I select C because it's the solution with the " FEWEST changes to the code"

upvoted 1 times

  **pentium75** 10 months, 1 week ago

Gateway Endpoint can provide access to S3 or DynamoDB, not to API Gateway

upvoted 2 times

  **awsgeek75** 9 months, 2 weeks ago

Fewest changes to the code doesn't mean break the code by doing something irrelevant. Gateway endpoint is for S3 and DynamoDB

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: B**

An interface endpoint is powered by PrivateLink, and uses an elastic network interface (ENI) as an entry point for traffic destined to the service

upvoted 3 times

  **kprakashbehera** 1 year, 7 months ago

**Selected Answer: B**

BBBBBB

upvoted 2 times

  **siyam008** 1 year, 8 months ago

**Selected Answer: C**

<https://www.linkedin.com/pulse/aws-interface-endpoint-vs-gateway-alex-chang>

upvoted 1 times

  **siyam008** 1 year, 8 months ago

Correct answer is B. Incorrectly selected C

upvoted 3 times

  **DASBOL** 1 year, 8 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-private-apis.html>

upvoted 5 times

  **Sherif\_Abbas** 1 year, 8 months ago

**Selected Answer: C**

The only time where an Interface Endpoint may be preferable (for S3 or DynamoDB) over a Gateway Endpoint is if you require access from on-premises, for example you want private access from your on-premise data center

upvoted 2 times

  **Steve\_4542636** 1 year, 8 months ago

The RESTful services is neither an S3 or DynamDB service, so a VPC Gateway endpoint isn't available here.

upvoted 5 times

  **bdp123** 1 year, 8 months ago

**Selected Answer: B**

fewest changes to code and below link:

<https://gkzz.medium.com/what-is-the-differences-between-vpc-endpoint-gateway-endpoint-ae97bfab97d8>

upvoted 3 times

  **PoisonBlack** 1 year, 6 months ago

This really helped me understand the difference between the two. Thx

upvoted 2 times

A company hosts a multiplayer gaming application on AWS. The company wants the application to read data with sub-millisecond latency and run one-time queries on historical data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon RDS for data that is frequently accessed. Run a periodic custom script to export the data to an Amazon S3 bucket.
- B. Store the data directly in an Amazon S3 bucket. Implement an S3 Lifecycle policy to move older data to S3 Glacier Deep Archive for long-term storage. Run one-time queries on the data in Amazon S3 by using Amazon Athena.
- C. Use Amazon DynamoDB with DynamoDB Accelerator (DAX) for data that is frequently accessed. Export the data to an Amazon S3 bucket by using DynamoDB table export. Run one-time queries on the data in Amazon S3 by using Amazon Athena.
- D. Use Amazon DynamoDB for data that is frequently accessed. Turn on streaming to Amazon Kinesis Data Streams. Use Amazon Kinesis Data Firehose to read the data from Kinesis Data Streams. Store the records in an Amazon S3 bucket.

**Correct Answer: C**

*Community vote distribution*

C (100%)

- lexotan**

Highly Voted

2 years ago

Selected Answer: C

would be nice to have an explanation on why examtopic selects its answers.

upvoted 12 times
- ale\_brd\_111**

1 year, 4 months ago

exam topic does not select anything, these are questions from the free forum topics, the only thing exam topic does is to aggregate them all under one single point of view and if you pay you get to see them all aggregated else you can still scroll topic by topic for free

upvoted 7 times
- TariqKipkemei**

Highly Voted

1 year, 6 months ago

Selected Answer: C

DAX delivers up to a 10 times performance improvement—from milliseconds to microseconds.

Using DynamoDB export to S3, you can export data from an Amazon DynamoDB table to an Amazon S3 bucket. This feature enables you to perform analytics and complex queries on your data using other AWS services such as Athena, AWS Glue.

upvoted 9 times
- Omshanti**

Most Recent

7 months, 1 week ago

Selected Answer: C

test test

upvoted 2 times
- sandordini**

1 year ago

Selected Answer: C

Sub-millisecond: DynamoDB (DAX), onetime query, Least operational overhead: Athena

upvoted 4 times
- Uzbekistan**

1 year, 2 months ago

Selected Answer: C

Dynamo DB + DAX = low latency.

upvoted 6 times
- fabiomarrocolo**

1 year, 2 months ago

Scusate io ho pagato contributor perchè vedo ancora + votati invece di vedere solo la risposta corretta? Grazie.Fabio

upvoted 2 times
- LoXoL**

1 year, 2 months ago



Vedrai sempre e comunque sia la risposta della community ("Most Voted") che la risposta degli admin (rettangolo verde). Occhio perche' la risposta degli admin non sempre e' corretta.

upvoted 1 times
- Mikado211**

1 year, 4 months ago

Sub-millisecond latency == DAX

upvoted 4 times

  **Mikado211** 1 year, 4 months ago  
So C !  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**



Amazon DynamoDB with DynamoDB Accelerator (DAX) is a fully managed, in-memory caching solution for DynamoDB. DAX can improve the performance of DynamoDB by up to 10x. This makes it a good choice for data that needs to be accessed with sub-millisecond latency. DynamoDB table export allows you to export data from DynamoDB to an S3 bucket. This can be useful for running one-time queries on historical data. Amazon Athena is a serverless, interactive query service that makes it easy to analyze data in Amazon S3. Athena can be used to run one-time queries on the data in the S3 bucket.  
upvoted 6 times

  **aaroncelestin** 1 year, 8 months ago

A NoSQL isn't even mentioned in the question and yet we are supposed to just imagine this fictional customer is using a NoSql DB  
upvoted 3 times

  **marufxplorer** 1 year, 10 months ago

C  
Amazon DynamoDB with DynamoDB Accelerator (DAX): DynamoDB is a fully managed NoSQL database service provided by AWS. It is designed for low-latency access to frequently accessed data. DynamoDB Accelerator (DAX) is an in-memory cache for DynamoDB that can significantly reduce read latency, making it suitable for achieving sub-millisecond read times.  
upvoted 4 times

  **lucdt4** 1 year, 11 months ago

**Selected Answer: C**

C is correct  
A don't meets a requirement (LEAST operational overhead) because use script  
B: Not regarding to require  
D: Kinesis for near-real-time (Not for read)  
-> C is correct  
upvoted 3 times

  **DagsH** 2 years, 1 month ago

**Selected Answer: C**

Agreed C will be best because of DynamoDB DAX  
upvoted 2 times

  **BeeKayEnn** 2 years, 1 month ago

Option C will be the best fit.  
As they would like to retrieve the data with sub-millisecond, DynamoDB with DAX is the answer.  
DynamoDB supports some of the world's largest scale applications by providing consistent, single-digit millisecond response times at any scale. You can build applications with virtually unlimited throughput and storage.  
upvoted 3 times

  **Grace83** 2 years, 1 month ago

C is the correct answer  
upvoted 2 times

  **KAUS2** 2 years, 1 month ago


**Selected Answer: C**

Option C is the right one. The questions clearly states "sub-millisecond latency "  
upvoted 3 times

  **smgsi** 2 years, 1 month ago

**Selected Answer: C**

[https://aws.amazon.com/dynamodb/dax/?nc1=h\\_ls](https://aws.amazon.com/dynamodb/dax/?nc1=h_ls)  
upvoted 4 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: C**

Ccccccccccc  
upvoted 3 times

A company uses a payment processing system that requires messages for a particular payment ID to be received in the same order that they were sent. Otherwise, the payments might be processed incorrectly.

Which actions should a solutions architect take to meet this requirement? (Choose two.)



- A. Write the messages to an Amazon DynamoDB table with the payment ID as the partition key.
- B. Write the messages to an Amazon Kinesis data stream with the payment ID as the partition key.
- C. Write the messages to an Amazon ElastiCache for Memcached cluster with the payment ID as the key.
- D. Write the messages to an Amazon Simple Queue Service (Amazon SQS) queue. Set the message attribute to use the payment ID.
- E. Write the messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the message group to use the payment ID.


**Correct Answer:** *BE*

*Community vote distribution*

BE (79%)

AE (18%)

-   **Ashkan\_10**

Highly Voted 



 2 years, 1 month ago


Selected Answer: BE

Option B is preferred over A because Amazon Kinesis Data Streams inherently maintain the order of records within a shard, which is crucial for the given requirement of preserving the order of messages for a particular payment ID. When you use the payment ID as the partition key, all messages for that payment ID will be sent to the same shard, ensuring that the order of messages is maintained.

On the other hand, Amazon DynamoDB is a NoSQL database service that provides fast and predictable performance with seamless scalability. While it can store data with partition keys, it does not guarantee the order of records within a partition, which is essential for the given use case. Hence, using Kinesis Data Streams is more suitable for this requirement.

As DynamoDB does not keep the order, I think BE is the correct answer here.

upvoted 30 times
-   **awsgeek75**

Highly Voted 



 1 year, 3 months ago

I don't understand the question. The only requirement is: " system that requires messages for a particular payment ID to be received in the same order that they were sent"



SQS FIFO (E) meets this requirement.


Why would you "write the message" to Kinesis or DynamoDB anymore. There is no streaming or DB storage requirement in the question. Between A/B, B is better logically but it doesn't meet any stated requirement.

Happy to understand what I'm missing

upvoted 9 times
-   **MatAlves** 7 months, 3 weeks ago

Instead of "what actions...", the question should say "what are the alternatives/options that meet this requirement".



upvoted 4 times
-   **Manimgh**

Most Recent 

 2 weeks, 5 days ago

Selected Answer: BE



same order = FIFO first in first out

upvoted 1 times
-   **FlyingHawk** 5 months ago

Selected Answer: BE



Both B and E can ensure the messages being processed in the order they were received, however SQS FIFO queues (E) are simpler, more cost-effective, and better aligned with the needs of this use case compared to Kinesis. Unless you have additional real-time analytics or high-throughput requirements, E is the superior choice, the big reason why KDS can preserve the order is due to sequence number assigned to data record when it was written to shard, Dynamo DB does not have a sequence number associated with the record.

<https://docs.aws.amazon.com/streams/latest/dev/key-concepts.html>

upvoted 2 times
-   **pentium75** 1 year, 3 months ago

Selected Answer: BE

Both Kinesis and SQS FIFO queue guarantee the order, other answers don't.

upvoted 4 times
-   **meowruki** 1 year, 5 months ago

Option B (Write the messages to an Amazon Kinesis data stream with the payment ID as the partition key): Kinesis can provide ordered processing within a shard  
Write the messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the message group to use the payment ID.

SQS FIFO (First-In-First-Out) queues preserve the order of messages within a message group.  
upvoted 4 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: BE**

Technically both B and E will ensure processing order, but SQS FIFO was specifically built to handle this requirement.  
There is no ask on how to store the data so A and C are out.  
upvoted 2 times

  **Pritam228** 1 year, 7 months ago

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.Partitions.html>  
upvoted 2 times


  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: DE**

options D and E are better because they mimic a real-world queue system and ensure that payments are processed in the correct order, just like customers in a store would be served in the order they arrived. This is crucial for a payment processing system where order matters to avoid mistakes in payment processing.  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

Amazon Kinesis Data Streams Overkill for Ordering  
Overkill for Ordering: While Kinesis can maintain order within a partition key, it might be seen as overkill for a scenario where your primary concern is maintaining the order of payments. SQS FIFO queues (option E) are specifically designed for this purpose and provide an easier and more cost-effective solution.  
upvoted 1 times

  **omoakin** 1 year, 11 months ago

AAAAAAAAA EEEEEEEEEEEEEEE  
upvoted 3 times

  **Konb** 1 year, 11 months ago

**Selected Answer: AE**

IF the question would be "Choose all the solutions that fulfill these requirements" I would chosen BE.

But it is:  
"Which actions should a solutions architect take to meet this requirement? "

For this reason I chose AE, because we don't need both Kinesis AND SQS for this solution. Both choices complement to order processing: order stored in DB, work item goes to the queue.  
upvoted 3 times

  **Smart** 1 year, 9 months ago

Incorrect, AWS will clarify it by using the phrase - "combination of actions".  
upvoted 2 times

  **luisgu** 1 year, 11 months ago

**Selected Answer: BE**

E --> no doubt  
B --> see <https://docs.aws.amazon.com/streams/latest/dev/key-concepts.html>  
upvoted 2 times

  **kruasan** 2 years ago

**Selected Answer: BE**

1) SQS FIFO queues guarantee that messages are received in the exact order they are sent. Using the payment ID as the message group ensures all messages for a payment ID are received sequentially.  
2) Kinesis data streams can also enforce ordering on a per partition key basis. Using the payment ID as the partition key will ensure strict ordering of messages for each payment ID.  
upvoted 3 times

  **kruasan** 2 years ago

The other options do not guarantee message ordering. DynamoDB and ElastiCache are not message queues. SQS standard queues deliver messages in approximate order only.  
upvoted 3 times

  **mrgeee** 2 years ago



**Selected Answer: BE**

BE no doubt.  
upvoted 2 times

  **nosense** 2 years ago

**Selected Answer: BE**

Option A, writing the messages to an Amazon DynamoDB table, would not necessarily preserve the order of messages for a particular payment ID  
upvoted 2 times

  **MssP** 2 years, 1 month ago

**Selected Answer: BE**

I don't understand A, How you can guarantee the order with DynamoDB?? The order is guaranteed with SQS FIFO and Kinesis Data Stream in 1 shard...  
upvoted 5 times

  **pentium75** 1 year, 4 months ago

If it really means "combination of actions" then A+E would work, because you'd use the FIFO queue (E) to guarantee the order. Then the order in the database doesn't matter. If they want to alternative solutions then obviously B and E would work while A alone doesn't.  
upvoted 2 times

  **Grace83** 2 years, 1 month ago

AE is the answer  
upvoted 2 times

A company is building a game system that needs to send unique events to separate leaderboard, matchmaking, and authentication services concurrently. The company needs an AWS event-driven system that guarantees the order of the events.

Which solution will meet these requirements?

- A. Amazon EventBridge event bus
- B. Amazon Simple Notification Service (Amazon SNS) FIFO topics
- C. Amazon Simple Notification Service (Amazon SNS) standard topics
- D. Amazon Simple Queue Service (Amazon SQS) FIFO queues

**Correct Answer: B**

*Community vote distribution*

B (79%)

D (16%)

5%

- bella**

Highly Voted

2 years ago

Selected Answer: B

I don't honestly / can't understand why people go to ChapGPT to ask for the answers.... if I recall correctly they only consolidated their DB until 2021...

upvoted 18 times
- LeonSauveterre**

5 months ago

Yes, it used to be just like that. Fortunately things have changed as of today - Nov 29 2024 in my timezone. Personally I find ChatGPT + Amazon Q to be quite awesome for SAA tests

upvoted 2 times
- aaroncelestin**

1 year, 8 months ago

Yup, ChatGPT doesn't //know// anything about AWS services. It only repeats what other people have said about it, which could be nonsense or hyperbole or some combination thereof.

upvoted 5 times
- LazyTs**

Highly Voted

1 year, 8 months ago

Selected Answer: B

The answer is B la. SNS FIFO topics queue should be used combined with SQS FIFO queue in this case. The question asked for correct order to different event, so asking for SNS fan out here to send to individual SQS.

<https://docs.aws.amazon.com/sns/latest/dg/fifo-example-use-case.html>

upvoted 14 times
- dkw2342**

1 year, 2 months ago

B is correct, but this is not about SNS -> SQS fan-out, it's not necessary. Just SNS FIFO for ordered pub/sub messaging.

upvoted 3 times
- Po\_chih**

1 year, 6 months ago

The best answer!

upvoted 1 times
- FlyingHawk**

Most Recent

5 months ago

Selected Answer: B

Base on this doc <https://docs.aws.amazon.com/sns/latest/dg/fifo-topic-message-ordering.html>, the answer should B and D, however, if you can only select one, then B.

upvoted 1 times
- MatAlves**

7 months, 3 weeks ago

Selected Answer: B

SNS can have many-to-many relations, while SQS supports only one consumer at a time (many-to-one).

upvoted 5 times
- 1e22522**

9 months ago

First time in my life that the answer is actually SNS

upvoted 4 times
- richiexamaws**

11 months ago

Selected Answer: D

AWS does not currently offer FIFO topics for SNS. SNS only supports standard topics, which do not guarantee message order.

upvoted 1 times

  **elmyth** 8 months ago

I'm creating a topic right now and I have both types to choose.

upvoted 3 times

  **Darshan07** 1 year, 2 months ago

**Selected Answer: B**

Even chat gpt said B

upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: B**

Yes, you can technically do this with SQS FIFO partitioned queue by giving separate group ID's to leaderboard, matchmaking etc but this is not as useful as SNS FIFO and is overkill as no need for storage etc. B is more elegant and concise solution,

upvoted 5 times

  **foha2012** 1 year, 4 months ago

Guys, ChatGPT sucks !. Try removing [most voted] from choice B and it will choose D. And if you put [most voted] in front of A, it will select A. LOL !


upvoted 3 times

  **Marco\_St** 1 year, 4 months ago

**Selected Answer: B**

just know SNS FIFO also can send events or messages cocurrently to many subscribers while maintaining the order it receives. SNS fanout pattern is set in standard SNS which is commonly used to fan out events to large number of subscribers and usually for duplicated messages.

upvoted 2 times

  **Mikado211** 1 year, 4 months ago

**Selected Answer: B**

SQS looks like a good idea first, but since we have to send the same message to multiple destination, even if SQS could do it, SNS is much more dedicated to this kind of usage.

upvoted 5 times

  **sparun1607** 1 year, 5 months ago

My Answer is B

<https://docs.aws.amazon.com/sns/latest/dg/sns-fifo-topics.html>

You can use Amazon SNS FIFO (first in, first out) topics with Amazon SQS FIFO queues to provide strict message ordering and message deduplication. The FIFO capabilities of each of these services work together to act as a fully managed service to integrate distributed applications that require data consistency in near-real time. Subscribing Amazon SQS standard queues to Amazon SNS FIFO topics provides best-effort ordering and at least once delivery.

upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

bbbbbbbbbbbbbbbbbb



upvoted 2 times

  **jaydesai8** 1 year, 9 months ago

**Selected Answer: D**

SQS FIFO maintains the order of the events - Answer is D

upvoted 2 times

  **jayce5** 1 year, 11 months ago

**Selected Answer: B**

It should be the fan-out pattern, and the pattern starts with Amazon SNS FIFO for the orders.



upvoted 3 times

  **danielklein09** 1 year, 11 months ago

**Selected Answer: D**

Answer is D. You are so lazy because instead of searching in documentation or your notes, you are asking ChatGPT. Do you really think you will take this exam ? Hint: ask ChatGPT

upvoted 5 times

  **lucdt4** 1 year, 11 months ago

**Selected Answer: D**

D is correct (SQS FIFO)

Because B can't send event concurrently though it can send in the order of the events

upvoted 1 times

A hospital is designing a new application that gathers symptoms from patients. The hospital has decided to use Amazon Simple Queue Service (Amazon SQS) and Amazon Simple Notification Service (Amazon SNS) in the architecture.

A solutions architect is reviewing the infrastructure design. Data must be encrypted at rest and in transit. Only authorized personnel of the hospital should be able to access the data.

Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

- A. Turn on server-side encryption on the SQS components. Update the default key policy to restrict key usage to a set of authorized principals.
- B. Turn on server-side encryption on the SNS components by using an AWS Key Management Service (AWS KMS) customer managed key. Apply a key policy to restrict key usage to a set of authorized principals.
- C. Turn on encryption on the SNS components. Update the default key policy to restrict key usage to a set of authorized principals. Set a condition in the topic policy to allow only encrypted connections over TLS.
- D. Turn on server-side encryption on the SQS components by using an AWS Key Management Service (AWS KMS) customer managed key. Apply a key policy to restrict key usage to a set of authorized principals. Set a condition in the queue policy to allow only encrypted connections over TLS.
- E. Turn on server-side encryption on the SQS components by using an AWS Key Management Service (AWS KMS) customer managed key. Apply an IAM policy to restrict key usage to a set of authorized principals. Set a condition in the queue policy to allow only encrypted connections over TLS.

**Correct Answer: BD**

*Community vote distribution*

BD (71%)

CD (18%)

11%

- awsgeek75**

Highly Voted

9 months, 2 weeks ago

My god! Every other question is about SQS! I thought this was AWS Solution Architect test not "How to solve any problem in AWS using SQS" test!

upvoted 16 times
- fkie4**

Highly Voted

1 year, 7 months ago

**Selected Answer: BD**

read this:  
<https://docs.aws.amazon.com/sns/latest/dg/sns-server-side-encryption.html>  
<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-server-side-encryption.html>

upvoted 15 times

**Gooniegoogoo**

1 year, 4 months ago

good call.. that confirms on that page:

Important

All requests to topics with SSE enabled must use HTTPS and Signature Version 4.

For information about compatibility of other services with encrypted topics, see your service documentation.

Amazon SNS only supports symmetric encryption KMS keys. You cannot use any other type of KMS key to encrypt your service resources. For help determining whether a KMS key is a symmetric encryption key, see Identifying asymmetric KMS keys.

upvoted 4 times

**pentium75**

Most Recent

10 months, 1 week ago

**Selected Answer: BD**

A and C involve 'updating the default key policy' which is not something you. Either you create a key policy, OR AWS assigns THE "default key policy".  
E 'applies an IAM policy to restrict key usage to a set of authorized principals' which is not how IAM policies work. You can 'apply an IAM policy to restrict key usage', but it would be restricted to the principals who have the policy attached; you can't specify them in the policy.

Leaves B and D. That B lacks the TLS statement is irrelevant because "all requests to topics with SSE enabled must use HTTPS" anyway.

upvoted 6 times

**dkw2342**

8 months ago

Yes, BD is correct.

"All requests to queues with SSE enabled must use HTTPS and Signature Version 4." -> valid for SNS and SQS alike:  
<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-server-side-encryption.html>

<https://docs.aws.amazon.com/sns/latest/dg/sns-server-side-encryption.html>

"Set a condition in the queue policy to allow only encrypted connections over TLS." refers to the "aws:SecureTransport" condition, but it's actually redundant.

upvoted 3 times

🗳️ 👤 **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: CD**

Its only options C and D that covers encryption on transit, encryption at rest and a restriction policy.

upvoted 3 times

🗳️ 👤 **Lalo** 1 year, 4 months ago

Answer is BD

SNS: AWS KMS, key policy, SQS: AWS KMS, Key policy

upvoted 4 times

🗳️ 👤 **luisgu** 1 year, 5 months ago

**Selected Answer: BD**

"IAM policies you can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached"

reference: [https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/security\\_iam\\_service-with-iam.html](https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/security_iam_service-with-iam.html)

that excludes E

upvoted 2 times

🗳️ 👤 **imvb88** 1 year, 6 months ago

**Selected Answer: CD**

Encryption at transit = use SSL/TLS -> rule out A,B

Encryption at rest = encryption on components -> keep C, D, E

KMS always need a key policy, IAM is optional -> E out

-> C, D left, one for SNS, one for SQS. TLS: checked, encryption on components: checked

upvoted 4 times

🗳️ 👤 **Lalo** 1 year, 4 months ago

Answer is BD

SNS: AWS KMS, key policy, SQS: AWS KMS, Key policy

upvoted 4 times

🗳️ 👤 **imvb88** 1 year, 6 months ago

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-data-encryption.html>

You can protect data in transit using Secure Sockets Layer (SSL) or client-side encryption. You can protect data at rest by requesting Amazon SQS to encrypt your messages before saving them to disk in its data centers and then decrypt them when the messages are received.

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html>

A key policy is a resource policy for an AWS KMS key. Key policies are the primary way to control access to KMS keys. Every KMS key must have exactly one key policy. The statements in the key policy determine who has permission to use the KMS key and how they can use it. You can also use IAM policies and grants to control access to the KMS key, but every KMS key must have a key policy.

upvoted 2 times

🗳️ 👤 **MarkGerwich** 1 year, 7 months ago

CD

B does not include encryption in transit.

upvoted 3 times

🗳️ 👤 **Bofi** 1 year, 7 months ago

That was my objection toward option B. CD cover both encryption at Rest and Server-Side\_Encryption

upvoted 1 times

🗳️ 👤 **MssP** 1 year, 7 months ago

in transit is included in D. With C, not include encryption at rest... Server-side will include it.

upvoted 1 times

🗳️ 👤 **Maximus007** 1 year, 7 months ago

ChatGPT returned AD as a correct answer)

upvoted 1 times

🗳️ 👤 **cegama543** 1 year, 7 months ago

**Selected Answer: BE**

B: To encrypt data at rest, we can use a customer-managed key stored in AWS KMS to encrypt the SNS components.

E: To restrict access to the data and allow only authorized personnel to access the data, we can apply an IAM policy to restrict key usage to a set of authorized principals. We can also set a condition in the queue policy to allow only encrypted connections over TLS to encrypt data in transit.

upvoted 3 times

  **Karlos99** 1 year, 7 months ago

**Selected Answer: BD**

For a customer managed KMS key, you must configure the key policy to add permissions for each queue producer and consumer.  
<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-key-management.html>

upvoted 4 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: BE**

bebebe

upvoted 1 times

  **[Removed]** 1 year, 7 months ago

bdbdbdbd

All KMS keys must have a key policy. IAM policies are optional.

upvoted 7 times

A company runs a web application that is backed by Amazon RDS. A new database administrator caused data loss by accidentally editing information in a database table. To help recover from this type of incident, the company wants the ability to restore the database to its state from 5 minutes before any change within the last 30 days.

Which feature should the solutions architect include in the design to meet this requirement?

- A. Read replicas
- B. Manual snapshots
- C. Automated backups
- D. Multi-AZ deployments

**Correct Answer: C**

*Community vote distribution*

C (100%)

- Uzbekistan** 8 months ago

**Selected Answer: C**

Amazon RDS provides automated backups, which can be configured to take regular snapshots of the database instance. By enabling automated backups and setting the retention period to 30 days, the company can ensure that it retains backups for up to 30 days. Additionally, Amazon RDS allows for point-in-time recovery within the retention period, enabling the restoration of the database to its state from any point within the last 30 days, including 5 minutes before any change. This feature provides the required capability to recover from accidental data loss incidents.

upvoted 4 times
- Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

Automated backups allow you to recover your database to any point in time within your specified retention period, which can be up to 35 days. The recovery process creates a new Amazon RDS instance with a new endpoint, and the process takes time proportional to the size of the database. Automated backups are enabled by default and occur daily during the backup window. This feature provides an easy and convenient way to recover from data loss incidents such as the one described in the scenario.

upvoted 4 times
- elearningtakai** 1 year, 7 months ago

**Selected Answer: C**

Option C, Automated backups, will meet the requirement. Amazon RDS allows you to automatically create backups of your DB instance. Automated backups enable point-in-time recovery (PITR) for your DB instance down to a specific second within the retention period, which can be up to 35 days. By setting the retention period to 30 days, the company can restore the database to its state from up to 5 minutes before any change within the last 30 days.

upvoted 4 times
- joechen2023** 1 year, 4 months ago

I selected C as well, but still don't know how the automatic backup could have a copy 5 minutes before any change. AWS doc states "Automated backups occur daily during the preferred backup window. "

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_WorkingWithAutomatedBackups.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html).

I think the answer maybe A, as read replica will be kept sync and then restore from the read replica. could an expert help?

upvoted 2 times
- awsgeek75** 10 months ago

"the company wants the ability to restore the database to its state from 5 minutes before any change"

The automatic backup takes a backup every 5 minutes. This means it can restore the database to 5 minutes in the past.

upvoted 2 times
- TheFivePips** 8 months, 1 week ago

Automated backups enable point-in-time recovery (PITR) for your DB instance down to a specific second within the retention period, which can be up to 35 days

upvoted 2 times
- gold4otas** 1 year, 7 months ago

**Selected Answer: C**

C: Automated Backups


<https://aws.amazon.com/rds/features/backup/>

upvoted 3 times
- WherecanIstart** 1 year, 7 months ago

**Selected Answer: C**

Automated Backups...  
upvoted 3 times



 **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

cccccccc  
upvoted 1 times

A company’s web application consists of an Amazon API Gateway API in front of an AWS Lambda function and an Amazon DynamoDB database. The Lambda function handles the business logic, and the DynamoDB table hosts the data. The application uses Amazon Cognito user pools to identify the individual users of the application. A solutions architect needs to update the application so that only users who have a subscription can access premium content.

Which solution will meet this requirement with the LEAST operational overhead?



- A. Enable API caching and throttling on the API Gateway API.
- B. Set up AWS WAF on the API Gateway API. Create a rule to filter users who have a subscription.
- C. Apply fine-grained IAM permissions to the premium content in the DynamoDB table.
- D. Implement API usage plans and API keys to limit the access of users who do not have a subscription.


**Correct Answer: D**

*Community vote distribution*

D (86%)

14%



-   **Guru4Cloud**


Highly Voted 

 1 year, 2 months ago

Selected Answer: D

Implementing API usage plans and API keys is a straightforward way to restrict access to specific users or groups based on subscriptions. It allows you to control access at the API level and doesn't require extensive changes to your existing architecture. This solution provides a clear and manageable way to enforce access restrictions without complicating other parts of the application

upvoted 12 times
-   **Rcosmos**



Most Recent 

 3 months, 1 week ago

Selected Answer: C



Segurança robusta:

A lógica de controle de acesso é aplicada no nível do banco de dados, garantindo que usuários não autorizados não consigam acessar os dados premium, mesmo que tentem burlar o aplicativo.

upvoted 1 times
-   **Uzbekistan** 8 months ago

Selected Answer: C



Chat GPT said:  
Option C, "Apply fine-grained IAM permissions to the premium content in the DynamoDB table," would likely involve the least operational overhead.  
Here's why:  
Granular Control: IAM permissions allow you to control access at a very granular level, including specific actions (e.g., GetItem, PutItem) on individual resources (e.g., DynamoDB tables).  
Integration with Cognito: IAM policies can be configured to allow access based on the identity of the user authenticated through Cognito. You can create IAM roles or policies that grant access to users with specific attributes or conditions, such as having a subscription.  
Minimal Configuration Changes: This solution primarily involves configuring IAM policies for access control in DynamoDB, which can be done with minimal changes to the existing application architecture.

upvoted 1 times
-   **awsgeek75** 10 months ago



Selected Answer: C

C is correct as per the link and doc:  
<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html#apigateway-usage-plans-best-practices>



D: API keys cannot be used to limit access and this can only be done via methods defined in above link

upvoted 2 times
-   **awsgeek75** 10 months ago

Also, option A is for performance and not for security  
option B, WAF cannot control access based on subscription without massive custom coding which will be a big operational overhead

upvoted 2 times
-   **awsgeek75** 9 months, 2 weeks ago

I had to chose D but must have clicked C incorrectly. It is D as my explanation is about D not C! C is the wrong answer.

upvoted 2 times
-   **lipi0035** 11 months, 2 weeks ago



In the same document <https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html> if you scroll down, it says `Don't use API keys for authentication or authorization to control access to your APIs. If you have multiple APIs in a usage plan, a user with a valid API key for one API in that usage plan can access all APIs in that usage plan. Instead, to control access to your API, use an IAM role, a Lambda authorizer, or an Amazon Cognito user pool.`



In the same document at the bottom, it says "If you're using a developer portal to publish your APIs, note that all APIs in a given usage plan are subscribable, even if you haven't made them visible to your customers."



I go with C  
upvoted 2 times



  **awsgeek75** 10 months ago  
<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html#apigateway-usage-plans-best-practices>



Correct link  
upvoted 2 times



  **TariqKipkemei** 1 year ago  
**Selected Answer: D**  
After you create, test, and deploy your APIs, you can use API Gateway usage plans to make them available as product offerings for your customers. You can configure usage plans and API keys to allow customers to access selected APIs, and begin throttling requests to those APIs based on defined limits and quotas. These can be set at the API, or API method level.  
<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html#:~:text=Creating%20and%20using-,usage%20plans,-with%20API%20keys>  
upvoted 2 times



  **marufxplorer** 1 year, 4 months ago  
D  
Option D involves implementing API usage plans and API keys. By associating specific API keys with users who have a valid subscription, you can control access to the premium content.  
upvoted 2 times

  **kruasan** 1 year, 6 months ago  
**Selected Answer: D**  
A. This would not actually limit access based on subscriptions. It helps optimize and control API usage, but does not address the core requirement.  
B. This could work by checking user subscription status in the WAF rule, but would require ongoing management of WAF and increases operational overhead.  
C. This is a good approach, using IAM permissions to control DynamoDB access at a granular level based on subscriptions. However, it would require managing IAM permissions which adds some operational overhead.  
D. This option uses API Gateway mechanisms to limit API access based on subscription status. It would require the least amount of ongoing management and changes, minimizing operational overhead. API keys could be easily revoked/changed as subscription status changes.  
upvoted 4 times

  **imvb88** 1 year, 6 months ago  
CD both possible but D is more suitable since it mentioned in <https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>  
  
A,B not relevant.  
upvoted 3 times

  **elearningtakai** 1 year, 7 months ago  
**Selected Answer: D**  
The solution that will meet the requirement with the least operational overhead is to implement API Gateway usage plans and API keys to limit access to premium content for users who do not have a subscription.  
Option A is incorrect because API caching and throttling are not designed for authentication or authorization purposes, and it does not provide access control.  
Option B is incorrect because although AWS WAF is a useful tool to protect web applications from common web exploits, it is not designed for authorization purposes, and it might require additional configuration, which increases the operational overhead.  
Option C is incorrect because although IAM permissions can restrict access to data stored in a DynamoDB table, it does not provide a mechanism for limiting access to specific content based on the user subscription. Moreover, it might require a significant amount of additional IAM permissions configuration, which increases the operational overhead.  
upvoted 3 times

  **klayytech** 1 year, 7 months ago  
**Selected Answer: D**  
To meet the requirement with the least operational overhead, you can implement API usage plans and API keys to limit the access of users who do not have a subscription. This way, you can control access to your API Gateway APIs by requiring clients to submit valid API keys with requests. You can associate usage plans with API keys to configure throttling and quota limits on individual client accounts.  
upvoted 3 times

  **techhb** 1 year, 7 months ago  
answer is D ,if looking for least overhead  
<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>  
C will achieve it but operational overhead is high.  
upvoted 3 times

  **quentin17** 1 year, 7 months ago

**Selected Answer: D**

Both C&D are valid solution

According to ChatGPT:

"Applying fine-grained IAM permissions to the premium content in the DynamoDB table is a valid approach, but it requires more effort in managing IAM policies and roles for each user, making it more complex and adding operational overhead."

upvoted 2 times

  **Karlos99** 1 year, 7 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

upvoted 4 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

cccccccc

upvoted 1 times

  **pentium75** 10 months, 1 week ago

"Fine-grained permissions" for only two groups of users, hell no.

"IAM permissions" for customers, also no.

upvoted 2 times

A company is using Amazon Route 53 latency-based routing to route requests to its UDP-based application for users around the world. The application is hosted on redundant servers in the company's on-premises data centers in the United States, Asia, and Europe. The company's compliance requirements state that the application must be hosted on premises. The company wants to improve the performance and availability of the application.

What should a solutions architect do to meet these requirements?

- A. Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the on-premises endpoints. Create an accelerator by using AWS Global Accelerator, and register the NLBs as its endpoints. Provide access to the application by using a CNAME that points to the accelerator DNS.
- B. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the on-premises endpoints. Create an accelerator by using AWS Global Accelerator, and register the ALBs as its endpoints. Provide access to the application by using a CNAME that points to the accelerator DNS.
- C. Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the on-premises endpoints. In Route 53, create a latency-based record that points to the three NLBs, and use it as an origin for an Amazon CloudFront distribution. Provide access to the application by using a CNAME that points to the CloudFront DNS.
- D. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the on-premises endpoints. In Route 53, create a latency-based record that points to the three ALBs, and use it as an origin for an Amazon CloudFront distribution. Provide access to the application by using a CNAME that points to the CloudFront DNS.

**Correct Answer: A**

*Community vote distribution*

A (98%)

🗳️ 👤 **Guru4Cloud** Highly Voted 👍 1 year, 8 months ago

**Selected Answer: A**

NLBs allow UDP traffic (ALBs don't support UDP)  
Global Accelerator uses Anycast IP addresses and its global network to intelligently route users to the optimal endpoint  
Using NLBs as Global Accelerator endpoints provides improved availability and DDoS protection.  
upvoted 12 times

🗳️ 👤 **lucdt4** Highly Voted 👍 1 year, 11 months ago

**Selected Answer: A**

C - D: Cloudfront don't support UDP/TCP  
B: Global accelerator don't support ALB  
A is correct  
upvoted 6 times

🗳️ 👤 **JA2018** Most Recent 🕒 5 months, 2 weeks ago

**Selected Answer: C**

CloudFront's origin can be on-premises sources.

Check this out:  
[https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html#concept\\_CustomOrigin](https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html#concept_CustomOrigin)

"A custom origin is an HTTP server, for example, a web server. The HTTP server can be an Amazon EC2 instance or an HTTP server that you host somewhere else. "

upvoted 1 times

🗳️ 👤 **sandordini** 1 year ago

**Selected Answer: A**

Non-HTTP, Massive performance: NLB, UDP: AWS Global Accelerator  
upvoted 3 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: A**

Neither ALB (B+D) nor CloudFront (C+D) do support UDP.  
upvoted 5 times

🗳️ 👤 **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: A**

UDP = NLB and Global Accelerator  
upvoted 4 times

  **live\_reply\_developers** 1 year, 10 months ago

**Selected Answer: A**

NLB + GA support UDP/TCP  
upvoted 4 times



  **Gooniegoogoo** 1 year, 10 months ago

good reference <https://blog.cloudcraft.co/alb-vs-nlb-which-aws-load-balancer-fits-your-needs/>  
upvoted 3 times

  **SkyZeroZx** 2 years ago

**Selected Answer: A**

UDP = NBL  
UDP = GLOBAL ACCELERATOR  
UPD NOT WORKING WITH CLOUDFRONT  
ANS IS A  
upvoted 4 times



  **MssP** 2 years, 1 month ago

**Selected Answer: A**

More discussions at: <https://www.examttopics.com/discussions/amazon/view/51508-exam-aws-certified-solutions-architect-associate-saa-c02/>  
upvoted 1 times

  **Grace83** 2 years, 1 month ago

Why is C not correct - does anyone know?  
upvoted 2 times

  **MssP** 2 years, 1 month ago

It could be valid but I think A is better. Uses the AWS global network to optimize the path from users to applications, improving the performance of TCP and UDP traffic  
upvoted 2 times

  **Shrestwt** 2 years ago

Latency based routing is already using in the application, so AWS global network will optimize the path from users to applications.  
upvoted 2 times

  **FourOfAKind** 2 years, 1 month ago

**Selected Answer: A**

UDP == NLB  
Must be hosted on-premises != CloudFront  
upvoted 4 times

  **imvb88** 2 years ago

actually CloudFront's origin can be on-premises. Source:  
[https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html#concept\\_CustomOrigin](https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html#concept_CustomOrigin)

"A custom origin is an HTTP server, for example, a web server. The HTTP server can be an Amazon EC2 instance or an HTTP server that you host somewhere else."  
upvoted 2 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: A**

aaaaaaaaa  
upvoted 3 times

A solutions architect wants all new users to have specific complexity requirements and mandatory rotation periods for IAM user passwords.

What should the solutions architect do to accomplish this?

- A. Set an overall password policy for the entire AWS account.
- B. Set a password policy for each IAM user in the AWS account.
- C. Use third-party vendor software to set password requirements.
- D. Attach an Amazon CloudWatch rule to the Create\_newuser event to set the password with the appropriate requirements.

**Correct Answer: A**

*Community vote distribution*

A (94%)

3%

- lostmagnet001**

Highly Voted

1 year, 2 months ago

**Selected Answer: A**

i get confused, the question saids "NEW" users... if you apply this password policy it would affect all the users in the AWS account....

upvoted 10 times
- angel\_marquina**

Highly Voted

1 year, 7 months ago

The question is for new users, answer A is not exact for that case.

upvoted 9 times
- Rcosmos**

Most Recent

3 months, 1 week ago

**Selected Answer: U**

Por que as outras opções não são adequadas?

B. Definir uma política de senha para cada usuário do IAM:

O IAM não permite políticas de senha configuradas individualmente para usuários. As políticas de senha são aplicadas a nível da conta.

upvoted 1 times
- LeonSauveterre**

5 months ago

**Selected Answer: A**

A: Overall password policy for the entire AWS account - affects new users immediately, but not old users immediately. Next time the old users change their passwords, the new rule will apply then.

B: Logically wrong. If you already have a user, then that user must've already have a password. So when you individually set a password policy for these users, they are only affected next time they change their passwords, just like I mentioned above about option A. Plus, this is really tiresome because this is user-wise, and there might be too many users, and the work you must do manually is increasing horribly.

upvoted 1 times
- [Removed]**

11 months, 1 week ago

**Selected Answer: B**

Because its mentioned "all new users"

upvoted 1 times
- [Removed]**

11 months, 1 week ago

Ignore the above, seems to be custom policy can help in the case, so A should be right

upvoted 1 times
- TariqKipkemei**

1 year, 6 months ago

**Selected Answer: A**

You can set a custom password policy on your AWS account to specify complexity requirements and mandatory rotation periods for your IAM users' passwords. When you create or change a password policy, most of the password policy settings are enforced the next time your users change their passwords. However, some of the settings are enforced immediately.

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_passwords\\_account-policy.html#:~:text=Setting%20an%20account-,password%20policy,-for%20IAM%20users](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_passwords_account-policy.html#:~:text=Setting%20an%20account-,password%20policy,-for%20IAM%20users)

upvoted 4 times
- klayytech**

2 years, 1 month ago

**Selected Answer: A**

To accomplish this, the solutions architect should set an overall password policy for the entire AWS account. This policy will apply to all IAM users in the account, including new users.

upvoted 4 times

  **WherecanIstart** 2 years, 1 month ago

**Selected Answer: A**

Set overall password policy ...

upvoted 3 times

  **kampatra** 2 years, 1 month ago

**Selected Answer: A**

A is correct

upvoted 2 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: A**

aaaaaaa

upvoted 5 times

A company has migrated an application to Amazon EC2 Linux instances. One of these EC2 instances runs several 1-hour tasks on a schedule. These tasks were written by different teams and have no common programming language. The company is concerned about performance and scalability while these tasks run on a single instance. A solutions architect needs to implement a solution to resolve these concerns.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Batch to run the tasks as jobs. Schedule the jobs by using Amazon EventBridge (Amazon CloudWatch Events).
- B. Convert the EC2 instance to a container. Use AWS App Runner to create the container on demand to run the tasks as jobs.
- C. Copy the tasks into AWS Lambda functions. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events).
- D. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the tasks. Create an Auto Scaling group with the AMI to run multiple copies of the instance.



**Correct Answer: A**


*Community vote distribution*

A (73%)

C (16%)

9%



-   **fkie4**

Highly Voted 



 1 year, 7 months ago

Selected Answer: C



question said "These tasks were written by different teams and have no common programming language", and key word "scalable". Only Lambda can fulfil these. Lambda can be done in different programming languages, and it is scalable

upvoted 10 times
-   **wsdasdasdqwdaw** 1 year ago



AWS Batch - As a fully managed service, AWS Batch helps you to run batch computing workloads of any scale. AWS Batch automatically provisions compute resources and optimizes the workload distribution based on the quantity and scale of the workloads. With AWS Batch, there's no need to install or manage batch computing software, so you can focus your time on analyzing results and solving problems. <https://docs.aws.amazon.com/batch/latest/userguide/what-is-batch.html> ---> I am voting for A, C would have been OK if the time was within 15 minutes.

upvoted 8 times
-   **FourOfAKind** 1 year, 7 months ago



But the question states "several 1-hour tasks on a schedule", and the maximum runtime for Lambda is 15 minutes, so it can't be A.

upvoted 32 times
-   **FourOfAKind** 1 year, 7 months ago



can't be C


upvoted 9 times
-   **JTruong** 10 months, 1 week ago

Lambda can only execute job under 15 mins\* so C can't be the answer

upvoted 6 times
-   **smgsi** 1 year, 7 months ago



It's not because time limit of lambda is 15 minutes


upvoted 13 times
-   **pentium75**

Highly Voted 

 10 months, 1 week ago

"Running on a schedule" = Batch  
Not C due Lambda < 15 min  
Not D, auto-scaling doesn't make sense for things running on a schedule



upvoted 9 times
-   **Manimgh**

Most Recent 

 2 weeks, 5 days ago

Selected Answer: A

no common programming language = Batch

upvoted 1 times
-   **foha2012** 9 months, 2 weeks ago

Selected Answer: D

Answer = D  
"performance and scalability while these tasks run on a single instance" They gave me a legacy application and want it to autoscale for performance. They dont want it to run on a single EC2 instance. Shouldn't I make an AMI and provision multiple EC2 instances in an autoscaling group ? I could put an ALB in front of it. I wont have to deal with "uncommon programming languages" inside the application... Just a thought..


upvoted 2 times

  **awsgeek75** 10 months ago

**Selected Answer: A**

AWS Batch is for jobs running at schedule on EC2. so option A  
B is operational overhead  
C Lambda is 15 mins max execution  
D Scaling is not a requirement



upvoted 7 times

  **meowruki** 11 months, 1 week ago

**Selected Answer: A**

AWS Batch: AWS Batch is a fully managed service for running batch computing workloads. It dynamically provisions the optimal quantity and type of compute resources based on the volume and specific resource requirements of the batch jobs. It allows you to run tasks written in different programming languages with minimal operational overhead.

upvoted 4 times

  **hungta** 11 months, 2 weeks ago

**Selected Answer: A**

The task working for hour but lambda function timeout is 15 minutes. So vote A.  
upvoted 2 times

  **yodelin** 1 year ago

I know guys are stressed out trying to figure this exam out okay, but no matter what people say, with or without reasoning, at least put your mouth clean. Going like AAA is an issue, but talking shi\* on him just because he didn't write down the reasoning is your fault.

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

It can run heterogeneous workloads and tasks without needing to convert them to a common format.  
AWS Batch manages the underlying compute resources - no need to manage containers, Lambda functions or Auto Scaling groups.  
upvoted 6 times

  **zjcorpuz** 1 year, 3 months ago

AWS Lambda function can only be run for 15 mins  
upvoted 2 times

  **jaydesai8** 1 year, 3 months ago

**Selected Answer: A**

maximum runtime for Lambda is 15 minutes, hence A  
upvoted 3 times



  **antropaws** 1 year, 5 months ago

**Selected Answer: A**

I also go with A.  
upvoted 2 times

  **omoakin** 1 year, 5 months ago



C. Copy the tasks into AWS Lambda functions. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events)  
upvoted 1 times

  **ruqui** 1 year, 5 months ago  
wrong, Lambda maximum runtime is 15 minutes and the tasks run for an hour  
upvoted 3 times

  **KMohsoe** 1 year, 5 months ago

**Selected Answer: A**

B and D out!  
A and C let's think!  
AWS Lambda functions are time limited.  
So, Option A  
upvoted 2 times



  **lucdt4** 1 year, 5 months ago

AAAAAAAAAAAAAAAAAAAA  
because lambda only run within 15 minutes  
upvoted 3 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

Answer is A.  
Could have been C but AWS Lambda functions can be only configured to run up to 15 minutes per execution. While the task in question need an 1hour to run,  
upvoted 5 times

  **luisgu** 1 year, 5 months ago

**Selected Answer: D**

question is asking for the LEAST operational overhead. With batch, you have to create the compute environment, create the job queue, create the job definition and create the jobs --> more operational overhead than creating an ASG

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Things 'running on a schedule' = Batch, not autoscaling

upvoted 2 times

A company runs a public three-tier web application in a VPC. The application runs on Amazon EC2 instances across multiple Availability Zones. The EC2 instances that run in private subnets need to communicate with a license server over the internet. The company needs a managed solution that minimizes operational maintenance.

Which solution meets these requirements?

- A. Provision a NAT instance in a public subnet. Modify each private subnet's route table with a default route that points to the NAT instance.
- B. Provision a NAT instance in a private subnet. Modify each private subnet's route table with a default route that points to the NAT instance.
- C. Provision a NAT gateway in a public subnet. Modify each private subnet's route table with a default route that points to the NAT gateway.
- D. Provision a NAT gateway in a private subnet. Modify each private subnet's route table with a default route that points to the NAT gateway.

**Correct Answer:** C

*Community vote distribution*

C (100%)

**UnluckyDucky**

Highly Voted

 1 year, 7 months ago

Selected Answer: C

"The company needs a managed solution that minimizes operational maintenance"

Watch out for NAT instances vs NAT Gateways.

As the company needs a managed solution that minimizes operational maintenance - NAT Gateway in a public subnet is the answer.

upvoted 9 times

**lucdt4**

Highly Voted

 1 year, 5 months ago

C

Nat gateway can't deploy in a private subnet.

upvoted 5 times

**von\_himmlen**

Most Recent

 5 months, 3 weeks ago

C

<https://docs.aws.amazon.com/appstream2/latest/developerguide/managing-network-internet-NAT-gateway.html> ...and a NAT gateway in a public subnet.

upvoted 2 times

**Guru4Cloud** 1 year, 2 months ago

Selected Answer: C

This meets the requirements for a managed, low maintenance solution for private subnets to access the internet:

NAT gateway provides automatic scaling, high availability, and fully managed service without admin overhead. Placing the NAT gateway in a public subnet with proper routes allows private instances to use it for internet access. Minimal operational maintenance compared to NAT instances.

upvoted 3 times

**Guru4Cloud** 1 year, 2 months ago

No good:

NAT instances (A, B) require more hands-on management.

Placing a NAT gateway in a private subnet (D) would not allow internet access.

upvoted 4 times

**TariqKipkemei** 1 year, 5 months ago

Selected Answer: C

minimizes operational maintenance = NGW

upvoted 2 times

**WherecanIstart** 1 year, 7 months ago

Selected Answer: C

C..provision NGW in Public Subnet

upvoted 3 times

**cegama543** 1 year, 7 months ago

Selected Answer: C

cccc is the best  
upvoted 2 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

cccccccc  
upvoted 3 times

A company needs to create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to host a digital media streaming application. The EKS cluster will use a managed node group that is backed by Amazon Elastic Block Store (Amazon EBS) volumes for storage. The company must encrypt all data at rest by using a customer managed key that is stored in AWS Key Management Service (AWS KMS).

Which combination of actions will meet this requirement with the LEAST operational overhead? (Choose two.)

- A. Use a Kubernetes plugin that uses the customer managed key to perform data encryption.
- B. After creation of the EKS cluster, locate the EBS volumes. Enable encryption by using the customer managed key.
- C. Enable EBS encryption by default in the AWS Region where the EKS cluster will be created. Select the customer managed key as the default key.
- D. Create the EKS cluster. Create an IAM role that has a policy that grants permission to the customer managed key. Associate the role with the EKS cluster.
- E. Store the customer managed key as a Kubernetes secret in the EKS cluster. Use the customer managed key to encrypt the EBS volumes.

**Correct Answer:** CD

*Community vote distribution*

CD (53%)BD (44%)

- asoli**

Highly Voted

2 years, 1 month ago

Selected Answer: CD

<https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html#:~:text=encrypted%20Amazon%20EBS%20volumes%20without%20using%20a%20launch%20template%2C%20encrypt%20all%20new%20Amazon%20EBS%20volumes%20created%20in%20your%20account.>

upvoted 16 times
- bujuman**

1 year, 1 month ago

If you want to encrypt Amazon EBS volumes for your nodes, you can deploy the nodes using a launch template. To deploy managed nodes with encrypted Amazon EBS volumes without using a launch template, encrypt all new Amazon EBS volumes created in your account. For more information, see Encryption by default in the Amazon EC2 User Guide for Linux Instances.

upvoted 3 times
- imvb88**

Highly Voted

2 years ago

Selected Answer: BD

Quickly rule out A (which plugin? > overhead) and E because of bad practice

Among B,C,D: B and C are functionally similar > choice must be between B or C, D is fixed

Between B and C: C is out since it set default for all EBS volume in the region, which is more than required and even wrong, say what if other EBS volumes of other applications in the region have different requirement?

upvoted 12 times
- NSA\_Poker**

11 months, 1 week ago

(C) is correct; the EBS volumes of other applications in the region will not be affected bc an IAM role will limit the encryption key to the EKS cluster.

upvoted 2 times
- Yak\_Yeti**

Most Recent

4 days, 7 hours ago

Selected Answer: CD

D is clear and as to B vs C? The trade-off is the potential for encrypting other volumes in the Region, but the manual effort and complexity of Option B outweigh this concern when considering operational efficiency.

upvoted 1 times
- babayomi**

6 months, 1 week ago

CD

This options are wrongly stated in my understanding. B is wrong because you can not encrypt an unencrypted created ebs volume. You need to take a snapshot of the volume, encrypt the snapshot ,creat new ebs volume from snapshot. You can also enable encryption during creation. So the statement is wrong. The next available option can only be C, except that in that account all ebs volumes created would be encrypted, this is also questionable. Because if another person create a new ebs volumes it's automatically encrypted.

upvoted 2 times
- scaredSquirrel**

8 months, 2 weeks ago

Selected Answer: CD

A and E are obvious nos. D is a shoo-in.

The difference between B&C is basically EBS encryption by default vs encryption. Encryption by default is by region, and encrypt everything in that region going forward, versus simple encryption is volume by volume, C is less operational overhead. Check doc & chatGPT.

upvoted 2 times

🗳️ 👤 **jjcode** 1 year, 2 months ago

this one is going on my skip list

upvoted 7 times

🗳️ 👤 **Mahmoudddddddd** 1 year, 1 month ago

Don't it came for me in my exam today xd

upvoted 4 times

🗳️ 👤 **JA2018** 5 months, 2 weeks ago

Hi Mahmoudddddddd, can you share what were your chosen answers?

upvoted 1 times

🗳️ 👤 **jaswantn** 1 year, 2 months ago

If question is giving a requirement related to a particular case and asking to encrypt all data at rest; it is clear that encryption is for this case only and not for other projects in entire region. so option B is more appropriate along with option D.

upvoted 2 times

🗳️ 👤 **frmrkc** 1 year, 3 months ago

**Selected Answer: CD**

It says: 'The company must encrypt ALL data at rest', so there is nothing wrong with 'enabling EBS encryption by default' . C & D

upvoted 5 times

🗳️ 👤 **LeonSauveterre** 5 months ago

Exactly. Option B is out of the question. Not to mention option C barely has any operational overhead.

upvoted 1 times

🗳️ 👤 **upliftinghut** 1 year, 3 months ago

**Selected Answer: BD**

B&D are correct. C is wrong because when you turn on encryption by default, AWS uses its own key while the requirement is using Customer key.

Detail is here: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html#encryption-by-default>

upvoted 3 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: BD**

Not A (avoid 3rd party plugins when there are native services)

Not C ("encryption by default" would impact other services)

Not E (Keys belong in KMS, not in EKS cluster)

upvoted 3 times

🗳️ 👤 **awsgeek75** 1 year, 4 months ago

"The company must encrypt all data at rest by using a customer managed key that is stored in AWS Key Management Service (AWS KMS)."

I am just a bit concerned that the question does not put any limits on not encrypting all the EBS by default in the account. Both B and C can work. C is a hack but it is definitely LEAST operational overhead. Also, we don't know if there are other services or not that may be impacted. What do you think?

upvoted 2 times

🗳️ 👤 **Marco\_St** 1 year, 4 months ago

**Selected Answer: CD**

EBS encryption is set regionally. AWS account is global but it does not mean EBS encryption is enable by default at account level. default EBS encryption is a regional setting within your AWS account. Enabling it in a specific region ensures that all new EBS volumes created in that region are encrypted by default, using either the default AWS managed key or a customer managed key that you specify.

upvoted 2 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

"Enabling it in a specific region ensures that all new EBS volumes created in that region are encrypted by default" which is not what we want. We want to encrypt the EBS volumes used by this EKS cluster, NOT "all new EBS volumes created in that region."

upvoted 2 times

🗳️ 👤 **maudsha** 1 year, 6 months ago

**Selected Answer: CD**

IF you need to encrypt an unencrypted volume,

- Create an EBS snapshot of the volume
- Encrypt the EBS snapshot ( using copy )
- Create new EBS volume from the snapshot ( the volume will also be encrypted )

so it has an operational overhead.

So assuming they won't use this account for anything else we can use C. Enable EBS encryption by default in the AWS Region where the EKS cluster will be created. Select the customer managed key as the default key.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

"Assuming they won't use this account for anything else" how could we assume that?

upvoted 2 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: CD**

Option D is required wither way.

Technically both option B and C would work, but with B you would have to enable encryption node by node, while with option C provides a onetime action of enabling encryption on all nodes.

The requirement is the option with LEAST operational overhead.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

B created some deployment work, but NOT "operational (!) overhead" once it's deployed. C enables encryption by default for all new EBS volumes which is not what we want.

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: CD**

These options allow EBS encryption with the customer managed KMS key with minimal operational overhead:

C) Setting the KMS key as the regional EBS encryption default automatically encrypts new EKS node EBS volumes.

D) The IAM role grants the EKS nodes access to use the key for encryption/decryption operations.

upvoted 1 times

  **jaydesai8** 1 year, 9 months ago

**Selected Answer: CD**

C - enable EBS encryption by default in a region -<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>


D - Provides key access permission just to the EKS cluster without changing broader IAM permissions

upvoted 1 times

  **pentium75** 1 year, 4 months ago

We're not asked to enable EBS encryption by default.

upvoted 2 times

  **pedroso** 1 year, 10 months ago

**Selected Answer: BD**

I was in doubt between B and C.

You can't "Enable EBS encryption by default in the AWS Region". Enable EBS encryption by default is only possible at Account level, not Region.

B is the right option once you can enable encryption on the EBS volume with KMS and custom KMS.

upvoted 2 times

  **antropaws** 1 year, 10 months ago

Not accurate: "Encryption by default is a Region-specific setting":



<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html#encryption-by-default>

upvoted 4 times

  **pentium75** 1 year, 4 months ago

Still C is wrong because "encryption by default" is not what we want.

upvoted 2 times

  **jayce5** 1 year, 11 months ago

**Selected Answer: CD**

It's C and D. I tried it in my AWS console.

C seems to have fewer operations ahead compared to B.

upvoted 5 times

A company wants to migrate an Oracle database to AWS. The database consists of a single table that contains millions of geographic information systems (GIS) images that are high resolution and are identified by a geographic code.

When a natural disaster occurs, tens of thousands of images get updated every few minutes. Each geographic code has a single image or row that is associated with it. The company wants a solution that is highly available and scalable during such events.

Which solution meets these requirements MOST cost-effectively?

- A. Store the images and geographic codes in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance.
- B. Store the images in Amazon S3 buckets. Use Amazon DynamoDB with the geographic code as the key and the image S3 URL as the value.
- C. Store the images and geographic codes in an Amazon DynamoDB table. Configure DynamoDB Accelerator (DAX) during times of high load.
- D. Store the images in Amazon S3 buckets. Store geographic codes and image S3 URLs in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance.

Correct Answer: B

Community vote distribution

B (72%)

D (28%)

- Wayne23Fang

Highly Voted

1 year, 7 months ago

Selected Answer: B

Amazon prefers people to move from Oracle to its own services like DynamoDB and S3.  
upvoted 16 times
- Karlos99

Highly Voted

2 years, 1 month ago

Selected Answer: D

The company wants a solution that is highly available and scalable  
upvoted 9 times
- EllenLiu

4 months, 2 weeks ago

another thing is the design of dynamoDB, "Use Amazon DynamoDB with the geographic code as the key and the image S3 URL as the value."  
geographic code can not be the only key for any image  
upvoted 1 times
- [Removed]

2 years, 1 month ago

But DynamoDB is also highly available and scalable  
<https://aws.amazon.com/dynamodb/faqs/#:~:text=DynamoDB%20automatically%20scales%20throughput%20capacity,high%20availability%20and%20data%20durability.>  
upvoted 5 times
- pbpally

1 year, 12 months ago

Yes but has a size limit at 400kb so theoretically it could store images but it's not a plausible solution.  
upvoted 2 times
- ruqui

1 year, 11 months ago

It doesn't matter the size limit of DynamoDB!!!! The images are saved in S3 buckets. Right answer is B  
upvoted 8 times
- jaydesai8

1 year, 9 months ago

but would it be easy and cost-effective to migrate Oracle (relational db) to (Dynamodb)NoSQL?  
upvoted 5 times
- pentium75

1 year, 4 months ago

Yes because it's a single table with two records, for which Oracle or any relation database has been a bad choice in the first place.  
upvoted 5 times
- upliftinghut

Most Recent

1 year, 3 months ago

Selected Answer: B

DynamoDB with its HA and built-in scalability. The nature of the table also resonates with NoSQL than SQL DB such as Oracle. Only 1 table so migration is just a script from Oracle to DynamoDB  
  
D is workable but more expensive with Oracle licenses and other setups for HA and scalability  
upvoted 3 times

- upliftinghut

1 year, 3 months ago

HA & built-in scalability of Amazon DynamoDB :  
[https://aws.amazon.com/dynamodb/features/#:~:text=Amazon%20DynamoDB%20is%20a%20fully,for%20the%20most%20demanding%20appli](https://aws.amazon.com/dynamodb/features/#:~:text=Amazon%20DynamoDB%20is%20a%20fully,for%20the%20most%20demanding%20applications.)  
cations.  
upvoted 3 times
- awsgeek75

1 year, 4 months ago

Selected Answer: B

A puts images in Oracle, not a good idea  
C DAX is not going to help with images  
D It is doable but RDS on multi AZ does not give you more performance or write scalability. It gives more availability and read scalability which is not required here.  
B works as Geographic code is the key in DynamoDB and S3 image URL is the data so DynamoDB can handle tens of thousands such record and S3 can scale for writing  
upvoted 4 times
- pentium75

1 year, 4 months ago

Selected Answer: B

They are currently using Oracle, but only for one simple table with a single key-value pair. This is a typical use case for a NoSQL database like DynamoDB (and whoever decided to use Oracle for this in the first place should be fired). Oracle is expensive as hell, so options A and D might work but are surely not cost-effective. C won't work because the images are too big for the database. Leaves B which would be the ideal solution and meet the availability and scalability requirements.  
upvoted 6 times
- wsdasdasdqwdaw

1 year, 6 months ago

For D - Oracle is not cheap as well. RDS with Oracle vs DynamoDB, I would go for pure AWS provided option. In each exam there is a lot of marketing => B  
upvoted 3 times
- jubolano

1 year, 6 months ago

Selected Answer: D

Cost effective, D  
upvoted 2 times
- awsgeek75

1 year, 3 months ago

How is Oracle more cost effective than other options?  
upvoted 2 times
- wsdasdasdqwdaw

1 year, 6 months ago

B or D, but the question is MOST cost-effectively DynamoDB is more expensive than RDS, I am going for D  
upvoted 3 times
- JA2018

5 months, 2 weeks ago

Key in sTEM states that the customer's current Oracle DB table is setup with a single key-value pair.  
  
A typical use case for a NoSQL database like DynamoDB.  
upvoted 1 times
- gouranga45

1 year, 7 months ago

Selected Answer: B

Answer is B, DynamoDB is Highly available and scalable  
upvoted 2 times
- baba365

1 year, 7 months ago

A single table in a relational db can have items that are related ? e.g. 'select \* from Faculty where department\_id in (10, 20) and dept\_name = AWS'. In the sql query example above, \* means all and Faculty is name of the table.  
upvoted 2 times
- Eminenza22

1 year, 8 months ago

Selected Answer: B

B option offers a cost-effective solution for storing and accessing high-resolution GIS images during natural disasters. Storing the images in Amazon S3 buckets provides scalable and durable storage, while using Amazon DynamoDB allows for quick and efficient retrieval of images based on geographic codes. This solution leverages the strengths of both S3 and DynamoDB to meet the requirements of high availability, scalability, and cost-effectiveness.  
upvoted 2 times
- cd93



1 year, 8 months ago



Selected Answer: B



What were the company thinking using the most expensive DB on the planet FOR ONE SINGLE TABLE???  
Migrate a single table from SQL to NoSQL should be easy enough I guess...  
upvoted 3 times
- vini15



1 year, 9 months ago



Should be D.  
the question says company wants to migrate oracle to AWS. Oracle is a relational db hence RDS makes more sense whereas Dynamodb is non relational db.  
upvoted 3 times



  **pentium75** 1 year, 4 months ago  
But relational DB does not make sense for the use case. It's a single table.  
upvoted 2 times

  **iBanan** 1 year, 9 months ago  
I hate these questions:) I can't choose between B and D  
upvoted 6 times

  **ces\_9999** 1 year, 9 months ago  
Guys the answer is B the oracle database only has one table without any relationships so why we should use a relational database in the first place, second we are storing the images in S3 not in the database why not use this alongside dynamo  
upvoted 6 times



  **Kp88** 1 year, 9 months ago  
You can't do migration of Oracle to Dynmodb without SCT. I am not the DB guy but since its saying oracle I would go with D otherwise B makes more sense if a company is starting out from scratch.  
upvoted 2 times

  **Kp88** 1 year, 9 months ago  
Actually now that I think about it , B sounds ok as well. Company just need to use SCT and that would be more cost effective.  
upvoted 2 times

  **joehong** 1 year, 10 months ago  

Selected Answer: D

  
"A company wants to migrate an Oracle database to AWS"  
upvoted 2 times

  **pentium75** 1 year, 4 months ago  
Yeah, per my understanding that doesn't implicate that the destination must be an Oracle database.  
upvoted 2 times

  **secdgs** 1 year, 10 months ago  
D: Wrorng  
if you caluate License Oracle Database, It is not cost-effectively. Multi-AZ is not scalable and if you set scalable, you need more license for Oracle database.  
upvoted 2 times

A company has an application that collects data from IoT sensors on automobiles. The data is streamed and stored in Amazon S3 through Amazon Kinesis Data Firehose. The data produces trillions of S3 objects each year. Each morning, the company uses the data from the previous 30 days to retrain a suite of machine learning (ML) models.

Four times each year, the company uses the data from the previous 12 months to perform analysis and train other ML models. The data must be available with minimal delay for up to 1 year. After 1 year, the data must be retained for archival purposes.

Which storage solution meets these requirements MOST cost-effectively?

- A. Use the S3 Intelligent-Tiering storage class. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year.
- B. Use the S3 Intelligent-Tiering storage class. Configure S3 Intelligent-Tiering to automatically move objects to S3 Glacier Deep Archive after 1 year.
- C. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Create an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 1 year.
- D. Use the S3 Standard storage class. Create an S3 Lifecycle policy to transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days, and then to S3 Glacier Deep Archive after 1 year.

**Correct Answer: D**

*Community vote distribution*

D (91%)

7%

- UnluckyDucky**

Highly Voted

1 year, 7 months ago

Selected Answer: D

Access patterns is given, therefore D is the most logical answer.

Intelligent tiering is for random, unpredictable access.

upvoted 14 times
- ealpuche**

1 year, 5 months ago

You are missing: <<The data must be available with minimal delay for up to 1 year. After one year, the data must be retained for archival purposes.>> You are secure that data after 1 year is not accessible anymore.

upvoted 2 times
- TariqKipkemei**

Highly Voted

1 year, 5 months ago

Selected Answer: D

First 30 days data accessed every morning = S3 Standard  
Beyond 30 days data accessed quarterly = S3 Standard-Infrequent Access  
Beyond 1 year data retained = S3 Glacier Deep Archive

upvoted 5 times
- jjcode**

Most Recent

8 months, 3 weeks ago

I dont get how its A

  1. Each morning, the company uses the data from the previous 30 days
  2. Four times each year, the company uses the data from the previous 12 months to perform analysis and train other ML models
  3. The data must be available with minimal delay for up to 1 year. After 1 year, the data must be retained for archival purposes

The data ingestion happens 4 times a year, that means that after the initial 30 days it still needs to be pulled 3 more times, why would you put the data in standard infrequent if you were going to use it 3 more times and speed is a requirement? Makes more sense to put it in S3 standard, or intelligent then straight to glacier.

upvoted 2 times
- upliftinghut**

9 months, 2 weeks ago

Selected Answer: D

Clear access pattern. data in Standard-Infrequent Access is for data requires rapid access when needed

upvoted 2 times
- awsgeek75**

10 months ago

Selected Answer: D

A and B, Intelligent Tiering cannot be configured. It is managed by AWS.  
C SIA does not allow immediate access for "each morning"  
D is best for 30 day standard access, SIA after 30 days and archive after 1 year

upvoted 4 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: D**

See reasoning below, just accidentally voted A

upvoted 2 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: A**

The data is used every day (typical use case for Standard) for 30 days, for the remaining 12 months it is used 3 or 4 times (typical use case for IA), after 12 months it is not used at all but must be kept (typical use case for Glacier Deep Archive).

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Sorry, D!!!!!!!!! Not A!!!! D!

upvoted 4 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: D**

This option optimizes costs while meeting the data access requirements:

Store new data in S3 Standard for first 30 days of frequent access  
Transition to S3 Standard-IA after 30 days for infrequent access up to 1 year  
Archive to Glacier Deep Archive after 1 year for long-term archival

upvoted 3 times

  **ealpuche** 1 year, 5 months ago

**Selected Answer: A**


Option A meets the requirements most cost-effectively. The S3 Intelligent-Tiering storage class provides automatic tiering of objects between the S3 Standard and S3 Standard-Infrequent Access (S3 Standard-IA) tiers based on changing access patterns, which helps optimize costs. The S3 Lifecycle policy can be used to transition objects to S3 Glacier Deep Archive after 1 year for archival purposes. This solution also meets the requirement for minimal delay in accessing data for up to 1 year. Option B is not cost-effective because it does not include the transition of data to S3 Glacier Deep Archive after 1 year. Option C is not the best solution because S3 Standard-IA is not designed for long-term archival purposes and incurs higher storage costs. Option D is also not the most cost-effective solution as it transitions objects to the S3 Standard-IA tier after 30 days, which is unnecessary for the requirement to retrain the suite of ML models each morning using data from the previous 30 days.

upvoted 1 times

  **pentium75** 10 months, 1 week ago

I can't follow. The data is used every day (typical use case for Standard) for 30 days, for the remaining 12 months it is used 3 or 4 times (typical use case for IA), after 12 months it is not used at all but must be kept (typical use case for Glacier Deep Archive).



upvoted 3 times

  **KAUS2** 1 year, 7 months ago

**Selected Answer: D**

Agree with UnluckyDucky , the correct option is D

upvoted 2 times

  **fkie4** 1 year, 7 months ago

**Selected Answer: D**

Should be D. see this:  
<https://www.examttopics.com/discussions/amazon/view/68947-exam-aws-certified-solutions-architect-associate-saa-c02/>



upvoted 3 times

  **Nithin1119** 1 year, 7 months ago

**Selected Answer: B**

Bbbbbbbbbb

upvoted 1 times

  **fkie4** 1 year, 7 months ago

hello!!??

upvoted 2 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: D**

ddddddd

upvoted 4 times

  **[Removed]** 1 year, 7 months ago

D because:

- First 30 days- data access every morning ( predictable and frequently) – S3 standard
- After 30 days, accessed 4 times a year – S3 infrequently access
- Data preserved- S3 Gllacier Deep Archive

upvoted 9 times



A company is running several business applications in three separate VPCs within the us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds of gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.

A solutions architect needs to design a network connectivity solution that maximizes cost-effectiveness.

Which solution meets these requirements?

- A. Configure three AWS Site-to-Site VPN connections from the data center to AWS. Establish connectivity by configuring one VPN connection for each VPC.
- B. Launch a third-party virtual network appliance in each VPC. Establish an IPsec VPN tunnel between the data center and each virtual appliance.
- C. Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway in us-east-1. Establish connectivity by configuring each VPC to use one of the Direct Connect connections.
- D. Set up one AWS Direct Connect connection from the data center to AWS. Create a transit gateway, and attach each VPC to the transit gateway. Establish connectivity between the Direct Connect connection and the transit gateway.

**Correct Answer:** D

*Community vote distribution*

D (100%)

- TariqKipkemei**

Highly Voted

 1 year ago

**Selected Answer: D**

AWS Transit Gateway connects your Amazon Virtual Private Clouds (VPCs) and on-premises networks through a central hub. This connection simplifies your network and puts an end to complex peering relationships. Transit Gateway acts as a highly scalable cloud router—each new connection is made only once.

<https://aws.amazon.com/transit-gateway/#:~:text=AWS-,Transit%20Gateway,-connects%20your%20Amazon>

upvoted 6 times
- Guru4Cloud**

Highly Voted

 1 year, 2 months ago

**Selected Answer: D**

This option leverages a single Direct Connect for consistent, private connectivity between the data center and AWS. The transit gateway allows each VPC to share the Direct Connect while keeping the VPCs isolated. This provides a cost-effective architecture to meet the requirements.

upvoted 5 times
- upliftinghut**

Most Recent

 9 months, 2 weeks ago

**Selected Answer: D**

AWS Direct connect is costly but the saving comes from less data transfer cost with Direct Connect and Transit gateway

upvoted 4 times
- alexandercamachop** 1 year, 5 months ago

**Selected Answer: D**

Transit GW, is a hub for connecting all VPCs.  
Direct Connect is expensive, therefor only 1 of them connected to the Transit GW (Hub for all our VPCs that we connect to it)

upvoted 4 times
- KMohsoe** 1 year, 5 months ago

**Selected Answer: D**

Option D

upvoted 4 times
- Sivasaa** 1 year, 6 months ago

Can someone tell why option C will not work here

upvoted 3 times
- pentium75** 10 months, 1 week ago

And besides the cost, C does not allow the applications "to communicate between VPCs".

upvoted 2 times
- jdamian** 1 year, 6 months ago

cost-effectiveness, 3 DC are more than 1 (more expensive). There is no need to connect more than 1 DC.  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

Using multiple Site-to-Site VPNs (A) or Direct Connects (C) incurs higher costs without providing significant benefits.  
upvoted 2 times

  **SkyZeroZx** 1 year, 6 months ago

**Selected Answer: D**

cost-effectiveness  
D  
upvoted 2 times

  **WherecanIstart** 1 year, 7 months ago

**Selected Answer: D**

Transit Gateway will achieve this result..  
upvoted 4 times

  **Karlos99** 1 year, 7 months ago

**Selected Answer: D**

maximizes cost-effectiveness  
upvoted 3 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: D**

ddddddddd  
upvoted 3 times

An ecommerce company is building a distributed application that involves several serverless functions and AWS services to complete order-processing tasks. These tasks require manual approvals as part of the workflow. A solutions architect needs to design an architecture for the order-processing application. The solution must be able to combine multiple AWS Lambda functions into responsive serverless applications. The solution also must orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Step Functions to build the application.
- B. Integrate all the application components in an AWS Glue job.
- C. Use Amazon Simple Queue Service (Amazon SQS) to build the application.
- D. Use AWS Lambda functions and Amazon EventBridge events to build the application.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **kinglong12** Highly Voted 1 year, 1 month ago

**Selected Answer: A**

AWS Step Functions is a fully managed service that makes it easy to build applications by coordinating the components of distributed applications and microservices using visual workflows. With Step Functions, you can combine multiple AWS Lambda functions into responsive serverless applications and orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers. Step Functions also allows for manual approvals as part of the workflow. This solution meets all the requirements with the least operational overhead.

upvoted 15 times

  **Guru4Cloud** Highly Voted 8 months ago

**Selected Answer: A**

AWS Step Functions allow you to easily coordinate multiple Lambda functions and services into serverless workflows with visual workflows. Step Functions are designed for building distributed applications that combine services and require human approval steps.

Using Step Functions provides a fully managed orchestration service with minimal operational overhead.

upvoted 7 times

  **TariqKipkemei** Most Recent 6 months, 2 weeks ago

**Selected Answer: A**

involves several serverless functions and AWS services, require manual approvals as part of the workflow, combine the Lambda functions into responsive serverless applications, orchestrate data and services = AWS Step Functions

upvoted 4 times

  **capino** 8 months, 3 weeks ago

**Selected Answer: A**

Serverless && workflow service that need human approval::::step functions

upvoted 3 times

  **BeeKayEnn** 1 year, 1 month ago

Key: Distributed Application Processing, Microservices orchestration (Orchestrate Data and Services)

A would be the best fit.

AWS Step Functions is a visual workflow service that helps developers use AWS services to build distributed applications, automate processes, orchestrate microservices, and create data and machine learning (ML) pipelines.

Reference: [https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20\(ML\)%20pipelines](https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20(ML)%20pipelines).

upvoted 4 times

  **COTIT** 1 year, 1 month ago

**Selected Answer: A**

Approval is explicit for the solution. -> "A common use case for AWS Step Functions is a task that requires human intervention (for example, an approval process). Step Functions makes it easy to coordinate the components of distributed applications as a series of steps in a visual workflow called a state machine. You can quickly build and run state machines to execute the steps of your application in a reliable and scalable fashion. (<https://aws.amazon.com/pt/blogs/compute/implementing-serverless-manual-approval-steps-in-aws-step-functions-and-amazon-api-gateway/>)"

upvoted 6 times

  **ktulu2602** 1 year, 1 month ago

**Selected Answer: A**

Option A: Use AWS Step Functions to build the application.

AWS Step Functions is a serverless workflow service that makes it easy to coordinate distributed applications and microservices using visual workflows. It is an ideal solution for designing architectures for distributed applications that involve multiple AWS services and serverless functions, as it allows us to orchestrate the flow of our application components using visual workflows. AWS Step Functions also integrates with other AWS services like AWS Lambda, Amazon EC2, and Amazon ECS, and it has built-in error handling and retry mechanisms. This option provides a serverless solution with the least operational overhead for building the application.

upvoted 5 times

A company has launched an Amazon RDS for MySQL DB instance. Most of the connections to the database come from serverless applications. Application traffic to the database changes significantly at random intervals. At times of high demand, users report that their applications experience database connection rejection errors.



Which solution will resolve this issue with the LEAST operational overhead?

- A. Create a proxy in RDS Proxy. Configure the users’ applications to use the DB instance through RDS Proxy.
- B. Deploy Amazon ElastiCache for Memcached between the users’ applications and the DB instance.
- C. Migrate the DB instance to a different instance class that has higher I/O capacity. Configure the users’ applications to use the new DB instance.
- D. Configure Multi-AZ for the DB instance. Configure the users’ applications to switch between the DB instances.

Correct Answer: A

Community vote distribution

A (100%)

  **TariqKipkemei** Highly Voted 6 months, 2 weeks ago

Selected Answer: A

database connection rejection errors = RDS Proxy  
upvoted 7 times

  **ktulu2602** Highly Voted 1 year, 1 month ago

Selected Answer: A


The correct solution for this scenario would be to create a proxy in RDS Proxy. RDS Proxy allows for managing thousands of concurrent database connections, which can help reduce connection errors. RDS Proxy also provides features such as connection pooling, read/write splitting, and retries. This solution requires the least operational overhead as it does not involve migrating to a different instance class or setting up a new cache layer. Therefore, option A is the correct answer.  
upvoted 5 times

  **Guru4Cloud** Most Recent 8 months ago

Selected Answer: A

RDS Proxy provides a proxy layer that pools and shares database connections to improve scalability. This allows the proxy to handle connection spikes to the database gracefully.

Using RDS Proxy requires minimal operational overhead - just create the proxy and reconfigure applications to use it. No code changes needed.  
upvoted 4 times

  **antropaws** 11 months, 1 week ago

Wait, why not B?????  
upvoted 2 times

  **live\_reply\_developers** 10 months ago



Amazon ElastiCache tends to have a lower operational overhead compared to Amazon RDS Proxy. BUT we already have " Amazon RDS for MySQL DB instance"  
upvoted 2 times

  **Guru4Cloud** 8 months ago



ElastiCache (B) and larger instance type (C) help performance but don't resolve connection issues.  
upvoted 2 times

  **Guru4Cloud** 8 months ago

ElastiCache (B) and larger instance type (C) help performance but don't resolve connection issues.  
upvoted 4 times

  **roxx529** 11 months, 2 weeks ago

To reduce application failures resulting from database connection timeouts, the best solution is to enable RDS Proxy on the RDS DB instances  
upvoted 2 times

  **COTIT** 1 year, 1 month ago

Selected Answer: A

Many applications, including those built on modern serverless architectures, can have a large number of open connections to the database server and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows

applications to pool and share connections established with the database, improving database efficiency and application scalability.  
(<https://aws.amazon.com/pt/rds/proxy/>)  
upvoted 4 times

A company recently deployed a new auditing system to centralize information about operating system versions, patching, and installed software for Amazon EC2 instances. A solutions architect must ensure all instances provisioned through EC2 Auto Scaling groups successfully send reports to the auditing system as soon as they are launched and terminated.

Which solution achieves these goals MOST efficiently?

- A. Use a scheduled AWS Lambda function and run a script remotely on all EC2 instances to send data to the audit system.
- B. Use EC2 Auto Scaling lifecycle hooks to run a custom script to send data to the audit system when instances are launched and terminated.
- C. Use an EC2 Auto Scaling launch configuration to run a custom script through user data to send data to the audit system when instances are launched and terminated.
- D. Run a custom script on the instance operating system to send data to the audit system. Configure the script to be invoked by the EC2 Auto Scaling group when the instance starts and is terminated.

**Correct Answer: B**

*Community vote distribution*

B (100%)

🗳️ 👤 **ktulu2602** Highly Voted 👍 2 years, 1 month ago

**Selected Answer: B**

The most efficient solution for this scenario is to use EC2 Auto Scaling lifecycle hooks to run a custom script to send data to the audit system when instances are launched and terminated. The lifecycle hook can be used to delay instance termination until the script has completed, ensuring that all data is sent to the audit system before the instance is terminated. This solution is more efficient than using a scheduled AWS Lambda function, which would require running the function periodically and may not capture all instances launched and terminated within the interval. Running a custom script through user data is also not an optimal solution, as it may not guarantee that all instances send data to the audit system. Therefore, option B is the correct answer.

upvoted 12 times

🗳️ 👤 **COTIT** Highly Voted 👍 2 years, 1 month ago

**Selected Answer: B**

Amazon EC2 Auto Scaling offers the ability to add lifecycle hooks to your Auto Scaling groups. These hooks let you create solutions that are aware of events in the Auto Scaling instance lifecycle, and then perform a custom action on instances when the corresponding lifecycle event occurs. (<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>)

upvoted 6 times

🗳️ 👤 **LeonSauveterre** Most Recent ⌚ 5 months ago

**Selected Answer: B**

C: a **\*\*launch\*\*** configuration CANNOT handle instance termination automatically.

upvoted 1 times

🗳️ 👤 **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: B**

Use EC2 Auto Scaling lifecycle hooks to run a custom script to send data to the audit system when instances are launched and terminated

upvoted 2 times

🗳️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

EC2 Auto Scaling lifecycle hooks allow you to perform custom actions as instances launch and terminate. This is the most efficient way to trigger the auditing script execution at instance launch and termination.

upvoted 5 times

🗳️ 👤 **WherecanIstart** 2 years, 1 month ago

**Selected Answer: B**

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

upvoted 4 times

🗳️ 👤 **fkie4** 2 years, 1 month ago

it is B. read this:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

upvoted 3 times

A company is developing a real-time multiplayer game that uses UDP for communications between the client and servers in an Auto Scaling group. Spikes in demand are anticipated during the day, so the game server platform must adapt accordingly. Developers want to store gamer scores and other non-relational data in a database solution that will scale without intervention.

Which solution should a solutions architect recommend?

- A. Use Amazon Route 53 for traffic distribution and Amazon Aurora Serverless for data storage.
- B. Use a Network Load Balancer for traffic distribution and Amazon DynamoDB on-demand for data storage.
- C. Use a Network Load Balancer for traffic distribution and Amazon Aurora Global Database for data storage.
- D. Use an Application Load Balancer for traffic distribution and Amazon DynamoDB global tables for data storage.

**Correct Answer:** B

*Community vote distribution*

B (100%)

- TariqKipkemei

Highly Voted

11 months, 3 weeks ago

Selected Answer: B

UDP = NLB  
Non-relational data = Dynamo DB  
upvoted 17 times
- fruto123

Highly Voted

1 year, 1 month ago

Selected Answer: B

key words - UDP, non-relational data  
answers - NLB for UDP application, DynamoDB for non-relational data  
upvoted 6 times
- Manimgh

Most Recent

2 weeks, 4 days ago

Selected Answer: B

non-relational data = DynamoDB  
upvoted 1 times
- Guru4Cloud

8 months ago

Selected Answer: B

This option provides the most scalable and optimized architecture for the real-time multiplayer game:  
  
Network Load Balancer efficiently distributes UDP gaming traffic to the Auto Scaling group of game servers. DynamoDB On-Demand mode provides auto-scaling non-relational data storage for gamer scores and other game data. DynamoDB is optimized for fast, high-scale access patterns seen in gaming. Together, the Network Load Balancer and DynamoDB On-Demand provide an architecture that can smoothly scale up and down to match spikes in gaming demand.  
upvoted 5 times
- elearningtakai

1 year, 1 month ago

Selected Answer: B

Option B is a good fit because a Network Load Balancer can handle UDP traffic, and Amazon DynamoDB on-demand can provide automatic scaling without intervention  
upvoted 3 times
- KAUS2

1 year, 1 month ago

Selected Answer: B

Correct option is "B"  
upvoted 2 times
- aragon\_saa

1 year, 1 month ago

B

<https://www.examttopics.com/discussions/amazon/view/29756-exam-aws-certified-solutions-architect-associate-saa-c02/>  
upvoted 2 times
- Kenp1192

1 year, 1 month ago

B

Because NLB can handle UDP and DynamoDB is Non-Relational

upvoted 3 times

A company hosts a frontend application that uses an Amazon API Gateway API backend that is integrated with AWS Lambda. When the API receives requests, the Lambda function loads many libraries. Then the Lambda function connects to an Amazon RDS database, processes the data, and returns the data to the frontend application. The company wants to ensure that response latency is as low as possible for all its users with the fewest number of changes to the company's operations.

Which solution will meet these requirements?

- A. Establish a connection between the frontend application and the database to make queries faster by bypassing the API.
- B. Configure provisioned concurrency for the Lambda function that handles the requests.
- C. Cache the results of the queries in Amazon S3 for faster retrieval of similar datasets.
- D. Increase the size of the database to increase the number of connections Lambda can establish at one time.

**Correct Answer:** *B*

*Community vote distribution*

B (100%)

  **UnluckyDucky** Highly Voted  1 year, 1 month ago

**Selected Answer: B**

Key: the Lambda function loads many libraries

Configuring provisioned concurrency would get rid of the "cold start" of the function therefore speeding up the process.  
upvoted 17 times

  **kampatra** Highly Voted  1 year, 1 month ago

**Selected Answer: B**

Provisioned concurrency – Provisioned concurrency initializes a requested number of execution environments so that they are prepared to respond immediately to your function's invocations. Note that configuring provisioned concurrency incurs charges to your AWS account.  
upvoted 11 times

  **TariqKipkemei** Most Recent  6 months, 2 weeks ago

**Selected Answer: B**

Provisioned concurrency pre-initializes execution environments which are prepared to respond immediately to incoming function requests.  
upvoted 7 times

  **Guru4Cloud** 8 months ago

**Selected Answer: B**

Provisioned concurrency ensures a configured number of execution environments are ready to serve requests to the Lambda function. This avoids cold starts where the function would otherwise need to load all the libraries on each invocation.  
upvoted 4 times

  **Guru4Cloud** 8 months ago

**Selected Answer: B**

Provisioned concurrency ensures a configured number of execution environments are ready to serve requests to the Lambda function. This avoids cold starts where the function would otherwise need to load all the libraries on each invocation.  
upvoted 2 times

  **elearningtakai** 1 year, 1 month ago

**Selected Answer: B**

Answer B is correct  
<https://docs.aws.amazon.com/lambda/latest/dg/provisioned-concurrency.html>  
Answer C: need to modify the application  
upvoted 5 times

  **elearningtakai** 1 year, 1 month ago

This is relevant to "cold start" with keywords: "Lambda function loads many libraries"  
upvoted 2 times

  **Karlos99** 1 year, 1 month ago

**Selected Answer: B**

<https://docs.aws.amazon.com/lambda/latest/dg/provisioned-concurrency.html>  
upvoted 4 times

A company is migrating its on-premises workload to the AWS Cloud. The company already uses several Amazon EC2 instances and Amazon RDS DB instances. The company wants a solution that automatically starts and stops the EC2 instances and DB instances outside of business hours. The solution must minimize cost and infrastructure maintenance.

Which solution will meet these requirements?

- A. Scale the EC2 instances by using elastic resize. Scale the DB instances to zero outside of business hours.
- B. Explore AWS Marketplace for partner solutions that will automatically start and stop the EC2 instances and DB instances on a schedule.
- C. Launch another EC2 instance. Configure a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule.
- D. Create an AWS Lambda function that will start and stop the EC2 instances and DB instances. Configure Amazon EventBridge to invoke the Lambda function on a schedule.

**Correct Answer:** *D*

*Community vote distribution*

D (100%)

  **ktulu2602** Highly Voted  2 years, 1 month ago

**Selected Answer: D**

The most efficient solution for automatically starting and stopping EC2 instances and DB instances on a schedule while minimizing cost and infrastructure maintenance is to create an AWS Lambda function and configure Amazon EventBridge to invoke the function on a schedule.

Option A, scaling EC2 instances by using elastic resize and scaling DB instances to zero outside of business hours, is not feasible as DB instances cannot be scaled to zero.

Option B, exploring AWS Marketplace for partner solutions, may be an option, but it may not be the most efficient solution and could potentially add additional costs.

Option C, launching another EC2 instance and configuring a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule, adds unnecessary infrastructure and maintenance.

upvoted 17 times

  **Guru4Cloud** Highly Voted  1 year, 8 months ago

**Selected Answer: D**

This option leverages AWS Lambda and EventBridge to automatically schedule the starting and stopping of resources.

Lambda provides the script/code to stop/start instances without managing servers.

EventBridge triggers the Lambda on a schedule without cronjobs.

No additional code or third party tools needed.

Serverless, maintenance-free solution

upvoted 6 times

  **1e22522** Most Recent  9 months ago

**Selected Answer: D**

its d but nowadays u use system manager me thinks

upvoted 3 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: D**

Create an AWS Lambda function that will start and stop the EC2 instances and DB instances. Configure Amazon EventBridge to invoke the Lambda function on a schedule.

upvoted 4 times

  **WheracanIstart** 2 years, 1 month ago

**Selected Answer: D**

Minimize cost and maintenance...

upvoted 3 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: D**

DDDDDDDDDDDD

upvoted 1 times

A company hosts a three-tier web application that includes a PostgreSQL database. The database stores the metadata from documents. The company searches the metadata for key terms to retrieve documents that the company reviews in a report each month. The documents are stored in Amazon S3. The documents are usually written only once, but they are updated frequently.

The reporting process takes a few hours with the use of relational queries. The reporting process must not prevent any document modifications or the addition of new documents. A solutions architect needs to implement a solution to speed up the reporting process.

Which solution will meet these requirements with the LEAST amount of change to the application code?

- A. Set up a new Amazon DocumentDB (with MongoDB compatibility) cluster that includes a read replica. Scale the read replica to generate the reports.
- B. Set up a new Amazon Aurora PostgreSQL DB cluster that includes an Aurora Replica. Issue queries to the Aurora Replica to generate the reports.
- C. Set up a new Amazon RDS for PostgreSQL Multi-AZ DB instance. Configure the reporting module to query the secondary RDS node so that the reporting module does not affect the primary node.
- D. Set up a new Amazon DynamoDB table to store the documents. Use a fixed write capacity to support new document entries. Automatically scale the read capacity to support the reports.

Correct Answer: B

Community vote distribution

B (96%)

2%

 **Guru4Cloud** Highly Voted 1 year, 7 months ago

**Selected Answer: B**

The key reasons are:

Aurora PostgreSQL provides native PostgreSQL compatibility, so minimal code changes would be required. Using an Aurora Replica separates the reporting workload from the main workload, preventing any slowdown of document updates/inserts. Aurora can auto-scale read replicas to handle the reporting load. This allows leveraging the existing PostgreSQL database without major changes. DynamoDB would require more significant rewrite of data access code. RDS Multi-AZ alone would not fully separate the workloads, as the secondary is for HA/failover more than scaling read workloads.


upvoted 12 times

 **TariqKipkemei** Highly Voted 1 year, 11 months ago

**Selected Answer: B**

Load balancing = Read replica  
High availability = Multi AZ

upvoted 8 times

 **BillaRanga** 1 year, 2 months ago

No Modifications allowed = Read Replica

upvoted 4 times

 **terminator69** Most Recent 8 months, 2 weeks ago


How in the bloody hell it's D?????

upvoted 1 times

 **TruthWS** 1 year, 1 month ago

B is correct

upvoted 2 times

 **ExamGuru727** 1 year, 1 month ago

**Selected Answer: B**

We also have a requirement for the Least amount of change to the code. Since our DB is PostgreSQL, A & D are immediately out. Multi-AZ won't help with offloading read requests, hence the answer is B ;)

upvoted 5 times

 **Buck12345** 1 year, 2 months ago

It is B

upvoted 2 times

  **Cyberkayu** 1 year, 4 months ago

**Selected Answer: C**

D. Reporting process Must not prevent = allow modification and addition of new document.

all read replica were wrong.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

How would 'issuing queries to the read replica' prevent modifications or updates?



upvoted 2 times

  **KMohsoe** 1 year, 11 months ago

**Selected Answer: A**

Why not A? :(

upvoted 1 times

  **wRhIH** 1 year, 10 months ago

"The reporting process takes a few hours with the use of RELATIONAL queries."

upvoted 4 times

  **Murtadhaceit** 1 year, 4 months ago

DocumentDB (For MongoDB) is no SQL. DynamoDB is also No SQL. Therefore, options A and D are out.

upvoted 5 times

  **lexotan** 2 years ago

**Selected Answer: B**

B is the right one. why admin does not correct these wrong answers?

upvoted 4 times

  **imvb88** 2 years ago

**Selected Answer: B**

The reporting process queries the metadata (not the documents) and use relational queries-> A, D out

C: wrong since secondary RDS node in MultiAZ setup is in standby mode, not available for querying

B: reporting using a Replica is a design pattern. Using Aurora is an exam pattern.

upvoted 5 times

  **WheracanIstart** 2 years, 1 month ago

**Selected Answer: B**

B is right..

upvoted 2 times

  **Maximus007** 2 years, 1 month ago

**Selected Answer: B**

While both B&D seems to be a relevant, ChatGPT suggest B as a correct one

upvoted 2 times

  **cegama543** 2 years, 1 month ago

**Selected Answer: B**

Option B (Set up a new Amazon Aurora PostgreSQL DB cluster that includes an Aurora Replica. Issue queries to the Aurora Replica to generate the reports) is the best option for speeding up the reporting process for a three-tier web application that includes a PostgreSQL database storing metadata from documents, while not impacting document modifications or additions, with the least amount of change to the application code.

upvoted 3 times



  **UnluckyDucky** 2 years, 1 month ago

**Selected Answer: B**

"LEAST amount of change to the application code"

Aurora is a relational database, it supports PostgreSQL and with the help of read replicas we can issue the reporting process that take several hours to the replica, therefore not affecting the primary node which can handle new writes or document modifications.

upvoted 2 times

  **Ashukaushal619** 2 years, 1 month ago



its D only ,reorrected

upvoted 2 times

  **Murtadhaceit** 1 year, 4 months ago

DynamoDB is no SQL. A and D are out!

upvoted 2 times

  **Ashukaushal619** 2 years, 1 month ago

**Selected Answer: B**

bbbbbbbbb  
upvoted 2 times

A company has a three-tier application on AWS that ingests sensor data from its users' devices. The traffic flows through a Network Load Balancer (NLB), then to Amazon EC2 instances for the web tier, and finally to EC2 instances for the application tier. The application tier makes calls to a database.

What should a solutions architect do to improve the security of the data in transit?

- A. Configure a TLS listener. Deploy the server certificate on the NLB.
- B. Configure AWS Shield Advanced. Enable AWS WAF on the NLB.
- C. Change the load balancer to an Application Load Balancer (ALB). Enable AWS WAF on the ALB.
- D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances by using AWS Key Management Service (AWS KMS).

**Correct Answer:** A

*Community vote distribution*

A (100%)

  **fruto123** Highly Voted  1 year, 1 month ago

**Selected Answer: A**

Network Load Balancers now support TLS protocol. With this launch, you can now offload resource intensive decryption/encryption from your application servers to a high throughput, and low latency Network Load Balancer. Network Load Balancer is now able to terminate TLS traffic and set up connections with your targets either over TCP or TLS protocol.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/create-tls-listener.html>

[https://exampleloadbalancer.com/nlbtls\\_demo.html](https://exampleloadbalancer.com/nlbtls_demo.html)

upvoted 23 times

  **imvb88** Highly Voted  1 year ago

**Selected Answer: A**

security of data in transit -> think of SSL/TLS. Check: NLB supports TLS

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/create-tls-listener.html>

B (DDoS), C (SQL Injection), D (EBS) is for data at rest.

upvoted 18 times

  **TariqKipkemei** Most Recent  6 months, 2 weeks ago

**Selected Answer: A**

secure data in transit = TLS

upvoted 2 times

  **Guru4Cloud** 8 months ago

**Selected Answer: A**

TLS provides encryption for data in motion over the network, protecting against eavesdropping and tampering. A valid server certificate signed by a trusted CA will provide further security.

upvoted 6 times

  **klayytech** 1 year, 1 month ago

**Selected Answer: A**

To improve the security of data in transit, you can configure a TLS listener on the Network Load Balancer (NLB) and deploy the server certificate on it. This will encrypt traffic between clients and the NLB. You can also use AWS Certificate Manager (ACM) to provision, manage, and deploy SSL/TLS certificates for use with AWS services and your internal connected resources<sup>1</sup>.

You can also change the load balancer to an Application Load Balancer (ALB) and enable AWS WAF on it. AWS WAF is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources<sup>3</sup>.

the A and C correct without transit but the need to improve the security of the data in transit? so he need SSL/TLS certificates

upvoted 3 times

  **Maximus007** 1 year, 1 month ago

**Selected Answer: A**

agree with fruto123

upvoted 4 times

A company is planning to migrate a commercial off-the-shelf application from its on-premises data center to AWS. The software has a software licensing model using sockets and cores with predictable capacity and uptime requirements. The company wants to use its existing licenses, which were purchased earlier this year.

Which Amazon EC2 pricing option is the MOST cost-effective?

- A. Dedicated Reserved Hosts
- B. Dedicated On-Demand Hosts
- C. Dedicated Reserved Instances
- D. Dedicated On-Demand Instances

**Correct Answer: A**

*Community vote distribution*

A (92%)8%

- fkie4**

Highly Voted

2 years, 1 month ago

Selected Answer: A

"predictable capacity and uptime requirements" means "Reserved"  
"sockets and cores" means "dedicated host"

upvoted 21 times
- imvb88**

Highly Voted

2 years ago

Selected Answer: A

Bring custom purchased licenses to AWS -> Dedicated Host -> C,D out  
Need cost effective solution -> "reserved" -> A

upvoted 7 times
- imvb88**

2 years ago

<https://aws.amazon.com/ec2/dedicated-hosts/>

Amazon EC2 Dedicated Hosts allow you to use your eligible software licenses from vendors such as Microsoft and Oracle on Amazon EC2, so that you get the flexibility and cost effectiveness of using your own licenses, but with the resiliency, simplicity and elasticity of AWS.

upvoted 3 times
- Hrishi\_707**

Most Recent

1 year, 1 month ago

BYOL >>> Dedicated Hosts

upvoted 4 times
- Uzbekistan**

1 year, 2 months ago

Selected Answer: A

A. Dedicated Reserved Hosts

Here's why:

License Flexibility: Dedicated Reserved Hosts allow the company to bring their existing licenses to AWS. This option enables them to continue using their purchased licenses without any additional cost or licensing changes.

Cost Optimization: Reserved Hosts offer significant cost savings compared to On-Demand pricing. By purchasing Reserved Hosts, the company can benefit from discounted hourly rates for the entire term of the reservation, which typically spans one or three years.

upvoted 3 times
- jjcode**

1 year, 2 months ago

I work with COTS applications they require a three tier architecture, its completely irrelevant and confusing to add that to the question, the key word here is licenses, since AWS wants your to use their solutions the answer to this is which of one the options solves this particular problem, in this case its dedicated hosts.

upvoted 1 times
- BillaRanga**

1 year, 2 months ago

Selected Answer: A

What is difference between dedicated host and reserved instance?

Dedicated Instance: The physical machine or underlying hardware is reserved for use for the whole account. You can have instances for different purposes on this hardware. Dedicated Host: The physical machine or the underlying hardware is reserved for "Single Use" only, eg. a certain application.

upvoted 2 times

  **BillaRanga** 1 year, 2 months ago

What is the difference between a dedicated instance and a dedicated host tenancy?

Dedicated Instance ( dedicated ) — Your instance runs on single-tenant hardware. Dedicated Host ( host ) — Your instance runs on a physical server with EC2 instance capacity fully dedicated to your use, an isolated server with configurations that you can control.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: A**

Actually the question is a bit ambiguous because there ARE "software licensing model using sockets and cores" that accept virtual sockets are cores as the base, for which C would work. But most of these license models are based on PHYSICAL sockets, thus A.

upvoted 4 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: A**

Dedicated Hosts give you visibility and control over how instances are placed on a physical server and also enable you to use your existing server-bound software licenses like Windows Server

upvoted 3 times

  **wsdasdasdqwda** 1 year, 6 months ago

Easy with one, but only 79% up to now answered correctly. It is A. Reserved because of the predictable and sockets and cores means dedicated host.

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

The correct answer is C. Dedicated Reserved Instances.

Dedicated Reserved Instances (DRIs) are the most cost-effective option for workloads that have predictable capacity and uptime requirements. DRIs offer a significant discount over On-Demand Instances, and they can be used to lock in a price for a period of time.

In this case, the company has predictable capacity and uptime requirements because the software has a software licensing model using sockets and cores. The company also wants to use its existing licenses, which were purchased earlier this year. Therefore, DRIs are the most cost-effective option.

upvoted 3 times



  **riccardoto** 1 year, 8 months ago

**Selected Answer: C**

I don't agree with people voting "A". The question reference that the COTS Application has a licensing model based on "sockets and cores". The question does not specify if it means TCP sockets (= open connections) or hardware sockets, so I assume that "TCP sockets are intended". If this is the case, sockets and cores can also remain stable with reserved instances - which are cheaper than reserved hosts.

I would go with "A" only if the question would clearly state that the COTS application has some strong dependency on physical hardware.

upvoted 1 times

  **JA2018** 5 months, 2 weeks ago

Hi Riccardoto, perhaps you might to enlighten us which software licensing model will bill customer on a "per-TCP socket" basis? :-)

Cheers.

upvoted 1 times

  **riccardoto** 1 year, 8 months ago

note: instead, if by socket we mean "CPU sockets", then A would be the right one.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

Even if "sockets" mean TCP sockets there are still the cores, thus A

upvoted 3 times

  **aragon\_saa** 2 years, 1 month ago

A

<https://www.examttopics.com/discussions/amazon/view/35818-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 2 times

  **fruto123** 2 years, 1 month ago

**Selected Answer: A**

Dedicated Host Reservations provide a billing discount compared to running On-Demand Dedicated Hosts. Reservations are available in three payment options.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html>

upvoted 4 times

  **Kenp1192** 2 years, 1 month ago

A

is the most cost effective

upvoted 2 times



A company runs an application on Amazon EC2 Linux instances across multiple Availability Zones. The application needs a storage layer that is highly available and Portable Operating System Interface (POSIX)-compliant. The storage layer must provide maximum data durability and must be shareable across the EC2 instances. The data in the storage layer will be accessed frequently for the first 30 days and will be accessed infrequently after that time.

Which solution will meet these requirements MOST cost-effectively?

- A. Use the Amazon S3 Standard storage class. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Glacier.
- B. Use the Amazon S3 Standard storage class. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Standard-Infrequent Access (S3 Standard-IA).
- C. Use the Amazon Elastic File System (Amazon EFS) Standard storage class. Create a lifecycle management policy to move infrequently accessed data to EFS Standard-Infrequent Access (EFS Standard-IA).
- D. Use the Amazon Elastic File System (Amazon EFS) One Zone storage class. Create a lifecycle management policy to move infrequently accessed data to EFS One Zone-Infrequent Access (EFS One Zone-IA).

**Correct Answer: C**

*Community vote distribution*

C (93%)

7%

- TariqKipkemei**

Highly Voted

1 year, 11 months ago

Selected Answer: C

Multi AZ = both EFS and S3 support  
Storage classes = both EFS and S3 support  
POSIX file system access = only Amazon EFS supports  
upvoted 17 times
- LazyTs**

Highly Voted

1 year, 8 months ago

Selected Answer: C

POSIX => EFS  
<https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html>  
upvoted 8 times
- Danilus**

Most Recent

5 months, 3 weeks ago

Selected Answer: C

key-Portable Operating System Interface (POSIX)-compliant.  
key -shareable across the ec2 instances  
key-accessed frequently for the first 30 days and will be accessed infrequently after  
key-storage layer that is highly available  
  
so the only service that work with POSIX is EFS in this case  
s3 can not use posix so A and B are wrong answers  
is not D because is not highly available becuse is One Zone so the answer is c  
upvoted 2 times
- zinabu**

1 year, 1 month ago

Answer:B

cause there is no life cycle policy for EFS that will work in S3 only.  
upvoted 2 times
- jjcode**

1 year, 2 months ago

"storage layer will be accessed frequently for the first 30 days and will be accessed infrequently after that time" Was the only reason they added this to trick you?

upvoted 3 times
- pentium75**

1 year, 4 months ago

Selected Answer: C

POSIX -> EFS, "maximum data durability" rules out One Zone  
upvoted 4 times
- maudsha**



1 year, 6 months ago



Selected Answer: C

Both standard and one zone have same durability.  
<https://docs.aws.amazon.com/efs/latest/ug/storage-classes.html>



Also EFS one zone can work with multiple EC2s in different AZs. But there will be a cost involved when you are accessing the EFS from a different AZ EC2. (EC2 data access charges)  
<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html>  
So if "all" EC2 instances accessing the files frequently there will be a storage cost + EC2 data access charges if you choose one zone.



So i would choose C.  
upvoted 2 times



  **beast2091** 1 year, 6 months ago  
Ans: C  
upvoted 2 times



  **baba365** 1 year, 7 months ago  
Ans: D, one-zone IA for 'most cost effective' .



<https://aws.amazon.com/efs/features/infrequent-access/>  
upvoted 1 times



  **AAAWrekng** 1 year, 6 months ago  
How does D fulfill the data durability requirement? Requirements must be met first, then consider 'most cost effective' - if you go to a tire shop, and say you want 4 new tires as cheap as possible. And they take off 4 tires and put on 2... Then they say you wanted it as cheap as possible...  
upvoted 4 times



  **Gajendr** 1 year, 4 months ago  
What about “ The application needs a storage layer that is highly available” and “application on Amazon EC2 Linux instances across multiple Availability Zones ” ?  
upvoted 2 times



  **Guru4Cloud** 1 year, 8 months ago  
**Selected Answer: C**  
Use the Amazon Elastic File System (Amazon EFS) Standard storage class. Create a lifecycle management policy to move infrequently accessed data to EFS Standard-Infrequent Access (EFS Standard-IA).  
upvoted 3 times



  **[Removed]** 1 year, 10 months ago  
**Selected Answer: D**  
Amazon Elastic File System (Amazon EFS) Standard storage class = "maximum data durability"  
upvoted 2 times



  **pentium75** 1 year, 4 months ago  
"ONE ZONE-IA" does not meet the "maximum data durability" requirement  
upvoted 2 times



  **Yadav\_Sanjay** 1 year, 10 months ago  
**Selected Answer: D**  
D - It should be cost-effective  
upvoted 2 times

  **pentium75** 1 year, 4 months ago  
But D does meet the durability requirement.  
upvoted 1 times

  **Abrar2022** 1 year, 10 months ago  
**Selected Answer: C**  
POSIX file system access = only Amazon EFS supports  
upvoted 4 times

  **imvb88** 2 years ago  
**Selected Answer: C**  
POSIX + sharable across EC2 instances --> EFS --> A, B out  
  
Instances run across multiple AZ -> C is needed.  
upvoted 2 times

  **WherecanIstart** 2 years, 1 month ago  
**Selected Answer: C**  
Linux based system points to EFS plus POSIX-compliant is also EFS related.  
upvoted 3 times

  **fkie4** 2 years, 1 month ago  
**Selected Answer: C**

"POSIX-compliant" means EFS.  
also, file system can be shared with multiple EC2 instances means "EFS"  
upvoted 5 times

  **KAUS2** 2 years, 1 month ago

**Selected Answer: C**

Option C is the correct answer .  
upvoted 2 times

A solutions architect is creating a new VPC design. There are two public subnets for the load balancer, two private subnets for web servers, and two private subnets for MySQL. The web servers use only HTTPS. The solutions architect has already created a security group for the load balancer allowing port 443 from 0.0.0.0/0. Company policy requires that each resource has the least access required to still be able to perform its tasks.



Which additional configuration strategy should the solutions architect use to meet these requirements?


- A. Create a security group for the web servers and allow port 443 from 0.0.0.0/0. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.
- B. Create a network ACL for the web servers and allow port 443 from 0.0.0.0/0. Create a network ACL for the MySQL servers and allow port 3306 from the web servers security group.
- C. Create a security group for the web servers and allow port 443 from the load balancer. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.
- D. Create a network ACL for the web servers and allow port 443 from the load balancer. Create a network ACL for the MySQL servers and allow port 3306 from the web servers security group.

**Correct Answer:** C

*Community vote distribution*

C (100%)



-   **WherecanIstart**


Highly Voted 

 1 year, 7 months ago

Selected Answer: C

Load balancer is public facing accepting all traffic coming towards the VPC (0.0.0.0/0). The web server needs to trust traffic originating from the ALB. The DB will only trust traffic originating from the Web server on port 3306 for Mysql

upvoted 6 times
-   **TheFivePips**



Most Recent 

 8 months, 1 week ago

Selected Answer: C



Option C aligns with the least access principle and provides a clear and granular control over the communication between different components in the architecture.

Option D suggests using network ACLs, but security groups are more suitable for controlling access to individual instances based on their security group membership, which is why Option C is the more appropriate choice in this contex

upvoted 3 times
-   **TariqKipkemei** 1 year ago

Selected Answer: C

Create a security group for the web servers and allow port 443 from the load balancer. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.



upvoted 3 times
-   **Guru4Cloud** 1 year, 2 months ago

Selected Answer: C

C) Create a security group for the web servers and allow port 443 from the load balancer. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.



This option follows the principle of least privilege by only allowing necessary access:

Web server SG allows port 443 from load balancer SG (not open to world)  
MySQL SG allows port 3306 only from web server SG

upvoted 4 times
-   **Guru4Cloud** 1 year, 2 months ago

Selected Answer: C



Create a security group for the web servers and allow port 443 from the load balancer. Create a security group for the MySQL servers and allow port 3306 from the web servers security group

upvoted 2 times
-   **elearningtakai** 1 year, 7 months ago

Selected Answer: C

Option C is the correct choice.

upvoted 2 times

  **fkie4** 1 year, 7 months ago

**Selected Answer: C**

Just C. plain and simple

upvoted 2 times

  **aragon\_saa** 1 year, 7 months ago

C

<https://www.examttopics.com/discussions/amazon/view/43796-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 3 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

cccccc

upvoted 1 times

An ecommerce company is running a multi-tier application on AWS. The front-end and backend tiers both run on Amazon EC2, and the database runs on Amazon RDS for MySQL. The backend tier communicates with the RDS instance. There are frequent calls to return identical datasets from the database that are causing performance slowdowns.

Which action should be taken to improve the performance of the backend?

- A. Implement Amazon SNS to store the database calls.
- B. Implement Amazon ElastiCache to cache the large datasets.
- C. Implement an RDS for MySQL read replica to cache database calls.
- D. Implement Amazon Kinesis Data Firehose to stream the calls to the database.

Correct Answer: B

Community vote distribution

B (100%)

  **elearningtakai** Highly Voted 1 year, 7 months ago

Selected Answer: B

the best solution is to implement Amazon ElastiCache to cache the large datasets, which will store the frequently accessed data in memory, allowing for faster retrieval times. This can help to alleviate the frequent calls to the database, reduce latency, and improve the overall performance of the backend tier.

upvoted 15 times

  **fruto123** Highly Voted 1 year, 7 months ago

Selected Answer: B

Key term is identical datasets from the database it means caching can solve this issue by cached in frequently used dataset from DB

upvoted 6 times

  **Manimgh** Most Recent 2 weeks, 4 days ago

Selected Answer: B

return identical datasets = Cash

upvoted 1 times

  **Ucy** 7 months ago

Selected Answer: B

Answer is B This will help reduce the frequency of calls to the database and improve overall performance by serving frequently accessed data from the cache instead of fetching it from the database every time. It's is not option C as it suggests implementing an RDS for MySQL read replica to cache database calls. While read replicas can offload read operations from the primary database instance and improve read scalability, they are primarily used for read scaling and high availability rather than caching. Read replicas are intended to handle read-heavy workloads by distributing read requests across multiple instances. However, they do not inherently cache data like ElastiCache does.

upvoted 4 times

  **Ucy** 7 months ago

Answer is B

This will help reduce the frequency of calls to the database and improve overall performance by serving frequently accessed data from the cache instead of fetching it from the database every time.

It's is not option C as it suggests implementing an RDS for MySQL read replica to cache database calls. While read replicas can offload read operations from the primary database instance and improve read scalability, they are primarily used for read scaling and high availability rather than caching.

Read replicas are intended to handle read-heavy workloads by distributing read requests across multiple instances. However, they do not inherently cache data like ElastiCache does.

upvoted 2 times

  **Bhanu1992** 7 months, 1 week ago

Keyword is identical datasets

upvoted 2 times

  **thewalker** 9 months ago

Selected Answer: B

As per Amazon Q:

ElastiCache can be used to cache datasets from queries to RDS databases. Some key points:

While creating an ElastiCache cluster from the RDS console provides convenience, the application is still responsible for leveraging the cache.

Caching query results in ElastiCache can significantly improve performance by allowing high-volume read operations to be served from cache versus hitting the database.

This is especially useful for applications with high read throughput needs, as scaling the database can become more expensive compared to scaling the cache as needs increase. ElastiCache nodes can support up to 400,000 queries per second.

Cost savings are directly proportional to read throughput - higher throughput applications see greater savings.

upvoted 2 times

🗲️ 👤 **Murtadhaceit** 11 months ago

**Selected Answer: B**

The best scenario to implement caching, identical calls to the same data sets.

upvoted 3 times

🗲️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

B) Implement Amazon ElastiCache to cache the large datasets.

The key issue is repeated calls to return identical datasets from the RDS database causing performance slowdowns.

Implementing Amazon ElastiCache for Redis or Memcached would allow these repeated query results to be cached, improving backend performance by reducing load on the database.

upvoted 4 times

🗲️ 👤 **Guru4Cloud** 1 year, 2 months ago

B) Implement Amazon ElastiCache to cache the large datasets.

The key issue is repeated calls to return identical datasets from the RDS database causing performance slowdowns.

Implementing Amazon ElastiCache for Redis or Memcached would allow these repeated query results to be cached, improving backend performance by reducing load on the database.

upvoted 2 times

🗲️ 👤 **Abrar2022** 1 year, 4 months ago

**Selected Answer: B**

Thanks Tariq for the simplified answer below:

frequent identical calls = ElastiCache

upvoted 3 times

🗲️ 👤 **TariqKipkemei** 1 year, 5 months ago

frequent identical calls = ElastiCache

upvoted 2 times

🗲️ 👤 **Mikebonsi70** 1 year, 7 months ago

Tricky question, anyway.

upvoted 2 times

🗲️ 👤 **Mikebonsi70** 1 year, 7 months ago

Yes, caching is the solution but is Elasticache compatible with RDS MySQL DB? So, what about the answer C with a DB read replica? For me it's C.

upvoted 1 times

🗲️ 👤 **aragon\_saa** 1 year, 7 months ago

B

<https://www.examttopics.com/discussions/amazon/view/27874-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 2 times

A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources. A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege.

Which combination of actions should the solutions architect take to accomplish this goal? (Choose two.)

- A. Have the deployment engineer use AWS account root user credentials for performing AWS CloudFormation stack operations.
- B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
- C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the AdministratorAccess IAM policy attached.
- D. Create a new IAM user for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
- E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using that IAM role.

**Correct Answer:** DE

Community vote distribution

DE (100%)

  **truongtx8** Highly Voted 9 months, 2 weeks ago

**Selected Answer:** DE

The answers inside the question: CloudFormation.

A is excluded since root account is never a choice for the principle of least privilege.

D, E left are the correct ones.



upvoted 5 times

  **aragon\_saa** Highly Voted 1 year, 7 months ago

D, E

<https://www.examttopics.com/discussions/amazon/view/46428-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 5 times

  **mwwt2022** 10 months ago

thank you my friend

upvoted 2 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

**Selected Answer:** DE

ABC are just giving too much access so CD are logical choices

upvoted 3 times

  **TariqKipkemei** 1 year ago

**Selected Answer:** DE

Create a new IAM user for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.

Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using that IAM role.

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer:** DE

The two actions that should be taken to follow the principle of least privilege are:

D) Create a new IAM user for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.

E) Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using that IAM role.

The principle of least privilege states that users should only be given the minimal permissions necessary to perform their job function.

upvoted 4 times

  **alexandercamachop** 1 year, 5 months ago

**Selected Answer:** DE

Option D, creating a new IAM user and adding them to a group with an IAM policy that allows AWS CloudFormation actions only, ensures that the deployment engineer has the necessary permissions to perform AWS CloudFormation operations while limiting access to other resources and actions. This aligns with the principle of least privilege by providing the minimum required permissions for their job activities.

Option E, creating an IAM role with specific permissions for AWS CloudFormation stack operations and allowing the deployment engineer to assume that role, is another valid approach. By using an IAM role, the deployment engineer can assume the role when necessary, granting them temporary permissions to perform CloudFormation actions. This provides a level of separation and limits the permissions granted to the engineer to only the required CloudFormation operations.

upvoted 3 times

  **Babaaaaa** 1 year, 5 months ago

**Selected Answer: DE**

Dddd,Eeee

upvoted 2 times

  **elearningtakai** 1 year, 7 months ago

**Selected Answer: DE**

D & E are a good choices

upvoted 2 times

  **fruto123** 1 year, 7 months ago

**Selected Answer: DE**

I agree DE

upvoted 3 times

A company is deploying a two-tier web application in a VPC. The web tier is using an Amazon EC2 Auto Scaling group with public subnets that span multiple Availability Zones. The database tier consists of an Amazon RDS for MySQL DB instance in separate private subnets. The web tier requires access to the database to retrieve product information.

The web application is not working as intended. The web application reports that it cannot connect to the database. The database is confirmed to be up and running. All configurations for the network ACLs, security groups, and route tables are still in their default states.

What should a solutions architect recommend to fix the application?

- A. Add an explicit rule to the private subnet’s network ACL to allow traffic from the web tier’s EC2 instances.
- B. Add a route in the VPC route table to allow traffic between the web tier’s EC2 instances and the database tier.
- C. Deploy the web tier's EC2 instances and the database tier’s RDS instance into two separate VPCs, and configure VPC peering.
- D. Add an inbound rule to the security group of the database tier’s RDS instance to allow traffic from the web tiers security group.

Correct Answer: D

Community vote distribution

D (97%)

- TariqKipkemei

Highly Voted

1 year, 5 months ago

Selected Answer: D

Security group defaults block all inbound traffic..Add an inbound rule to the security group of the database tier’s RDS instance to allow traffic from the web tiers security group

upvoted 15 times
- smartegnine

Highly Voted

1 year, 4 months ago

Selected Answer: D

Security Groups are tied on instance where as network ACL are tied to Subnet.

upvoted 5 times
- ExamGuru727

Most Recent

7 months, 1 week ago

Selected Answer: D

For those questioning why the answer is not A:

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-network-acls.html>

Default NACLs allow all traffic, and in this question NACLs, SGs and route tables are in their default states.

upvoted 5 times
- hgjdsh

7 months, 2 weeks ago

Selected Answer: A

I think the answer should be A. Sine the services are in different subnets, the NACL would by default block all the incoming traffic to the subnet. Security group rule wouldn't be able to override NACL rule.

upvoted 1 times
- njufi

7 months, 2 weeks ago

I selected option D as well, but I have a question regarding option A. Considering that the EC2 instances and the RDS are located in different subnets, shouldn't the network ACLs for each subnet allow traffic from one another as well? Given that the default settings for network ACLs typically block all traffic, wouldn't it be necessary to explicitly permit communication between the subnets?

upvoted 2 times
- elearningtakai

1 year, 7 months ago

Selected Answer: D

By default, all inbound traffic to an RDS instance is blocked. Therefore, an inbound rule needs to be added to the security group of the RDS instance to allow traffic from the security group of the web tier's EC2 instances.

upvoted 4 times
- Russs99

1 year, 7 months ago



Selected Answer: D

D is the correct answer

upvoted 2 times
- aragon\_saa

1 year, 7 months ago

D  
<https://www.examttopics.com/discussions/amazon/view/81445-exam-aws-certified-solutions-architect-associate-saa-c02/>  
upvoted 2 times

  **KAUS2** 1 year, 7 months ago

**Selected Answer: D**

D is correct option  
upvoted 2 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: D**

ddddddd  
upvoted 3 times

A company has a large dataset for its online advertising business stored in an Amazon RDS for MySQL DB instance in a single Availability Zone. The company wants business reporting queries to run without impacting the write operations to the production DB instance.

Which solution meets these requirements?

- A. Deploy RDS read replicas to process the business reporting queries.
- B. Scale out the DB instance horizontally by placing it behind an Elastic Load Balancer.
- C. Scale up the DB instance to a larger instance type to handle write operations and queries.
- D. Deploy the DB instance in multiple Availability Zones to process the business reporting queries.

**Correct Answer: A**

*Community vote distribution*

A (100%)

- mwwt2022**

Highly Voted

10 months ago

**Selected Answer: A**

reporting queries to run without impacting the write operations -> read replicas  
upvoted 5 times
- Guru4Cloud**

Most Recent

1 year, 2 months ago

**Selected Answer: A**

A) Deploy RDS read replicas to process the business reporting queries.  
  
The key points are:  
  
RDS read replicas allow read-only copies of the production DB instance to be created  
Queries to the read replica don't affect the source DB instance performance  
This isolates reporting queries from production traffic and write operations  
So using RDS read replicas is the best way to meet the requirements of running reporting queries without impacting production write operations.  
upvoted 4 times
- james2033**

1 year, 3 months ago

**Selected Answer: A**

"single AZ", "large dataset", "Amazon RDS for MySQL database". Want "business report queries". --> Solution "Read replicas", choose A.  
upvoted 2 times
- antropaws**

1 year, 5 months ago

**Selected Answer: A**

No doubt A.  
upvoted 3 times
- TariqKipkemei**

1 year, 5 months ago

Load balance read operations = read replicas  
upvoted 2 times
- TariqKipkemei**

1 year ago

reports=read replica  
upvoted 2 times
- KAUS2**

1 year, 7 months ago


**Selected Answer: A**

Option "A" is the right answer . Read replica use cases - You have a production database that is taking on normal load & You want to run a reporting application to run some analytics
  - You create a Read Replica to run the new workload there
  - The production application is unaffected
  - Read replicas are used for SELECT (=read) only kind of statements (not INSERT, UPDATE, DELETE)upvoted 3 times
- [Removed]**

1 year, 7 months ago

**Selected Answer: A**

aaaaaaaaaaaa  
upvoted 2 times

  **cegama543** 1 year, 7 months ago

**Selected Answer: A**

option A is the best solution for ensuring that business reporting queries can run without impacting write operations to the production DB instance.

upvoted 4 times

A company hosts a three-tier ecommerce application on a fleet of Amazon EC2 instances. The instances run in an Auto Scaling group behind an Application Load Balancer (ALB). All ecommerce data is stored in an Amazon RDS for MariaDB Multi-AZ DB instance.

The company wants to optimize customer session management during transactions. The application must store session data durably.



Which solutions will meet these requirements? (Choose two.)


- A. Turn on the sticky sessions feature (session affinity) on the ALB.
- B. Use an Amazon DynamoDB table to store customer session information.
- C. Deploy an Amazon Cognito user pool to manage user session information.
- D. Deploy an Amazon ElastiCache for Redis cluster to store customer session information.
- E. Use AWS Systems Manager Application Manager in the application to manage user session information.

**Correct Answer: AD**

*Community vote distribution*

AD (38%)	BD (32%)	AB (28%)
----------	----------	----------

-   **fruto123**



Highly Voted 

 2 years, 1 month ago

Selected Answer: AD



It is A and D. Proof is in link below.

<https://aws.amazon.com/caching/session-management/>



upvoted 28 times
-   **babayomi** 6 months, 1 week ago

Thanks,



First consider the durability requirement of a session. It's never meant to be kept after the session terminates, tha5 said. You wouldn't need a pamanent database for a session. You would only require a cache, so that it can expire after a while when the session has already completed. That is why D is correct. A is required in the first place in order to initialize the use of the cache.


upvoted 2 times
-   **pentium75** 1 year, 3 months ago

This doesn't say anything about durability

upvoted 4 times
-   **babayomi** 6 months, 1 week ago

First consider the durability requirement of a session. It's never meant to be kept after the session terminates, tha5 said. You wouldn't need a pamanent database for a session. You would only require a cache, so that it can expire after a while when the session has already completed. That is why D is correct. A is required in the first place in order to initialize the use of the cache.



upvoted 2 times
-   **Marco\_St**

Highly Voted 

 1 year, 4 months ago



Selected Answer: BD


I did not get why A is most voted? The question did not mention anything about fixed routing target so the ALB should route traffic randomly to each server. Then we just need to provide cache session management to avoid session lost issue instead of using sticky session.

upvoted 14 times
-   **babayomi** 6 months, 1 week ago

AD

First consider the durability requirement of a session. It's never meant to be kept after the session terminates, tha5 said. You wouldn't need a pamanent database for a session. You would only require a cache, so that it can expire after a while when the session has already completed. That is why D is correct. A is required in the first place in order to initialize the use of the cache.



upvoted 1 times
-   **chest\_jd**

Most Recent 

 3 weeks, 2 days ago

Selected Answer: BD

First of all - application MUST STORE session DATA. We should get some storage on backend side. Just Band D can do it.

upvoted 1 times
-   **bigjft** 2 months ago

Selected Answer: BD

Avoid sticky sessions:

While using sticky sessions on an Application Load Balancer (ALB) might seem like a solution, it can hinder scalability as it ties a user to a specific

instance, which could become a bottleneck if traffic spikes.  
Dynamodb and Elastic cache are best i think

upvoted 1 times

  **zdi561** 3 months ago

**Selected Answer: BD**

A is not right, the question is to where to save user session, the sticky session is used for ALB to send the request of the same client to the same target.

upvoted 1 times

  **FlyingHawk** 5 months ago

**Selected Answer: BD**

People select AD because they think it is combine A and D to make a final solution. if the expectation is that each selected option should independently meet the requirements, then B and D. I think "Which solutions will meet these requirements?" should mean two independent solutions.

upvoted 1 times

  **Arad** 5 months, 2 weeks ago

**Selected Answer: BD**

session management during transaction ---> DynamoDB and ElastiCache

upvoted 1 times

  **Sergantus** 6 months ago

**Selected Answer: BD**

We are not looking here for a combination of options, just two independent solutions.

Because of the "durability" requirement -

B - DynamoDB is durable



D - ElastiCache for Redis although mostly used as in-memory cache, ElastiCache for Redis provides options for data persistence - RDB (Redis Database Backup) and AOF (Append-Only File). RDB periodically saves a snapshot of the dataset to disk, while AOF logs every write operation and replays it in case of a restart. ElastiCache for Redis also allows configuring replication groups with Multi-AZ and read replicas, providing additional redundancy and failover options.

upvoted 2 times

  **FlyingHawk** 5 months ago

I agree with your comments

upvoted 1 times

  **babayomi** 6 months, 1 week ago

AD

First consider the durability requirement of a session. It's never meant to be kept after the session terminates, tha5 said. You wouldn't need a pamanent database for a session. You would only require a cache, so that it can expire after a while when the session has already completed. That is why D is correct. A is required in the first place in order to initialize the use of the cache.

upvoted 1 times

  **Uzbekistan** 1 year, 2 months ago

**Selected Answer: BD**

Option A suggests using sticky sessions (session affinity) on the Application Load Balancer (ALB). While sticky sessions can help route requests from the same client to the same backend server, it doesn't directly address the requirement for durable storage of session data. Sticky sessions are typically used to maintain session state at the load balancer level, but they do not provide data durability in case of server failures or restarts. Option A - is not correct ! ! !


So answer is option B and D ! ! !

upvoted 3 times

  **babayomi** 6 months, 1 week ago



First consider the durability requirement of a session. It's never meant to be kept after the session terminates, tha5 said. You wouldn't need a pamanent database for a session. You would only require a cache, so that it can expire after a while when the session has already completed. That is why D is correct. A is required in the first place in order to initialize the use of the cache.

upvoted 1 times

  **jjcode** 1 year, 2 months ago

why does it matter to store user sessions durably? they EXPIRE, why would a company care about storing user sessions, thats not something thats done in the real world, those things are usually data dumped, or overwritten with new session tokens LOL, this whole question is &^%&\*\$^\$#@%^

upvoted 4 times

  **tuso** 1 year, 3 months ago

I think the question is intended to mean "Combination of services", as some answers say "to store" or "to manage". So i am going for A+B, as sticky sessions are intended to manage the sessions and DynamoDB to store durably.

upvoted 1 times



  **pentium75** 1 year, 4 months ago

**Selected Answer: AB**

Going for AB. Sticky Sessions to "optimize customer session management during transactions" and DynamoDB to "store session data durably".

D, ElastiCache does NOT allow "durable" storage. Just because there's an article that contains both words "ElastiCache" and "durable" does not prove the contrary.



C and E, Cognito and Systems Manager, have nothing to do with the issue.  
upvoted 5 times



  **dkw2342** 1 year, 2 months ago  
I agree that ElastiCache for Redis is not a durable KV store.  
  
But what about the phrasing?

"Which solutions will meet these requirements? (Choose two.)" Solutions (plural) implies two ways to \*independently\* fulfill the requirements. If you're supposed to select a combination of options, it's usually phrased like this: "Which combination of solutions ..."  
upvoted 3 times



  **avdxeqtr** 1 year, 3 months ago  
<https://aws.amazon.com/blogs/developer/elasticache-as-an-asp-net-session-store/>



Amazon ElastiCache for Redis is highly suited as a session store to manage session information such as user authentication tokens, session state, and more. Simply use ElastiCache for Redis as a fast key-value store with appropriate TTL on session keys to manage your session information. Session management is commonly required for online applications, including games, e-commerce websites, and social media platforms.  
upvoted 2 times



  **avdxeqtr** 1 year, 3 months ago  
Correct link: <https://aws.amazon.com/elasticache/redis/>  
upvoted 3 times

  **m\_y\_s** 1 year, 5 months ago  
**Selected Answer: BD**



I don't understand what Sticky Session has to do with session storage. For the intent of the problem, I think DynamoDB and Redis are appropriate.  
upvoted 6 times



  **pentium75** 1 year, 4 months ago  
"Session storage" is not the only requirement here. It is about 'optimizing customer session management during transactions', obviously it makes sense to host customer sessions on same node to easy the session management.  
upvoted 4 times



  **Sergantus** 6 months ago  
Up until that node fails  
upvoted 2 times


  **daniel1** 1 year, 6 months ago  
**Selected Answer: BD**



Chatgpt4 says B and D  
Option A (Sticky sessions) is more for ensuring that a client's requests are sent to the same target once a session is established, but it doesn't provide a mechanism for durable session data storage across multiple instances. Option C (Amazon Cognito) is more for user identity management rather than session data storage during transactions. Option E (AWS Systems Manager Application Manager) is not a suitable or standard choice for session management in applications.  
upvoted 4 times

  **pentium75** 1 year, 4 months ago  
Answers starting with "ChatGPT says ..." are usually wrong.  
  
In that case, B and D solve the same part of the requirement (storing session data), just B is durable (as required) while D is not durable (thus failing to meet the requirement). We still need to 'optimize customer session management'.  
upvoted 4 times

  **TariqKipkemei** 1 year, 6 months ago  
**Selected Answer: AD**  
Well, this documentation says it all. Option A is obvious, and D ElastiCache for Redis, can even support replication in case of node failure/session data loss.  
<https://aws.amazon.com/caching/session-management/>  
upvoted 4 times

  **pentium75** 1 year, 3 months ago  
ElastiCache can be HA and supports replication, but it remains a cache, which is by definition not durable.  
upvoted 3 times

  **Sergantus** 6 months ago  
You can configure Redis with persistence options  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago  
**Selected Answer: AD**

It is A and D. Proof is in link below.  
  
<https://aws.amazon.com/caching/session-management/>

upvoted 2 times

  **pentium75** 1 year, 3 months ago

That does not say anything about durability.

upvoted 2 times

A company needs a backup strategy for its three-tier stateless web application. The web application runs on Amazon EC2 instances in an Auto Scaling group with a dynamic scaling policy that is configured to respond to scaling events. The database tier runs on Amazon RDS for PostgreSQL. The web application does not require temporary local storage on the EC2 instances. The company’s recovery point objective (RPO) is 2 hours.

The backup strategy must maximize scalability and optimize resource utilization for this environment.

Which solution will meet these requirements?



- A. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances and database every 2 hours to meet the RPO.
- B. Configure a snapshot lifecycle policy to take Amazon Elastic Block Store (Amazon EBS) snapshots. Enable automated backups in Amazon RDS to meet the RPO.
- C. Retain the latest Amazon Machine Images (AMIs) of the web and application tiers. Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO.
- D. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances every 2 hours. Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO.


**Correct Answer:** C

*Community vote distribution*

C (88%)

10%

-   **elearningtakai**



Highly Voted 

 2 years, 1 month ago



Selected Answer: C

that if there is no temporary local storage on the EC2 instances, then snapshots of EBS volumes are not necessary. Therefore, if your application does not require temporary storage on EC2 instances, using AMIs to back up the web and application tiers is sufficient to restore the system after a failure.



Snapshots of EBS volumes would be necessary if you want to back up the entire EC2 instance, including any applications and temporary data stored on the EBS volumes attached to the instances. When you take a snapshot of an EBS volume, it backs up the entire contents of that volume. This ensures that you can restore the entire EC2 instance to a specific point in time more quickly. However, if there is no temporary data stored on the EBS volumes, then snapshots of EBS volumes are not necessary.

upvoted 35 times
-   **MssP** 2 years, 1 month ago



I think "temporal local storage" refers to "instance store", no instance store is required. EBS is durable storage, not temporal.


upvoted 4 times
-   **MssP** 2 years, 1 month ago

Look at the first paragraph. <https://repost.aws/knowledge-center/instance-store-vs-ebs>

upvoted 2 times
-   **MatAlves** 7 months, 3 weeks ago

Considering it's a "stateless web application", that would still be no reason to back up the EBS volumes.



upvoted 2 times
-   **CloudForFun**


Highly Voted 

 2 years, 1 month ago

Selected Answer: C

The web application does not require temporary local storage on the EC2 instances => No EBS snapshot is required, retaining the latest AMI is enough.

upvoted 14 times
-   **Mikado211**



Most Recent 

 1 year, 4 months ago

Selected Answer: C

The web application does not require temporary local storage on the EC2 instances so we do not care about ECS. We only need two things here , the image of the instance (AMI) and a database backup.

C

upvoted 3 times
-   **TariqKipkemei** 1 year, 6 months ago

Selected Answer: C

"The web application does not require temporary local storage on the EC2 instances" rules out any option to back up the EC2 EBS volumes.

upvoted 2 times

  **darekw** 1 year, 9 months ago

Question says: ...stateless web application.. that means application doesn't store any data, so no EBS required  
upvoted 2 times

  **kruasan** 2 years ago

**Selected Answer: C**

Since the application has no local data on instances, AMIs alone can meet the RPO by restoring instances from the most recent AMI backup. When combined with automated RDS backups for the database, this provides a complete backup solution for this environment. The other options involving EBS snapshots would be unnecessary given the stateless nature of the instances. AMIs provide all the backup needed for the app tier.

This uses native, automated AWS backup features that require minimal ongoing management:  
- AMI automated backups provide point-in-time recovery for the stateless app tier.  
- RDS automated backups provide point-in-time recovery for the database.

upvoted 4 times


  **neosis91** 2 years ago

**Selected Answer: B**

BBBBBBBBBBB  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

Why back up EBS volumes of the autoscaled instances?  
upvoted 2 times

  **Rob1L** 2 years, 1 month ago

**Selected Answer: D**

I vote for D  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

Why back up EBS volumes of the autoscaled instances?  
upvoted 2 times

  **CapJackSparrow** 2 years, 1 month ago

**Selected Answer: C**

makes more sense.  
upvoted 3 times

  **nileshlg** 2 years, 1 month ago

**Selected Answer: C**

Answer is C. Keyword to notice "Stateless"  
upvoted 3 times

  **cra2yk** 2 years, 1 month ago

**Selected Answer: C**

why B? I mean "stateless" and "does not require temporary local storage" have indicate that we don't need to take snapshot for ec2 volume.  
upvoted 4 times

  **ktulu2602** 2 years, 1 month ago

**Selected Answer: B**

Option B is the most appropriate solution for the given requirements.



With this solution, a snapshot lifecycle policy can be created to take Amazon Elastic Block Store (Amazon EBS) snapshots periodically, which will ensure that EC2 instances can be restored in the event of an outage. Additionally, automated backups can be enabled in Amazon RDS for PostgreSQL to take frequent backups of the database tier. This will help to minimize the RPO to 2 hours.

Taking snapshots of Amazon EBS volumes of the EC2 instances and database every 2 hours (Option A) may not be cost-effective and efficient, as this approach would require taking regular backups of all the instances and volumes, regardless of whether any changes have occurred or not. Retaining the latest Amazon Machine Images (AMIs) of the web and application tiers (Option C) would provide only an image backup and not a data backup, which is required for the database tier. Taking snapshots of Amazon EBS volumes of the EC2 instances every 2 hours and enabling automated backups in Amazon RDS and using point-in-time recovery (Option D) would result in higher costs and may not be necessary to meet the RPO requirement of 2 hours.  
upvoted 4 times

  **pentium75** 1 year, 4 months ago

Why back up EBS volumes of the autoscaled instances?

"Retaining the latest Amazon Machine Images (AMIs) of the web and application tiers (Option C) would provide only an image backup and not a data backup, which is required for the database tier." False because option C also includes "automated backups in Amazon RDS".  
upvoted 2 times

  **cegama543** 2 years, 1 month ago

**Selected Answer: B**

B. Configure a snapshot lifecycle policy to take Amazon Elastic Block Store (Amazon EBS) snapshots. Enable automated backups in Amazon RDS to meet the RPO.

The best solution is to configure a snapshot lifecycle policy to take Amazon Elastic Block Store (Amazon EBS) snapshots, and enable automated backups in Amazon RDS to meet the RPO. An RPO of 2 hours means that the company needs to ensure that the backup is taken every 2 hours to minimize data loss in case of a disaster. Using a snapshot lifecycle policy to take Amazon EBS snapshots will ensure that the web and application tier can be restored quickly and efficiently in case of a disaster. Additionally, enabling automated backups in Amazon RDS will ensure that the database tier can be restored quickly and efficiently in case of a disaster. This solution maximizes scalability and optimizes resource utilization because it uses automated backup solutions built into AWS.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

No need to back up the EBS volumes of autoscaled instances.

upvoted 3 times

A company wants to deploy a new public web application on AWS. The application includes a web server tier that uses Amazon EC2 instances. The application also includes a database tier that uses an Amazon RDS for MySQL DB instance.

The application must be secure and accessible for global customers that have dynamic IP addresses.

How should a solutions architect configure the security groups to meet these requirements?

- A. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers.
- B. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers.
- C. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the IP addresses of the customers.
- D. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0. Configure the security group for the DB instance to allow inbound traffic on port 3306 from 0.0.0.0/0.

**Correct Answer: A**

*Community vote distribution*

A (86%)

14%

- Bhanu1992**

Highly Voted

7 months, 1 week ago

The keyword is dynamic IPs from the customer, then B, C out, D out due to 0.0.0.0/0

upvoted 6 times
- awsgeek75**

Highly Voted

9 months, 2 weeks ago

**Selected Answer: A**

"The application must be secure and accessible for global customers that have dynamic IP addresses." This just means "anyone" so BC are wrong as you cannot know in advance about the dynamic IP addresses. D is just opening the DB to the internet.

A is most secure as web is open to internet and db is open to web only.

upvoted 6 times
- Guru4Cloud**

Most Recent

1 year, 2 months ago

**Selected Answer: A**

It allows HTTPS access from any public IP address, meeting the requirement for global customer access. HTTPS provides encryption for secure communication. And for the database security group, only allowing inbound port 3306 from the web server security group properly restricts access to only the resources that need it.

upvoted 4 times
- jayce5**

1 year, 5 months ago

**Selected Answer: A**

Should be A since the customer IPs are dynamically.

upvoted 2 times
- antropaws**

1 year, 5 months ago

**Selected Answer: A**

A no doubt.

upvoted 3 times
- omoakin**

1 year, 5 months ago

BBBBBBBBBBBBBBBBBBBBBBB

from customers IPs

upvoted 1 times
- MostafaWardany**

1 year, 4 months ago

Correct answer A, customer dynamic IPs ==>> 443 from 0.0.0.0/0

upvoted 3 times
- TariqKipkemei**

1 year, 5 months ago

**Selected Answer: A**

dynamic source ips = allow all traffic - Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0. Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers.



upvoted 2 times

  **elearningtakai** 1 year, 7 months ago

**Selected Answer: A**

If the customers have dynamic IP addresses, option A would be the most appropriate solution for allowing global access while maintaining security.

upvoted 4 times

  **Kenzo** 1 year, 7 months ago

Correct answer is A.  
B and C are out.  
D is out because it is accepting traffic from every where instead of from web servers only

upvoted 4 times

  **Grace83** 1 year, 7 months ago

A is correct

upvoted 3 times

  **WherecanIstart** 1 year, 7 months ago

**Selected Answer: B**



Keyword dynamic ...A is the right answer. If the IP were static and specific, B would be the right answer

upvoted 4 times

  **pentium75** 10 months, 1 week ago

Then why voted B?

upvoted 2 times

  **boxu03** 1 year, 7 months ago

**Selected Answer: A**

aaaaaaa

upvoted 1 times

  **kprakashbehera** 1 year, 7 months ago

**Selected Answer: A**

Ans - A

upvoted 1 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: A**

aaaaaa

upvoted 1 times

A payment processing company records all voice communication with its customers and stores the audio files in an Amazon S3 bucket. The company needs to capture the text from the audio files. The company must remove from the text any personally identifiable information (PII) that belongs to customers.

What should a solutions architect do to meet these requirements?

- A. Process the audio files by using Amazon Kinesis Video Streams. Use an AWS Lambda function to scan for known PII patterns.
- B. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start an Amazon Textract task to analyze the call recordings.
- C. Configure an Amazon Transcribe transcription job with PII redaction turned on. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start the transcription job. Store the output in a separate S3 bucket.
- D. Create an Amazon Connect contact flow that ingests the audio files with transcription turned on. Embed an AWS Lambda function to scan for known PII patterns. Use Amazon EventBridge to start the contact flow when an audio file is uploaded to the S3 bucket.

Correct Answer: C

Community vote distribution

C (100%)

- TariqKipkemei

Highly Voted

6 months, 2 weeks ago

Selected Answer: C

speech to text = Amazon Transcribe

upvoted 5 times
- Guru4Cloud

Highly Voted

8 months, 1 week ago

Selected Answer: C

Amazon Transcribe is a service provided by Amazon Web Services (AWS) that converts speech to text using automatic speech recognition (ASR) technology

upvoted 5 times
- james2033

Most Recent

9 months, 3 weeks ago

Selected Answer: C

AWS Transcribe <https://aws.amazon.com/transcribe/> . Redacting or identifying (Personally identifiable instance) PII in real-time stream <https://docs.aws.amazon.com/transcribe/latest/dg/pii-redaction-stream.html> .

upvoted 2 times
- SimiTik

1 year ago

C

Amazon Transcribe is a service provided by Amazon Web Services (AWS) that converts speech to text using automatic speech recognition (ASR) technology. gtp

upvoted 4 times
- elearningtakai

1 year, 1 month ago

Selected Answer: C

Option C is the most suitable solution as it suggests using Amazon Transcribe with PII redaction turned on. When an audio file is uploaded to the S3 bucket, an AWS Lambda function can be used to start the transcription job. The output can be stored in a separate S3 bucket to ensure that the PII redaction is applied to the transcript. Amazon Transcribe can redact PII such as credit card numbers, social security numbers, and phone numbers.

upvoted 4 times
- WherecanIstart

1 year, 1 month ago

Selected Answer: C

C for sure.....

upvoted 2 times
- WherecanIstart

1 year, 1 month ago

C for sure

upvoted 2 times
- boxu03

1 year, 1 month ago

Selected Answer: C

CCCCCCCC

upvoted 2 times

  **Ruhi02** 1 year, 1 month ago

answer c

upvoted 2 times

  **KAUS2** 1 year, 1 month ago

**Selected Answer: C**

Option C is correct..

upvoted 2 times

A company is running a multi-tier ecommerce web application in the AWS Cloud. The application runs on Amazon EC2 instances with an Amazon RDS for MySQL Multi-AZ DB instance. Amazon RDS is configured with the latest generation DB instance with 2,000 GB of storage in a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume. The database performance affects the application during periods of high demand.

A database administrator analyzes the logs in Amazon CloudWatch Logs and discovers that the application performance always degrades when the number of read and write IOPS is higher than 20,000.

What should a solutions architect do to improve the application performance?

- A. Replace the volume with a magnetic volume.
- B. Increase the number of IOPS on the gp3 volume.
- C. Replace the volume with a Provisioned IOPS SSD (io2) volume.
- D. Replace the 2,000 GB gp3 volume with two 1,000 GB gp3 volumes.

**Correct Answer: C**

*Community vote distribution*

C (37%)

D (34%)

B (29%)

  **Bezha**  2 years, 1 month ago

**Selected Answer: D**

A - Magnetic Max IOPS 200 - Wrong  
B - gp3 Max IOPS 16000 per volume - Wrong  
C - RDS not supported io2 - Wrong  
D - Correct; 2 gp3 volume with 16 000 each  $2 \times 16000 = 32\ 000$  IOPS  
upvoted 37 times

  **FlyingHawk** 5 months ago

gp3 Max IOPS 16000 per volume is for SQL server, for MySQL server it is 64000 per volume, so B should work.  
[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)  
upvoted 1 times

  **dkw2342** 1 year, 1 month ago

I really wonder how this answer can be the top answer. How would it even be possible to provision multiple gp3 volumes for RDS? RDS manages the storage, we have no influence on the number of volumes.

\*Striping\* is something that RDS does automatically depending on storage class and volume size: "When you select General Purpose SSD or Provisioned IOPS SSD, depending on the engine selected and the amount of storage requested, Amazon RDS automatically stripes across multiple volumes to enhance performance (...)"

For MariaDB with 400 to 64,000 GiB of gp3 storage, RDS automatically provisions 4 volumes. This gives us 12,000 IOPS \*baseline\* and can be increased up to 64,000 \*provisioned\* IOPS.

RDS does not support io2.

Therefore: Option B  
upvoted 6 times

  **dkw2342** 1 year, 1 month ago



PS: [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)  
upvoted 2 times

  **zits88** 1 year ago

It must be that io2 was originally not supported by RDS, because I see this untruth reposted everywhere. It totally is.  
upvoted 4 times

  **joechen2023** 1 year, 10 months ago

<https://repost.aws/knowledge-center/ebs-volume-type-differences>  
RDS does support io2  
upvoted 3 times

  **wRhIH** 1 year, 10 months ago

that Link is to EBS instead of RDS

upvoted 6 times

  **baba365** 1 year, 7 months ago

'the application performance always degrades when the number of read and write IOPS is higher than 20,000' ... question didn't say read and write IOPs can't be higher than 32,000. Answer: C if it's based on performance and not cost related.

'Amazon RDS provides three storage types: General Purpose SSD (also known as gp2 and gp3), Provisioned IOPS SSD (also known as io1), and magnetic (also known as standard). They differ in performance characteristics and price, which means that you can tailor your storage performance and cost to the needs of your database workload.'

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)

upvoted 1 times

  **Michal\_L\_95** Highly Voted 2 years, 1 month ago

**Selected Answer: B**

It can not be option C as RDS does not support io2 storage type (only io1).

Here is a link to the RDS storage documentation: [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)

Also it is not the best option to take Magnetic storage as it supports max 1000 IOPS.

I vote for option B as gp3 storage type supports up to 64 000 IOPS where question mentioned with problem at level of 20 000.

upvoted 15 times

  **joechen2023** 1 year, 10 months ago

check the link below <https://repost.aws/knowledge-center/ebs-volume-type-differences>

it states:

General Purpose SSD volumes are good for a wide variety of transactional workloads that require less than the following:

16,000 IOPS

1,000 MiB/s of throughput



160-TiB volume size

upvoted 2 times

  **GalileoEC2** 2 years, 1 month ago

is this true? Amazon RDS (Relational Database Service) supports the Provisioned IOPS SSD (io2) storage type for its database instances. The io2 storage type is designed to deliver predictable performance for critical and highly demanding database workloads. It provides higher durability, higher IOPS, and lower latency compared to other Amazon EBS (Elastic Block Store) storage types. RDS offers the option to choose between the General Purpose SSD (gp3) and Provisioned IOPS SSD (io2) storage types for database instances.

upvoted 3 times

  **1rob** 1 year, 3 months ago

Please add a reference where it states that io2 is supported by RDS.

upvoted 1 times

  **zits88** 1 year ago

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html) - Right there in the first paragraph: "Amazon RDS provides three storage types: General Purpose SSD (also known as gp2 and gp3), Provisioned IOPS SSD (also known as io1 and io2 Block Express), and magnetic (also known as standard). They differ in performance characteristics and price, which means that you can tailor your storage performance and cost to the needs of your database workload. You can create Db2, MySQL, MariaDB, Oracle, and PostgreSQL RDS DB instances with up to 64 tebibytes (TiB) of storage. "

upvoted 7 times

  **zdi561** Most Recent 3 months ago

**Selected Answer: C**

RDS now support io2

upvoted 2 times

  **EllenLiu** 4 months, 2 weeks ago

**Selected Answer: C**

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html#Concepts.Storage.GeneralSSD](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html#Concepts.Storage.GeneralSSD)

upvoted 1 times

  **FlyingHawk** 5 months ago

**Selected Answer: B**

I selected C initially, but after read the doc [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html), noticed maximum IOPS for gp3 volume is 64K (16K for SQL server), For an RDS for MySQL workload, increasing the provisioned IOPS on the gp3 volume to exceed the application's peak demand (e.g., 30,000 IOPS) is a cost-effective and scalable solution.

upvoted 1 times

  **LeonSauveterre** 5 months ago

**Selected Answer: C**

How on earth would the answer be D? Splitting the gp3 volume into two does not solve the IOPS limitation because each volume would still be subject to the gp3 performance cap. The core of the crux is still there!

C - io2 volumes are exactly designed for workloads requiring consistent, high-performance storage.

upvoted 1 times

  **Mish** 5 months, 3 weeks ago

**Selected Answer: C**

Provisioned IOPS SSDs (io2) are specifically designed to deliver sustained high performance and low latency (RDS is supported in IO2). They can handle more than 20,000 IOPS.

upvoted 2 times

  **studydue** 6 months ago

Updated 2024

Answer: C



upvoted 4 times

  **ashishs174** 7 months, 2 weeks ago

Answer is C, io2 volumes are supported



<https://aws.amazon.com/blogs/aws/amazon-rds-now-supports-io2-block-express-volumes-for-mission-critical-database-workloads/>

upvoted 2 times

  **MatAlves** 7 months, 3 weeks ago

Nice to see that everyone just picked a different answer...



upvoted 3 times

  **ChymKuBoy** 8 months, 3 weeks ago

**Selected Answer: B**

B for sure

upvoted 1 times

  **example\_** 9 months ago

**Selected Answer: C**

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)

upvoted 5 times

  **FrozenCarrot** 9 months, 4 weeks ago

**Selected Answer: C**

Now EBS support io2.

upvoted 6 times

  **theamachine** 10 months, 2 weeks ago

**Selected Answer: C**



Provisioned IOPS SSDs (io2) are specifically designed to deliver sustained high performance and low latency (RDS is supported in IO2). They can handle more than 20,000 IOPS.

upvoted 7 times

  **FlyingHawk** 5 months ago

C is more expensive and overkill.

upvoted 1 times


  **Lin878** 10 months, 3 weeks ago

**Selected Answer: C**

It should be "C" right, now.

<https://aws.amazon.com/blogs/aws/amazon-rds-now-supports-io2-block-express-volumes-for-mission-critical-database-workloads/>

upvoted 5 times

  **[Removed]** 10 months, 3 weeks ago

C is the correct one

EBS Volume Types Use cases

Provisioned IOPS (PIOPS) SSD

- Critical business applications with sustained IOPS performance
- Or applications that need more than 16,000 IOPS
- Great for databases workloads (sensitive to storage perf and consistency)
- io1/io2 (4 GiB - 16 TiB):
- Max PIOPS: 64,000 for Nitro EC2 instances & 32,000 for other
- Can increase PIOPS independently from storage size
- io2 have more durability and more IOPS per GiB (at the same price as io1)
- io2 Block Express (4 GiB – 64 TiB):
- Sub-millisecond latency
- Max PIOPS: 256,000 with an IOPS:GiB ratio of 1,000:1

upvoted 3 times

  **Scheldon** 11 months, 2 weeks ago

**Selected Answer: C**

Per the newest info it should be C right now

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Storage.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html)

upvoted 4 times

An IAM user made several configuration changes to AWS resources in their company's account during a production deployment last week. A solutions architect learned that a couple of security group rules are not configured as desired. The solutions architect wants to confirm which IAM user was responsible for making changes.

Which service should the solutions architect use to find the desired information?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS CloudTrail
- D. AWS Config

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **cegama543** Highly Voted  1 year, 7 months ago

**Selected Answer: C**  
C. AWS CloudTrail

The best option is to use AWS CloudTrail to find the desired information. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of AWS account activities. CloudTrail can be used to log all changes made to resources in an AWS account, including changes made by IAM users, EC2 instances, AWS management console, and other AWS services. By using CloudTrail, the solutions architect can identify the IAM user who made the configuration changes to the security group rules.

upvoted 14 times

  **Manimgh** Most Recent  2 weeks, 4 days ago

**Selected Answer: C**  
which IAM user was responsible = CloudTrail

upvoted 1 times

  **BatVanyo** 6 months, 2 weeks ago

**Selected Answer: C**  
I was initially a bit confused on what Config and CloudTrail actually do, as both can be used to track configuration changes. However, this explanation is probably the best one I have come across so far:  
"Config reports on what has changed, whereas CloudTrail reports on who made the change, when, and from which location"



Since the question is which IAM user was responsible for making the changes, the answer is CloudTrail.

upvoted 4 times

  **d401c0d** 7 months, 1 week ago

**Selected Answer: C**  
CloudTrail = which user made which api calls. This is used for audit purpose.

upvoted 3 times

  **sheq** 10 months, 4 weeks ago

This question is the same with the question 388, isn't it?

upvoted 1 times

  **kambarami** 1 year, 1 month ago

This is how you know not to trust the moderators with their answers.

upvoted 1 times

  **Wayne23Fang** 1 year, 1 month ago

There is an article "How to use AWS Config and CloudTrail to find who made changes to a resource" in aws blog. Given CloudTrail provided AWS config original info, it seems for this particular one, C is better than AWS config.

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago



**Selected Answer: C**  
AWS CloudTrail is the correct service to use here to identify which user was responsible for the security group configuration changes

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: C**

AWS CloudTrail  
upvoted 2 times

  **Bezha** 1 year, 7 months ago

**Selected Answer: C**

AWS CloudTrail  
upvoted 2 times

  **[Removed]** 1 year, 7 months ago



**Selected Answer: C**

C. AWS CloudTrail  
upvoted 2 times

  **kprakashbehera** 1 year, 7 months ago

**Selected Answer: C**

CloudTrail logs will tell who did that  
upvoted 3 times

  **KAUS2** 1 year, 7 months ago

**Selected Answer: C**

Option "C" AWS CloudTrail is correct.  
upvoted 3 times

  **Nithin1119** 1 year, 7 months ago

cccccc  
upvoted 2 times

A company has implemented a self-managed DNS service on AWS. The solution consists of the following:

- Amazon EC2 instances in different AWS Regions
- Endpoints of a standard accelerator in AWS Global Accelerator

The company wants to protect the solution against DDoS attacks.

What should a solutions architect do to meet this requirement?

- A. Subscribe to AWS Shield Advanced. Add the accelerator as a resource to protect.
- B. Subscribe to AWS Shield Advanced. Add the EC2 instances as resources to protect.
- C. Create an AWS WAF web ACL that includes a rate-based rule. Associate the web ACL with the accelerator.
- D. Create an AWS WAF web ACL that includes a rate-based rule. Associate the web ACL with the EC2 instances.

**Correct Answer: A**

*Community vote distribution*

A (97%)

  **WherecanIstart** Highly Voted 1 year, 7 months ago

**Selected Answer: A**  
DDoS attacks = AWS Shield Advance  
Shield Advance protects Global Accelerator, NLB, ALB, etc  
upvoted 13 times

  **Abrar2022** Highly Voted 1 year, 4 months ago

**Selected Answer: A**  
DDoS attacks = AWS Shield Advance  
resource as Global Acc  
upvoted 5 times

  **pentium75** Most Recent 10 months, 1 week ago



**Selected Answer: A**  
Global Accelerator is what is exposed to the Internet = where DDoS attacks could land = what must be protected by Shield Advanced  
upvoted 5 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**  
So, the correct option is:  
  
B. Subscribe to AWS Shield Advanced. Add the EC2 instances as resources to protect.

Here's why this option is the most appropriate:

A. While you can add the accelerator as a resource to protect with AWS Shield Advanced, it's generally more effective to protect the individual resources (in this case, the EC2 instances) because AWS Shield Advanced will automatically protect resources associated with Global Accelerator  
upvoted 1 times

  **awsgeek75** 9 months, 2 weeks ago  
Which EC2 instance? Global Accelerator works by providing anycast IP addresses for the underlying resource (our EC2 in this case) so every end user trying to reach the EC2 server HAS to go through the Global Accelerator which is why the Global Accelerator needs to be protected and not the EC2.  
upvoted 5 times


  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**  
DDoS attacks = AWS Shield Advanced  
upvoted 3 times

  **nileshlg** 1 year, 7 months ago

**Selected Answer: A**  
Answer is A  
<https://docs.aws.amazon.com/waf/latest/developerguide/ddos-event-mitigation-logic-gax.html>

upvoted 2 times

  **ktulu2602** 1 year, 7 months ago

**Selected Answer: A**

AWS Shield is a managed service that provides protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. AWS Shield Standard is automatically enabled to all AWS customers at no additional cost. AWS Shield Advanced is an optional paid service. AWS Shield Advanced provides additional protections against more sophisticated and larger attacks for your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Route 53.

upvoted 4 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: A**

aaaaa  
accelator can not be attached to shield

upvoted 2 times

  **[Removed]** 1 year, 7 months ago

bbbbbbbbbb  
upvoted 1 times

  **enzomv** 1 year, 7 months ago

Your origin servers can be Amazon Simple Storage Service (S3), Amazon EC2, Elastic Load Balancing, or a custom server outside of AWS. You can also enable AWS Shield Advanced directly on Elastic Load Balancing or Amazon EC2 in the following AWS Regions - Northern Virginia, Ohio, Oregon, Northern California, Montreal, São Paulo, Ireland, Frankfurt, London, Paris, Stockholm, Singapore, Tokyo, Sydney, Seoul, Mumbai, Milan, and Cape Town.

My answer is B

upvoted 2 times

  **pentium75** 10 months, 1 week ago

You CAN enable Shield Advanced directly on EC2. You CAN also expose EC2 instances directly to the Internet. But in this case, what is exposed to the Internet (= where DDoS attacks could land) is the Global Accelerator, not your EC2 instances.

upvoted 3 times

  **enzomv** 1 year, 7 months ago

<https://docs.aws.amazon.com/waf/latest/developerguide/ddos-event-mitigation-logic-gax.html>

Sorry I meant A

upvoted 3 times

  **ktulu2602** 1 year, 7 months ago

Yes it can:  
AWS Shield is a managed service that provides protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. AWS Shield Standard is automatically enabled to all AWS customers at no additional cost. AWS Shield Advanced is an optional paid service. AWS Shield Advanced provides additional protections against more sophisticated and larger attacks for your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Route 53.

upvoted 2 times

An ecommerce company needs to run a scheduled daily job to aggregate and filter sales records for analytics. The company stores the sales records in an Amazon S3 bucket. Each object can be up to 10 GB in size. Based on the number of sales events, the job can take up to an hour to complete. The CPU and memory usage of the job are constant and are known in advance.

A solutions architect needs to minimize the amount of operational effort that is needed for the job to run.

Which solution meets these requirements?

- A. Create an AWS Lambda function that has an Amazon EventBridge notification. Schedule the EventBridge event to run once a day.
- B. Create an AWS Lambda function. Create an Amazon API Gateway HTTP API, and integrate the API with the function. Create an Amazon EventBridge scheduled event that calls the API and invokes the function.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an AWS Fargate launch type. Create an Amazon EventBridge scheduled event that launches an ECS task on the cluster to run the job.
- D. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an Amazon EC2 launch type and an Auto Scaling group with at least one EC2 instance. Create an Amazon EventBridge scheduled event that launches an ECS task on the cluster to run the job.

**Correct Answer:** C

Community vote distribution

C (100%)

🗲️ 👤 **ktulu2602** Highly Voted 👍 1 year, 7 months ago

**Selected Answer: C**

The requirement is to run a daily scheduled job to aggregate and filter sales records for analytics in the most efficient way possible. Based on the requirement, we can eliminate option A and B since they use AWS Lambda which has a limit of 15 minutes of execution time, which may not be sufficient for a job that can take up to an hour to complete.

Between options C and D, option C is the better choice since it uses AWS Fargate which is a serverless compute engine for containers that eliminates the need to manage the underlying EC2 instances, making it a low operational effort solution. Additionally, Fargate also provides instant scale-up and scale-down capabilities to run the scheduled job as per the requirement.

Therefore, the correct answer is:

C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an AWS Fargate launch type. Create an Amazon EventBridge scheduled event that launches an ECS task on the cluster to run the job.

upvoted 27 times

🗲️ 👤 **imvb88** Highly Voted 👍 1 year, 6 months ago

**Selected Answer: C**

"1-hour job" -> A, B out since max duration for Lambda is 15 min

Between C and D, "minimize operational effort" means Fargate -> C

upvoted 5 times

🗲️ 👤 **awsgeek75** Most Recent ⌚ 10 months ago

**Selected Answer: C**

A&B are out due to Lambda 15 min limits

C is less operationally complex than D so C is the right answer. Fargate is managed ECS cluster whereas EC2 based ECS will require more config overhead.

upvoted 4 times

🗲️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an AWS Fargate launch type. Create an Amazon EventBridge scheduled event that launches an ECS task on the cluster to run the job

upvoted 2 times

🗲️ 👤 **TariqKipkemei** 1 year, 5 months ago


**Selected Answer: C**

The best option is C.

'The job can take up to an hour to complete' rules out lambda functions as they only execute up to 15 mins. Hence option A and B are out.

'The CPU and memory usage of the job are constant and are known in advance' rules out the need for autoscaling. Hence option D is out.

upvoted 4 times

  **klayytech** 1 year, 7 months ago

**Selected Answer: C**

The solution that meets the requirements with the least operational overhead is to create a **\*\*Regional AWS WAF web ACL with a rate-based rule\*\*** and associate the web ACL with the API Gateway stage. This solution will protect the application from HTTP flood attacks by monitoring incoming requests and blocking requests from IP addresses that exceed the predefined rate.

Amazon CloudFront distribution with Lambda@Edge in front of the API Gateway Regional API endpoint is also a good solution but it requires more operational overhead than the previous solution.

Using Amazon CloudWatch metrics to monitor the Count metric and alerting the security team when the predefined rate is reached is not a solution that can protect against HTTP flood attacks.

Creating an Amazon CloudFront distribution in front of the API Gateway Regional API endpoint with a maximum TTL of 24 hours is not a solution that can protect against HTTP flood attacks.

upvoted 1 times

  **klayytech** 1 year, 7 months ago

**Selected Answer: C**

The solution that meets these requirements is C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an AWS Fargate launch type. Create an Amazon EventBridge scheduled event that launches an ECS task on the cluster to run the job. This solution will minimize the amount of operational effort that is needed for the job to run.

AWS Lambda which has a limit of 15 minutes of execution time,

upvoted 2 times

A company needs to transfer 600 TB of data from its on-premises network-attached storage (NAS) system to the AWS Cloud. The data transfer must be complete within 2 weeks. The data is sensitive and must be encrypted in transit. The company’s internet connection can support an upload speed of 100 Mbps.

Which solution meets these requirements MOST cost-effectively?

- A. Use Amazon S3 multi-part upload functionality to transfer the files over HTTPS.
- B. Create a VPN connection between the on-premises NAS system and the nearest AWS Region. Transfer the data over the VPN connection.
- C. Use the AWS Snow Family console to order several AWS Snowball Edge Storage Optimized devices. Use the devices to transfer the data to Amazon S3.
- D. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.

**Correct Answer: C**

*Community vote distribution*

C (100%)

- shanwford**

Highly Voted

1 year ago

**Selected Answer: C**

With the existing data link the transfer takes ~ 600 days in the best case. Thus, (A) and (B) are not applicable. Solution (D) could meet the target with a transfer time of 6 days, but the lead time for the direct connect deployment can take weeks! Thus, (C) is the only valid solution.

upvoted 14 times
- Guru4Cloud**

Most Recent

8 months, 1 week ago

**Selected Answer: C**

Use the AWS Snow Family console to order several AWS Snowball Edge Storage Optimized devices. Use the devices to transfer the data to Amazon S3.

upvoted 2 times
- TariqKipkemei**

11 months, 2 weeks ago

**Selected Answer: C**

C is the best option considering the time and bandwidth limitations

upvoted 2 times
- pbpally**

12 months ago

**Selected Answer: C**

We need the admin in here to tell us how they plan on this being achieved over connection with such a slow connection lol. It's C, folks.

upvoted 3 times
- KAUS2**

1 year, 1 month ago

**Selected Answer: C**

Best option is to use multiple AWS Snowball Edge Storage Optimized devices. Option "C" is the correct one.

upvoted 2 times
- ktulu2602**

1 year, 1 month ago

**Selected Answer: C**

All others are limited by the bandwidth limit

upvoted 1 times
- ktulu2602**

1 year, 1 month ago

Or provisioning time in the D case

upvoted 2 times
- KZM**

1 year, 1 month ago

It is C. Snowball (from Snow Family).

upvoted 2 times
- cegama543**

1 year, 1 month ago

**Selected Answer: C**

C. Use the AWS Snow Family console to order several AWS Snowball Edge Storage Optimized devices. Use the devices to transfer the data to Amazon S3.

The best option is to use the AWS Snow Family console to order several AWS Snowball Edge Storage Optimized devices and use the devices to transfer the data to Amazon S3. Snowball Edge is a petabyte-scale data transfer device that can help transfer large amounts of data securely and quickly. Using Snowball Edge can be the most cost-effective solution for transferring large amounts of data over long distances and can help meet the requirement of transferring 600 TB of data within two weeks.

upvoted 4 times

A financial company hosts a web application on AWS. The application uses an Amazon API Gateway Regional API endpoint to give users the ability to retrieve current stock prices. The company’s security team has noticed an increase in the number of API requests. The security team is concerned that HTTP flood attacks might take the application offline.

A solutions architect must design a solution to protect the application from this type of attack.

Which solution meets these requirements with the LEAST operational overhead?

- A. Create an Amazon CloudFront distribution in front of the API Gateway Regional API endpoint with a maximum TTL of 24 hours.
- B. Create a Regional AWS WAF web ACL with a rate-based rule. Associate the web ACL with the API Gateway stage.
- C. Use Amazon CloudWatch metrics to monitor the Count metric and alert the security team when the predefined rate is reached.
- D. Create an Amazon CloudFront distribution with Lambda@Edge in front of the API Gateway Regional API endpoint. Create an AWS Lambda function to block requests from IP addresses that exceed the predefined rate.

Correct Answer: B

Community vote distribution

B (100%)

- Guru4Cloud

Highly Voted

8 months, 1 week ago

Selected Answer: B

Regional AWS WAF web ACL is a managed web application firewall that can be used to protect your API Gateway API from a variety of attacks, including HTTP flood attacks.  
Rate-based rule is a type of rule that can be used to limit the number of requests that can be made from a single IP address within a specified period of time.  
API Gateway stage is a logical grouping of API resources that can be used to control access to your API.

upvoted 10 times
- elearningtakai

Highly Voted

1 year, 1 month ago

Selected Answer: B

A rate-based rule in AWS WAF allows the security team to configure thresholds that trigger rate-based rules, which enable AWS WAF to track the rate of requests for a specified time period and then block them automatically when the threshold is exceeded. This provides the ability to prevent HTTP flood attacks with minimal operational overhead.

upvoted 6 times
- TariqKipkemei

Most Recent

11 months, 2 weeks ago

Selected Answer: B

Answer is B

upvoted 2 times
- maxicalypse

1 year ago

B os correct

upvoted 2 times
- kampatra

1 year, 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/waf/latest/developerguide/web-acl.html>

upvoted 2 times
- [Removed]

1 year, 1 month ago

Selected Answer: B

bbbbbbbbb

upvoted 4 times

A meteorological startup company has a custom web application to sell weather data to its users online. The company uses Amazon DynamoDB to store its data and wants to build a new service that sends an alert to the managers of four internal teams every time a new weather event is recorded. The company does not want this new service to affect the performance of the current application.

What should a solutions architect do to meet these requirements with the LEAST amount of operational overhead?

- A. Use DynamoDB transactions to write new event data to the table. Configure the transactions to notify internal teams.
- B. Have the current application publish a message to four Amazon Simple Notification Service (Amazon SNS) topics. Have each team subscribe to one topic.
- C. Enable Amazon DynamoDB Streams on the table. Use triggers to write to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe.
- D. Add a custom attribute to each record to flag new items. Write a cron job that scans the table every minute for items that are new and notifies an Amazon Simple Queue Service (Amazon SQS) queue to which the teams can subscribe.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Buruguduystunstugudunstuy** Highly Voted 1 year, 1 month ago

**Selected Answer: C**

The best solution to meet these requirements with the least amount of operational overhead is to enable Amazon DynamoDB Streams on the table and use triggers to write to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe. This solution requires minimal configuration and infrastructure setup, and Amazon DynamoDB Streams provide a low-latency way to capture changes to the DynamoDB table. The triggers automatically capture the changes and publish them to the SNS topic, which notifies the internal teams.

upvoted 15 times

  **Buruguduystunstugudunstuy** 1 year, 1 month ago

Answer A is not a suitable solution because it requires additional configuration to notify the internal teams, and it could add operational overhead to the application.

Answer B is not the best solution because it requires changes to the current application, which may affect its performance, and it creates additional work for the teams to subscribe to multiple topics.

Answer D is not a good solution because it requires a cron job to scan the table every minute, which adds additional operational overhead to the system.

Therefore, the correct answer is C. Enable Amazon DynamoDB Streams on the table. Use triggers to write to a single Amazon SNS topic to which the teams can subscribe.

upvoted 6 times

  **Manimgh** Most Recent 2 weeks, 4 days ago

**Selected Answer: C**

Sends an alert = SNS



upvoted 1 times

  **Guru4Cloud** 8 months, 1 week ago

**Selected Answer: C**

Enable Amazon DynamoDB Streams on the table. Use triggers to write to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe

upvoted 4 times

  **james2033** 9 months, 2 weeks ago

**Selected Answer: C**

Question keyword: "sends an alert", a new weather event is recorded". Answer keyword C "Amazon DynamoDB Streams on the table", "Amazon Simple Notification Service" (Amazon SNS). Choose C. Easy question.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html>



<https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/>



upvoted 4 times



  **TariqKipkemei** 11 months, 2 weeks ago



**Selected Answer: C**

Best answer is C  
upvoted 3 times

  **TariqKipkemei** 6 months, 1 week ago  
DynamoDB Streams captures a time-ordered sequence of item-level modifications in any DynamoDB table and stores this information in a log for up to 24 hours. This capture activity can also invoke triggers to write the event to a single Amazon Simple Notification Service (Amazon SNS) topic to which the teams can subscribe to.  
upvoted 6 times



  **Hemanthgowda1932** 1 year, 1 month ago  
C is correct  
upvoted 2 times

  **Santosh43** 1 year, 1 month ago  
definitely C  
upvoted 2 times

  **Bezha** 1 year, 1 month ago  



Selected Answer: C

  
DynamoDB Streams  
upvoted 4 times

  **sitha** 1 year, 1 month ago  

Selected Answer: C

  
Answer : C  
upvoted 2 times

  **[Removed]** 1 year, 1 month ago  

Selected Answer: C

  
cccccccc  
upvoted 2 times

A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage.

The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand.

Which solution will meet these requirements?



- A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration.
- B. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zone. Deploy the database on an EC2 instance. Enable EC2 Auto Recovery.
- C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance with a read replica in a single Availability Zone. Promote the read replica to replace the primary DB instance if the primary DB instance fails.
- D. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones. Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

**Correct Answer: A**



*Community vote distribution*

A (93%)



3%



-   **pentium75** Highly Voted 10 months, 1 week ago



**Selected Answer: A**

B has app servers in a single AZ and a database on a single instance  
C has both DB replicas in a single AZ  
D does not work (EBS Multi-Attach requires EC2 instances in same AZ), and if it would work then the EBS volume would be an SPOF  
upvoted 7 times
-   **Guru4Cloud** Most Recent 1 year, 2 months ago



**Selected Answer: A**

Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration  
upvoted 2 times
-   **czyboi** 1 year, 2 months ago



Why is C incorrect ?  
upvoted 1 times
-   **Guru4Cloud** 1 year, 2 months ago

C is incorrect because the read replica also resides in a single AZ  
upvoted 4 times
-   **antropaws** 1 year, 5 months ago

**Selected Answer: A**

A most def.  
upvoted 3 times
-   **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration.  
upvoted 3 times
-   **Buruguduystunstugudunstuy** 1 year, 7 months ago



**Selected Answer: A**

The correct answer is A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon RDS DB instance in a Multi-AZ configuration.

To make an existing application highly available and resilient while avoiding any single points of failure and giving the application the ability to scale to meet user demand, the best solution would be to deploy the application servers using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones and use an Amazon RDS DB instance in a Multi-AZ configuration.

By using an Amazon RDS DB instance in a Multi-AZ configuration, the database is automatically replicated across multiple Availability Zones, ensuring that the database is highly available and can withstand the failure of a single Availability Zone. This provides fault tolerance and avoids any single points of failure.

upvoted 3 times

  **Thief** 1 year, 7 months ago

**Selected Answer: D**

Why not D?

upvoted 1 times

  **Guru4Cloud** 1 year, 2 months ago

D is incorrect because using Multi-Attach EBS adds complexity and doesn't provide automatic DB failover

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Multi-Attach does not work across Availability Zones.

upvoted 2 times

  **Buruguduystunstugudunstuy** 1 year, 7 months ago

Answer D, deploying the primary and secondary database servers on EC2 instances across multiple Availability Zones and using Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances, may provide high availability for the database but may introduce additional complexity, and management overhead, and potential performance issues.

upvoted 2 times

  **WheracanIstart** 1 year, 7 months ago

**Selected Answer: A**

Highly available = Multi-AZ approach

upvoted 3 times

  **nileshlg** 1 year, 7 months ago

**Selected Answer: A**

Answers is A

upvoted 1 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: A**

Option A is the correct solution. Deploying the application servers in an Auto Scaling group across multiple Availability Zones (AZs) ensures high availability and fault tolerance. An Auto Scaling group allows the application to scale horizontally to meet user demand. Using Amazon RDS DB instance in a Multi-AZ configuration ensures that the database is automatically replicated to a standby instance in a different AZ. This provides database redundancy and avoids any single point of failure.

upvoted 2 times

  **quentin17** 1 year, 7 months ago

**Selected Answer: C**



Highly available

upvoted 1 times

  **pentium75** 10 months, 1 week ago

No because instance and read replica "in a single Availability Zone"

upvoted 2 times

  **KAUS2** 1 year, 7 months ago

**Selected Answer: A**

Yes , agree with A

upvoted 2 times

  **cegama543** 1 year, 7 months ago

**Selected Answer: A**

agree with that

upvoted 2 times

A company needs to ingest and handle large amounts of streaming data that its application generates. The application runs on Amazon EC2 instances and sends data to Amazon Kinesis Data Streams, which is configured with default settings. Every other day, the application consumes the data and writes the data to an Amazon S3 bucket for business intelligence (BI) processing. The company observes that Amazon S3 is not receiving all the data that the application sends to Kinesis Data Streams.

What should a solutions architect do to resolve this issue?

- A. Update the Kinesis Data Streams default settings by modifying the data retention period.
- B. Update the application to use the Kinesis Producer Library (KPL) to send the data to Kinesis Data Streams.
- C. Update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams.
- D. Turn on S3 Versioning within the S3 bucket to preserve every version of every object that is ingested in the S3 bucket.

**Correct Answer: A**

*Community vote distribution*

A (70%)

C (27%)

  **WherecanIstart** Highly Voted 2 years, 1 month ago

**Selected Answer: A**

"A Kinesis data stream stores records from 24 hours by default, up to 8760 hours (365 days)."  
<https://docs.aws.amazon.com/streams/latest/dev/kinesis-extended-retention.html>

The question mentioned Kinesis data stream default settings and "every other day". After 24hrs, the data isn't in the Data stream if the default settings is not modified to store data more than 24hrs.

upvoted 33 times

  **babayomi** 6 months, 1 week ago

Thank you for the link

upvoted 2 times

  **cegama543** Highly Voted 2 years, 1 month ago

**Selected Answer: C**

C. Update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams.

The best option is to update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams. Kinesis Data Streams scales horizontally by increasing or decreasing the number of shards, which controls the throughput capacity of the stream. By increasing the number of shards, the application will be able to send more data to Kinesis Data Streams, which can help ensure that S3 receives all the data.

upvoted 17 times

  **Buruguduystunstugudunstuy** 2 years, 1 month ago

Answer C:

C. Update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams.

- Answer C updates the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams. By increasing the number of shards, the data is distributed across multiple shards, which allows for increased throughput and ensures that all data is ingested and processed by Kinesis Data Streams.

- Monitoring the Kinesis Data Streams and adjusting the number of shards as needed to handle changes in data throughput can ensure that the application can handle large amounts of streaming data.

upvoted 2 times

  **Buruguduystunstugudunstuy** 2 years, 1 month ago

@cegama543, my apologies. Moderator if you can disapprove of the post above? I made a mistake. It is supposed to be intended on the post that I submitted.

Thanks.

upvoted 2 times

  **CapJackSparrow** 2 years, 1 month ago

lets say you had infinity shards... if the retention period is 24 hours and you get the data every 48 hours, you will lose 24 hours of data no matter the amount of shards no?

upvoted 15 times

  **enzomv** 2 years, 1 month ago

Amazon Kinesis Data Streams supports changes to the data record retention period of your data stream. A Kinesis data stream is an ordered sequence of data records meant to be written to and read from in real time. Data records are therefore stored in shards in your stream

temporarily. The time period from when a record is added to when it is no longer accessible is called the retention period. A Kinesis data stream stores records from 24 hours by default, up to 8760 hours (365 days).

upvoted 6 times



  **Rcosmos** Most Recent 3 months, 1 week ago

**Selected Answer: U**

C. Atualizar o número de fragmentos do Kinesis:

Aumentar o número de fragmentos melhora a capacidade de processamento e a taxa de transferência do stream. No entanto, este não é o problema aqui, pois a perda de dados ocorre devido ao período de retenção insuficiente, e não devido a limitações de capacidade.

upvoted 1 times

  **abriggy** 9 months, 2 weeks ago

**Selected Answer: C**

Answer is C

Issue with A) Update the Kinesis Data Streams default settings by modifying the data retention period. is below

Limitation: Modifying the data retention period affects how long data is kept in the stream, but it does not address the issue of the stream's capacity to ingest data. If the stream is unable to handle the incoming data volume, extending the retention period will not resolve the data loss issue.

upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

Every other day, = 48 hours

Default settings = 24 hours

B: Development library so won't help

C: More shards may retain more data but they will have same limitation of 24 hours retention

D: Irrelevant

A: Increase the default limit from 24 hours to 48 hours

upvoted 6 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: A**

"Default settings" = 24 hour retention

upvoted 5 times

  **Murtadhaceit** 1 year, 4 months ago

**Selected Answer: A**

KDS has two modes:

1. Provisioned Mode: Answer C would be correct if KDS runs in this mode. We need to increase the number of shards.

2. On-Demand: Scales automatically, which means it doesn't need to adjust the number of shards based on observed throughput.

And since the question does not mention which type, I would go with On-demand. Therefore, A is the correct answer.

upvoted 3 times

  **TariqKipkemei** 1 year, 6 months ago


**Selected Answer: A**

Data records are stored in shards in a kinesis data stream temporarily. The time period from when a record is added, to when it is no longer accessible is called the retention period. This time period is 24 hours by default, but could be adjusted to 365 days.

Kinesis Data Streams automatically scales the number of shards in response to changes in data volume and traffic, so this rules out option C.

<https://docs.aws.amazon.com/streams/latest/dev/service-sizes-and-limits.html#:~:text=the%20number%20of-,shards,-in%20response%20to>

upvoted 2 times

  **Ramdi1** 1 year, 7 months ago

**Selected Answer: A**

I have only voted A because it mentions the default setting in Kinesis, if it did not mention that then I would look to increase the Shards. By default it is 24 hours and can go to 365 days. I think the question should be rephrased slightly. I had trouble deciding between A & C. Also apparently the most voted answer is the correct answer as per some advice I was given.

upvoted 3 times

  **BrijMohan08** 1 year, 8 months ago

**Selected Answer: A**

Default retention is 24 hrs, but the data read is every other day, so the S3 will never receive the data, Change the default retention period to 48 hours.

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

By default, a Kinesis data stream is created with one shard. If the data throughput to the stream is higher than the capacity of the single shard, the data stream may not be able to handle all the incoming data, and some data may be lost.

Therefore, to handle the high volume of data that the application sends to Kinesis Data Streams, the number of Kinesis shards should be increased to handle the required throughput.

Kinesis Data Streams shards are the basic units of scalability and availability. Each shard can process up to 1,000 records per second with a maximum of 1 MB of data per second. If the application is sending more data to Kinesis Data Streams than the shards can handle, then some of the data will be dropped.

upvoted 1 times

  **Guru4Cloud** 1 year, 8 months ago

If you have doubts, Please read about Kinesis Data Streams shards.

Ans: A is not the correct answer here

upvoted 1 times

  **Amycert** 1 year, 8 months ago

**Selected Answer: A**

the default retention period is 24 hours "The default retention period of 24 hours covers scenarios where intermittent lags in processing require catch-up with the real-time data. "

so we should increment this

upvoted 2 times

  **hsinchang** 1 year, 9 months ago

**Selected Answer: A**

As "Default settings" is mentioned here, I vote for A.

upvoted 2 times

  **jaydesai8** 1 year, 9 months ago



**Selected Answer: A**

keyword here is - default settings and every other day and since "A Kinesis data stream stores records from 24 hours by default, up to 8760 hours (365 days)."

<https://docs.aws.amazon.com/streams/latest/dev/kinesis-extended-retention.html>

Will go with A

upvoted 2 times

  **jayce5** 1 year, 11 months ago

**Selected Answer: A**

C is wrong because even if you update the number of Kinesis shards, you still need to change the default data retention period first. Otherwise, you would lose data after 24 hours.

upvoted 3 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: C**

A is unrelated to the issue. The correct answer is C.

upvoted 1 times

  **omoakin** 1 year, 11 months ago

Correct Ans. is B

upvoted 1 times

A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task. The developer already has an IAM user with valid IAM credentials required for Amazon S3.

What should a solutions architect do to grant the permissions?

- A. Add required IAM permissions in the resource policy of the Lambda function.
- B. Create a signed request using the existing IAM credentials in the Lambda function.
- C. Create a new IAM user and use the existing IAM credentials in the Lambda function.
- D. Create an IAM execution role with the required permissions and attach the IAM role to the Lambda function.

Correct Answer: D

Community vote distribution

D (100%)

- Buruguduystunstugudunstuy

Highly Voted

2 years, 1 month ago

To grant the necessary permissions to an AWS Lambda function to upload files to Amazon S3, a solutions architect should create an IAM execution role with the required permissions and attach the IAM role to the Lambda function. This approach follows the principle of least privilege and ensures that the Lambda function can only access the resources it needs to perform its specific task.

Therefore, the correct answer is D. Create an IAM execution role with the required permissions and attach the IAM role to the Lambda function.

upvoted 7 times
- AWSSURI

8 months, 1 week ago

Oh you're here

upvoted 3 times
- Guru4Cloud

Most Recent

1 year, 8 months ago

Selected Answer: D

Create Lambda execution role and attach existing S3 IAM role to the lambda function

upvoted 4 times
- Bilalglg93350

2 years, 1 month ago

D. Créez un rôle d'exécution IAM avec les autorisations requises et attachez le rôle IAM à la fonction Lambda.

L'architecte de solutions doit créer un rôle d'exécution IAM ayant les autorisations nécessaires pour accéder à Amazon S3 et effectuer les opérations requises (par exemple, charger des fichiers). Ensuite, le rôle doit être associé à la fonction Lambda, de sorte que la fonction puisse assumer ce rôle et avoir les autorisations nécessaires pour interagir avec Amazon S3.

upvoted 4 times
- nileshlg

2 years, 1 month ago

Selected Answer: D

Answer is D

upvoted 3 times
- kampatra

2 years, 1 month ago

Selected Answer: D

D - correct ans

upvoted 3 times
- sitha

2 years, 1 month ago

Selected Answer: D

Create Lambda execution role and attach existing S3 IAM role to the lambda function

upvoted 3 times
- ktulu2602

2 years, 1 month ago

Selected Answer: D

Definitely D

upvoted 3 times
- Nithin1119

2 years, 1 month ago

Selected Answer: D

ddddddd

upvoted 2 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: D**

ddddddd

upvoted 2 times

A company has deployed a serverless application that invokes an AWS Lambda function when new documents are uploaded to an Amazon S3 bucket. The application uses the Lambda function to process the documents. After a recent marketing campaign, the company noticed that the application did not process many of the documents.

What should a solutions architect do to improve the architecture of this application?

- A. Set the Lambda function's runtime timeout value to 15 minutes.
- B. Configure an S3 bucket replication policy. Stage the documents in the S3 bucket for later processing.
- C. Deploy an additional Lambda function. Load balance the processing of the documents across the two Lambda functions.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send the requests to the queue. Configure the queue as an event source for Lambda.

Correct Answer: D

Community vote distribution

D (100%)

  **Buruguduystunstugudunstuy** Highly Voted 2 years, 1 month ago

Selected Answer: D

To improve the architecture of this application, the best solution would be to use Amazon Simple Queue Service (Amazon SQS) to buffer the requests and decouple the S3 bucket from the Lambda function. This will ensure that the documents are not lost and can be processed at a later time if the Lambda function is not available.

This will ensure that the documents are not lost and can be processed at a later time if the Lambda function is not available. By using Amazon SQS, the architecture is decoupled and the Lambda function can process the documents in a scalable and fault-tolerant manner.

upvoted 5 times

  **Guru4Cloud** Most Recent 1 year, 8 months ago

Selected Answer: D

D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send the requests to the queue. Configure the queue as an event source for Lambd

upvoted 3 times

  **TariqKipkemei** 1 year, 11 months ago

Selected Answer: D

D is the best approach

upvoted 3 times

  **Russs99** 2 years, 1 month ago

Selected Answer: D

D is the correct answer



upvoted 3 times

  **Bilalglg93350** 2 years, 1 month ago

D. Créez une file d’attente Amazon Simple Queue Service (Amazon SQS). Envoyez les demandes à la file d’attente. Configurez la file d’attente en tant que source d’événement pour Lambda.

Cette solution permet de gérer efficacement les pics de charge et d'éviter la perte de documents en cas d'augmentation soudaine du trafic. Lorsque de nouveaux documents sont chargés dans le compartiment Amazon S3, les demandes sont envoyées à la file d'attente Amazon SQS, qui agit comme un tampon. La fonction Lambda est déclenchée en fonction des événements dans la file d'attente, ce qui permet un traitement équilibré et évite que l'application ne soit submergée par un grand nombre de documents simultanés.

upvoted 1 times

  **JA2018** 5 months, 2 weeks ago

From Google Translate:

Create an Amazon Simple Queue Service (Amazon SQS) queue. Send requests to the queue. Configure the queue as an event source for Lambda.

This solution makes it possible to effectively manage load peaks and avoid the loss of documents in the event of a sudden increase in traffic. When new documents are uploaded to the Amazon S3 bucket, requests are sent to the Amazon SQS queue, which acts as a buffer. The Lambda function is triggered based on events in the queue, which provides balanced processing and prevents the application from being overwhelmed by a large number of concurrent documents.

upvoted 1 times

  **Russs99** 2 years, 1 month ago

exactement. si je pouvais explique come cela en Francais aussi  
upvoted 1 times

  **WherecanIstart** 2 years, 1 month ago

**Selected Answer: D**

D is the correct answer.  
upvoted 1 times

  **kampatra** 2 years, 1 month ago

**Selected Answer: D**

D is correct  
upvoted 1 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: D**

D is correct  
upvoted 1 times

  **[Removed]** 2 years, 1 month ago

**Selected Answer: D**

ddddddddd  
upvoted 2 times

A solutions architect is designing the architecture for a software demonstration environment. The environment will run on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). The system will experience significant increases in traffic during working hours but is not required to operate on weekends.

Which combination of actions should the solutions architect take to ensure that the system can scale to meet demand? (Choose two.)

- A. Use AWS Auto Scaling to adjust the ALB capacity based on request rate.
- B. Use AWS Auto Scaling to scale the capacity of the VPC internet gateway.
- C. Launch the EC2 instances in multiple AWS Regions to distribute the load across Regions.
- D. Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization.
- E. Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. Revert to the default values at the start of the week.

Correct Answer: DE

Community vote distribution

DE (70%)

AD (15%)

Other

- pentium75

Highly Voted

1 year, 4 months ago

Selected Answer: DE

Not A - "AWS Auto Scaling" cannot adjust "ALB capacity" (<https://aws.amazon.com/autoscaling/faqs/>)  
Not B - VPC internet gateway has nothing to do with this  
Not C - Regions have nothing to do with scaling  
  
"The system will experience significant increases in traffic during working hours" -> addressed by D  
"But is not required to operate on weekends" -> addressed by E  
upvoted 15 times
- foha2012

1 year, 3 months ago

Good explanation!  
upvoted 2 times
- cd93

Highly Voted

1 year, 8 months ago

What does "ALB capacity" even means anyway? It should be "Target Group capacity" no?  
Answer should be DE, as D is a more comprehensive answer (and more practical in real life)  
upvoted 13 times
- bora4motion

Most Recent

4 days, 3 hours ago

Selected Answer: DE

who on this earth voted a?  
upvoted 1 times
- BigHammer

1 year, 8 months ago

AD  
E - the question doesn't ask about cost. Also, shutting it down during the weekend does nothing to improve scaling during the week. It doesn't address the requirements.  
upvoted 2 times
- JA2018

5 months, 2 weeks ago

Key in STEM:  
  
Nobody will use the system during weekend  
upvoted 1 times
- Guru4Cloud

1 year, 8 months ago

Selected Answer: DE

The solutions architect should take actions D and E:  
  
D) Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization. This will allow the Auto Scaling group to dynamically scale in and out based on demand.  
  
E) Use scheduled scaling to change the Auto Scaling group capacity to zero on weekends when traffic is expected to be low. This will minimize costs by terminating unused instances.  
upvoted 7 times

🗳️ 👤 **fuzzycr** 1 year, 9 months ago

**Selected Answer: AE**

Basado en los requerimientos la opción que se requiere para optimizar los costos de 0 operaciones en los fines de semana  
upvoted 2 times

🗳️ 👤 **jaydesai8** 1 year, 9 months ago

**Selected Answer: DE**

DE - This seems more close for the auto scaling -  
A - Its says auto scaling on ALB, but it should always be on EC2 instances and not ELB  
upvoted 7 times

🗳️ 👤 **XaviL** 1 year, 10 months ago

Hi guys, very simple  
\* A. because the question are asking about request rate!!!! This is a requirement!  
\* E. The weekend is not necessary to execute anything!

A&D. Is not possible, way you can put an ALB capacity based in cpu and in request rate???? You need to select one or another option (and this is for all questions here guys!)  
upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 10 months ago

**Selected Answer: AE**

ALBRequestCountPerTarget—Average Application Load Balancer request count per target.  
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html#target-tracking-choose-metrics>

It is possible to set to zero. "is not required to operate on weekends" means the instances are not required during the weekends.  
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/asg-capacity-limits.html>  
upvoted 3 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

A says to scale "ALB capacity", not number of EC2 instances. But "AWS Auto Scaling" cannot scale ALB capacity.  
upvoted 1 times

🗳️ 👤 **Uzi\_m** 1 year, 11 months ago

Option E is incorrect because the question specifically mentions an increase in traffic during working hours. Therefore, it is not advisable to schedule the instances for 24 hours using default settings throughout the entire week.  
E. Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. Revert to the default values at the start of the week.  
upvoted 1 times

🗳️ 👤 **omoakin** 1 year, 11 months ago

AD are the correct answs  
upvoted 3 times

🗳️ 👤 **TariqKipkemei** 1 year, 11 months ago

**Selected Answer: ADE**

Either one or two or all of these combinations will meet the need:  
Use AWS Auto Scaling to adjust the ALB capacity based on request rate.  
Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization.  
Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. Revert to the default values at the start of the week.  
upvoted 2 times

🗳️ 👤 **TariqKipkemei** 1 year, 6 months ago

Scheduled scaling was specifically designed to handle these kind of requirements.  
I therefore take out target scaling.  
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-scheduled-scaling.html#:~:text=RSS-,Scheduled%20scaling,-helps%20you%20to>  
upvoted 2 times

🗳️ 👤 **Joe94KR** 2 years ago

**Selected Answer: DE**

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html#target-tracking-choose-metrics>  
  
Based on docs, ASG can't track ALB's request rate, so the answer is D&E  
meanwhile ASG can track CPU rates.  
upvoted 5 times

🗳️ 👤 **[Removed]** 1 year, 11 months ago

The link shows:  
ALBRequestCountPerTarget—Average Application Load Balancer request count per target.  
upvoted 2 times

🗳️ 👤 **kraken21** 2 years, 1 month ago

**Selected Answer: DE**

Scaling should be at the ASG not ALB. So, not sure about "Use AWS Auto Scaling to adjust the ALB capacity based on request rate"  
upvoted 6 times

  **channn** 2 years, 1 month ago

**Selected Answer: AD**

A. Use AWS Auto Scaling to adjust the ALB capacity based on request rate: This will allow the system to scale up or down based on incoming traffic demand. The solutions architect should use AWS Auto Scaling to monitor the request rate and adjust the ALB capacity as needed.

D. Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization: This will allow the system to scale up or down based on the CPU utilization of the EC2 instances in the Auto Scaling group. The solutions architect should use a target tracking scaling policy to maintain a specific CPU utilization target and adjust the number of EC2 instances in the Auto Scaling group accordingly.

upvoted 9 times

  **pentium75** 1 year, 4 months ago

Auto scaling for ALB capacity?

upvoted 3 times

  **neosis91** 2 years, 1 month ago

**Selected Answer: AD**

A. Use a target tracking scaling policy to scale the Auto Scaling group based on instance CPU utilization. This approach allows the Auto Scaling group to automatically adjust the number of instances based on the specified metric, ensuring that the system can scale to meet demand during working hours.

D. Use scheduled scaling to change the Auto Scaling group minimum, maximum, and desired capacity to zero for weekends. Revert to the default values at the start of the week. This approach allows the Auto Scaling group to reduce the number of instances to zero during weekends when traffic is expected to be low. It will help the organization to save costs by not paying for instances that are not needed during weekends.

Therefore, options A and D are the correct answers. Options B and C are not relevant to the scenario, and option E is not a scalable solution as it would require manual intervention to adjust the group capacity every week.

upvoted 1 times

  **zooba72** 2 years, 1 month ago

**Selected Answer: DE**

This is why I don't believe A is correct use auto scaling to adjust the ALB .... D&E

upvoted 4 times

  **pentium75** 1 year, 4 months ago

Autoscaling can't scale the ALB

upvoted 2 times

A solutions architect is designing a two-tiered architecture that includes a public subnet and a database subnet. The web servers in the public subnet must be open to the internet on port 443. The Amazon RDS for MySQL DB instance in the database subnet must be accessible only to the web servers on port 3306.


Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)


- A. Create a network ACL for the public subnet. Add a rule to deny outbound traffic to 0.0.0.0/0 on port 3306.
- B. Create a security group for the DB instance. Add a rule to allow traffic from the public subnet CIDR block on port 3306.
- C. Create a security group for the web servers in the public subnet. Add a rule to allow traffic from 0.0.0.0/0 on port 443.
- D. Create a security group for the DB instance. Add a rule to allow traffic from the web servers' security group on port 3306.
- E. Create a security group for the DB instance. Add a rule to deny all traffic except traffic from the web servers' security group on port 3306.


**Correct Answer:** *CD*

*Community vote distribution*

CD (100%)

- 


 **Guru4Cloud**


Highly Voted 


 1 year, 2 months ago

Selected Answer: CD

Remember guys that SG is not used for Deny action, just Allow

upvoted 11 times
- 


 **datmd77**


Highly Voted 


 1 year, 6 months ago

Selected Answer: CD

Remember guys that SG is not used for Deny action, just Allow

upvoted 5 times
- 


 **Tendu**


Most Recent 

 2 months, 1 week ago

Selected Answer: CD

C and D

upvoted 1 times
- 


 **waldirlsantos** 6 months, 3 weeks ago


Selected Answer: CD

The following are the default rules for a security group that you create:

Allows no inbound traffic

Allows all outbound traffic

upvoted 3 times
- 


 **TariqKipkemei** 1 year ago


Selected Answer: CD

'must be accessible only to the web servers' is the key here.

Option B almost threw me off, but with this then all that exists in the public subnet would be able to access the DB security group.

Therefore C,D well applies the principle of least privilege.

upvoted 5 times
- 


 **Buruguduystunstugudunstuy** 1 year, 7 months ago


Selected Answer: CD

To meet the requirements of allowing access to the web servers in the public subnet on port 443 and the Amazon RDS for MySQL DB instance in the database subnet on port 3306, the best solution would be to create a security group for the web servers and another security group for the DB instance, and then define the appropriate inbound and outbound rules for each security group.

  1. Create a security group for the web servers in the public subnet. Add a rule to allow traffic from 0.0.0.0/0 on port 443.
  2. Create a security group for the DB instance. Add a rule to allow traffic from the web servers' security group on port 3306.



This will allow the web servers in the public subnet to receive traffic from the internet on port 443, and the Amazon RDS for MySQL DB instance in the database subnet to receive traffic only from the web servers on port 3306.



upvoted 3 times
- 



 **kampatra** 1 year, 7 months ago



Selected Answer: CD

CD - Correct ans.  
upvoted 3 times

  **Eden** 1 year, 7 months ago  
I choose CE  
upvoted 1 times



  **lili\_9** 1 year, 7 months ago  
CE support @sitha  
upvoted 1 times



  **sitha** 1 year, 7 months ago  
Answer: CE . The solution is to deny accessing DB from Internet and allow only access from webserver.  
upvoted 1 times



  **KAUS2** 1 year, 7 months ago  



Selected Answer: CD

  
C & D are the right choices. correct  
upvoted 2 times

  **KS2020** 1 year, 7 months ago  
why not CE?  
upvoted 2 times

  **kampatra** 1 year, 7 months ago  
By default Security Group deny all trafic and we need to configure to enable.  
upvoted 5 times

  **[Removed]** 1 year, 7 months ago  
Characteristics of security group rules  
  
You can specify allow rules, but not deny rules.  
[https://docs.aws.amazon.com/vpc/latest/userguide/VPC\\_SecurityGroups.html](https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html)  
upvoted 3 times

  **[Removed]** 1 year, 7 months ago  

Selected Answer: CD

  
cdcdcdcdcd  
upvoted 3 times

A company is implementing a shared storage solution for a gaming application that is hosted in the AWS Cloud. The company needs the ability to use Lustre clients to access data. The solution must be fully managed.

Which solution meets these requirements?

- A. Create an AWS DataSync task that shares the data as a mountable file system. Mount the file system to the application server.
- B. Create an AWS Storage Gateway file gateway. Create a file share that uses the required client protocol. Connect the application server to the file share.
- C. Create an Amazon Elastic File System (Amazon EFS) file system, and configure it to support Lustre. Attach the file system to the origin server. Connect the application server to the file system.
- D. Create an Amazon FSx for Lustre file system. Attach the file system to the origin server. Connect the application server to the file system.

Correct Answer: D

Community vote distribution

D (100%)

  **Buruguduystunstugudunstuy** Highly Voted 1 year, 1 month ago

Selected Answer: D

To meet the requirements of a shared storage solution for a gaming application that can be accessed using Lustre clients and is fully managed, the best solution would be to use Amazon FSx for Lustre.

Amazon FSx for Lustre is a fully managed file system that is optimized for compute-intensive workloads, such as high-performance computing, machine learning, and gaming. It provides a POSIX-compliant file system that can be accessed using Lustre clients and offers high performance, scalability, and data durability.

This solution provides a highly available, scalable, and fully managed shared storage solution that can be accessed using Lustre clients. Amazon FSx for Lustre is optimized for compute-intensive workloads and provides high performance and durability.

upvoted 6 times

  **Buruguduystunstugudunstuy** 1 year, 1 month ago

Answer A, creating an AWS DataSync task that shares the data as a mountable file system and mounting the file system to the application server, may not provide the required performance and scalability for a gaming application.

Answer B, creating an AWS Storage Gateway file gateway and connecting the application server to the file share, may not provide the required performance and scalability for a gaming application.

Answer C, creating an Amazon Elastic File System (Amazon EFS) file system and configuring it to support Lustre, may not provide the required performance and scalability for a gaming application and may require additional configuration and management overhead.


upvoted 3 times

  **Manimgh** Most Recent 1 week, 4 days ago

Selected Answer: D

Lustre = FSx

upvoted 1 times

  **Tendu** 2 months, 1 week ago

Selected Answer: D

D is correct answer.

upvoted 1 times

  **Guru4Cloud** 8 months, 1 week ago

Selected Answer: D

Lustre clients = Amazon FSx for Lustre file system

upvoted 4 times

  **TariqKipkemei** 11 months, 2 weeks ago

Selected Answer: D

Lustre clients = Amazon FSx for Lustre file system

upvoted 4 times

  **kampatra** 1 year, 1 month ago

Selected Answer: D

D - correct ans

upvoted 2 times

  **kprakashbehera** 1 year, 1 month ago

**Selected Answer: D**

FSx for Lustre  
DDDDDD

upvoted 2 times

  **KAUS2** 1 year, 1 month ago

**Selected Answer: D**

Amazon FSx for Lustre is the right answer

- Lustre is a type of parallel distributed file system, for large-scale computing, Machine Learning, High Performance Computing (HPC)
- Video Processing, Financial Modeling, Electronic Design Automatio

upvoted 2 times

  **cegama543** 1 year, 1 month ago

**Selected Answer: D**

Option D is the best solution because Amazon FSx for Lustre is a fully managed, high-performance file system that is designed to support compute-intensive workloads, such as those required by gaming applications. FSx for Lustre provides sub-millisecond access to petabyte-scale file systems, and supports Lustre clients natively. This means that the gaming application can access the shared data directly from the FSx for Lustre file system without the need for additional configuration or setup.

Additionally, FSx for Lustre is a fully managed service, meaning that AWS takes care of all maintenance, updates, and patches for the file system, which reduces the operational overhead required by the company.

upvoted 2 times

  **[Removed]** 1 year, 1 month ago

**Selected Answer: D**

ddddddddddddd  
upvoted 2 times

A company runs an application that receives data from thousands of geographically dispersed remote devices that use UDP. The application processes the data immediately and sends a message back to the device if necessary. No data is stored.

The company needs a solution that minimizes latency for the data transmission from the devices. The solution also must provide rapid failover to another AWS Region.

Which solution will meet these requirements?

- A. Configure an Amazon Route 53 failover routing policy. Create a Network Load Balancer (NLB) in each of the two Regions. Configure the NLB to invoke an AWS Lambda function to process the data.
- B. Use AWS Global Accelerator. Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the NLB. Process the data in Amazon ECS.
- C. Use AWS Global Accelerator. Create an Application Load Balancer (ALB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.
- D. Configure an Amazon Route 53 failover routing policy. Create an Application Load Balancer (ALB) in each of the two Regions. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.

**Correct Answer: B**

*Community vote distribution*

B (100%)

🗲️ 👤 **UnluckyDucky** Highly Voted 👍 1 year, 7 months ago

**Selected Answer: B**

Key words: geographically dispersed, UDP.

Geographically dispersed (related to UDP) - Global Accelerator - multiple entrances worldwide to the AWS network to provide better transfer rates.  
UDP - NLB (Network Load Balancer).

upvoted 15 times

🗲️ 👤 **Ruhi02** Highly Voted 👍 1 year, 7 months ago

Answer should be B.. there is typo mistake in B. Correct Answer is : Use AWS Global Accelerator. Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the NLB. Process the data in Amazon ECS.

upvoted 5 times

🗲️ 👤 **wizcloudifa** Most Recent 🕒 6 months, 3 weeks ago

**Selected Answer: B**

if its UDP it has to be Global Accelerator + NLB package, plus it has the provision for rapid failover as well, piece of cake.

upvoted 3 times

🗲️ 👤 **sandordini** 6 months, 3 weeks ago

**Selected Answer: B**

UDP: NLB + AWS Global Accelerator

upvoted 3 times

🗲️ 👤 **zinabu** 7 months ago

UDP/TCP=NLB

rapid failover= AWS global accelerator

upvoted 3 times

🗲️ 👤 **ferdzcruz** 10 months, 1 week ago



devices that use UDP = NLB

upvoted 2 times

🗲️ 👤 **ferdzcruz** 10 months, 1 week ago

minimizes latency = AWS Global Accelerator

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

This option meets the requirements:

Global Accelerator provides UDP support and minimizes latency using the AWS global network. Using NLBs allows the UDP traffic to be load balanced across Availability Zones. ECS Fargate provides rapid scaling and failover across Regions. NLB endpoints allow rapid failover if one Region goes down.  
upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago



**Selected Answer: B**

UDP = AWS Global Accelerator and Network Load Balancer  
upvoted 2 times

  **kraken21** 1 year, 7 months ago

**Selected Answer: B**

Global accelerator for multi region automatic failover. NLB for UDP.  
upvoted 3 times

  **MaxMa** 1 year, 7 months ago

why not A?  
upvoted 1 times

  **kraken21** 1 year, 7 months ago

NLBs do not support lambda target type. Tricky!!! <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html>  
upvoted 10 times

  **Buruguduystunstugudunstuy** 1 year, 7 months ago

**Selected Answer: B**

To meet the requirements of minimizing latency for data transmission from the devices and providing rapid failover to another AWS Region, the best solution would be to use AWS Global Accelerator in combination with a Network Load Balancer (NLB) and Amazon Elastic Container Service (Amazon ECS).

AWS Global Accelerator is a service that improves the availability and performance of applications by using static IP addresses (Anycast) to route traffic to optimal AWS endpoints. With Global Accelerator, you can direct traffic to multiple Regions and endpoints, and provide automatic failover to another AWS Region.  
upvoted 4 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: B**

bbbbbbbbb  
upvoted 2 times

A solutions architect must migrate a Windows Internet Information Services (IIS) web application to AWS. The application currently relies on a file share hosted in the user's on-premises network-attached storage (NAS). The solutions architect has proposed migrating the IIS web servers to Amazon EC2 instances in multiple Availability Zones that are connected to the storage solution, and configuring an Elastic Load Balancer attached to the instances.

Which replacement to the on-premises file share is MOST resilient and durable?

- A. Migrate the file share to Amazon RDS.
- B. Migrate the file share to AWS Storage Gateway.
- C. Migrate the file share to Amazon FSx for Windows File Server.
- D. Migrate the file share to Amazon Elastic File System (Amazon EFS).

**Correct Answer: C**

*Community vote distribution*

C (95%)

5%

- channn

Highly Voted

1 year, 7 months ago

Selected Answer: C

A) RDS is a database service  
B) Storage Gateway is a hybrid cloud storage service that connects on-premises applications to AWS storage services.  
D) provides shared file storage for Linux-based workloads, but it does not natively support Windows-based workloads.  
upvoted 9 times
- Buruguduystunstugudunstuy

Highly Voted

1 year, 7 months ago

Selected Answer: C

The most resilient and durable replacement for the on-premises file share in this scenario would be Amazon FSx for Windows File Server.  
  
Amazon FSx is a fully managed Windows file system service that is built on Windows Server and provides native support for the SMB protocol. It is designed to be highly available and durable, with built-in backup and restore capabilities. It is also fully integrated with AWS security services, providing encryption at rest and in transit, and it can be configured to meet compliance standards.  
upvoted 7 times
- Buruguduystunstugudunstuy

1 year, 7 months ago

Migrating the file share to Amazon RDS or AWS Storage Gateway is not appropriate as these services are designed for database workloads and block storage respectively, and do not provide native support for the SMB protocol.  
  
Migrating the file share to Amazon EFS (Linux ONLY) could be an option, but Amazon FSx for Windows File Server would be more appropriate in this case because it is specifically designed for Windows file shares and provides better performance for Windows applications.  
upvoted 6 times
- Manimgh

Most Recent

1 week, 4 days ago

Selected Answer: C

IIS = FSx  
upvoted 1 times
- com7

11 months, 1 week ago

Selected Answer: C

Windows Server to FSx For Windows  
upvoted 3 times
- Guru4Cloud

1 year, 2 months ago

Selected Answer: C

Windows client = Amazon FSx for Windows File Serve  
upvoted 2 times
- TariqKipkemei

1 year, 5 months ago

Selected Answer: C

Windows client = Amazon FSx for Windows File Server  
upvoted 3 times
- Grace83

1 year, 7 months ago

Obviously C is the correct answer - FSx for Windows - Windows

upvoted 5 times

  **UnluckyDucky** 1 year, 7 months ago

**Selected Answer: C**

FSx for Windows - Windows.  
EFS - Linux.

upvoted 6 times

  **mwwt2022** 10 months ago

good summary

upvoted 3 times

  **elearningtakai** 1 year, 7 months ago

**Selected Answer: D**

Amazon EFS is a scalable and fully-managed file storage service that is designed to provide high availability and durability. It can be accessed by multiple EC2 instances across multiple Availability Zones simultaneously. Additionally, it offers automatic and instantaneous data replication across different availability zones within a region, which makes it resilient to failures.

upvoted 2 times

  **asoli** 1 year, 7 months ago

EFS is a wrong choice because it can only work with Linux instances. That application has a Windows web server , so its OS is Windows and EFS cannot connect to it

upvoted 5 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

Amazon FSx


upvoted 1 times

  **sitha** 1 year, 7 months ago

Amazon FSx makes it easy and cost effective to launch, run, and scale feature-rich, high-performance file systems in the cloud.

Answer : C



upvoted 2 times

  **KAUS2** 1 year, 7 months ago

**Selected Answer: C**

FSx for Windows is a fully managed Windows file system share drive . Hence C is the correct answer.

upvoted 3 times

  **Ruhi02** 1 year, 7 months ago

FSx for Windows is ideal in this case. So answer is C.

upvoted 2 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

cccccccc

upvoted 1 times

A company is deploying a new application on Amazon EC2 instances. The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes. The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest.

Which solution will meet this requirement?

- A. Create an IAM role that specifies EBS encryption. Attach the role to the EC2 instances.
- B. Create the EBS volumes as encrypted volumes. Attach the EBS volumes to the EC2 instances.
- C. Create an EC2 instance tag that has a key of Encrypt and a value of True. Tag all instances that require encryption at the EBS level.
- D. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBS encryption in the account. Ensure that the key policy is active.

Correct Answer: B

Community vote distribution

B (100%)

  **Buruguduystunstugudunstuy** Highly Voted 1 year, 1 month ago

Selected Answer: B

The solution that will meet the requirement of ensuring that all data that is written to the EBS volumes is encrypted at rest is B. Create the EBS volumes as encrypted volumes and attach the encrypted EBS volumes to the EC2 instances.

When you create an EBS volume, you can specify whether to encrypt the volume. If you choose to encrypt the volume, all data written to the volume is automatically encrypted at rest using AWS-managed keys. You can also use customer-managed keys (CMKs) stored in AWS KMS to encrypt and protect your EBS volumes. You can create encrypted EBS volumes and attach them to EC2 instances to ensure that all data written to the volumes is encrypted at rest.

Answer A is incorrect because attaching an IAM role to the EC2 instances does not automatically encrypt the EBS volumes.

Answer C is incorrect because adding an EC2 instance tag does not ensure that the EBS volumes are encrypted.  
upvoted 12 times

  **Kds53829** Most Recent 6 months, 1 week ago

B is the answer  
upvoted 2 times

  **Guru4Cloud** 8 months, 1 week ago

Selected Answer: B

B. Create the EBS volumes as encrypted volumes. Attach the EBS volumes to the EC2 instances.  
upvoted 2 times

  **TariqKipkemei** 11 months, 2 weeks ago

Selected Answer: B

Windows client = Amazon FSx for Windows File Server  
upvoted 2 times


  **TariqKipkemei** 6 months, 1 week ago

ignore this, mind stuck on last question hhhhhh.  
Just create the EBS volumes as encrypted volumes then attach the EBS volumes to the EC2 instances.  
upvoted 5 times

  **elearningtakai** 1 year, 1 month ago

Selected Answer: B

The other options either do not meet the requirement of encrypting data at rest (A and C) or do so in a more complex or less efficient manner (D).  
upvoted 2 times

  **Bofi** 1 year, 1 month ago

Why not D, EBS encryption require the use of KMS key  
upvoted 1 times

  **Buruguduystunstugudunstuy** 1 year, 1 month ago



Answer D is incorrect because creating a KMS key policy that enforces EBS encryption does not automatically encrypt EBS volumes. You need to create encrypted EBS volumes and attach them to EC2 instances to ensure that all data written to the volumes are encrypted at rest.  
upvoted 10 times

  **WherecanIstart** 1 year, 1 month ago

**Selected Answer: B**

Create encrypted EBS volumes and attach encrypted EBS volumes to EC2 instances..

upvoted 3 times

  **sitha** 1 year, 1 month ago

Use Amazon EBS encryption as an encryption solution for your EBS resources associated with your EC2 instances. Select KMS Keys either default or custom

upvoted 2 times

  **Ruhi02** 1 year, 1 month ago

Answer B. You can enable encryption for EBS volumes while creating them.

upvoted 2 times

  **[Removed]** 1 year, 1 month ago

**Selected Answer: B**

bbbbbbbbb

upvoted 2 times

A company has a web application with sporadic usage patterns. There is heavy usage at the beginning of each month, moderate usage at the start of each week, and unpredictable usage during the week. The application consists of a web server and a MySQL database server running inside the data center. The company would like to move the application to the AWS Cloud, and needs to select a cost-effective database platform that will not require database modifications.

Which solution will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon RDS for MySQL
- C. MySQL-compatible Amazon Aurora Serverless
- D. MySQL deployed on Amazon EC2 in an Auto Scaling group

**Correct Answer: C**

*Community vote distribution*

C (92%)

8%

  **channn** Highly Voted 1 year, 7 months ago

**Selected Answer: C**

C: Aurora Serverless is a MySQL-compatible relational database engine that automatically scales compute and memory resources based on application usage. no upfront costs or commitments required.

A: DynamoDB is a NoSQL

B: Fixed cost on RDS class

D: More operation requires

upvoted 12 times

  **Buruguduystunstugudunstuy** Highly Voted 1 year, 7 months ago

**Selected Answer: C**

Answer C, MySQL-compatible Amazon Aurora Serverless, would be the best solution to meet the company's requirements.

Aurora Serverless can be a cost-effective option for databases with sporadic or unpredictable usage patterns since it automatically scales up or down based on the current workload. Additionally, Aurora Serverless is compatible with MySQL, so it does not require any modifications to the application's database code.

upvoted 6 times

  **TariqKipkemei** Most Recent 1 year ago

**Selected Answer: C**

The is a huge demand for auto-scaling which Amazon RDS cannot do. This contributes to the cost savings as Aurora serverless would scale done in low peak times, this contributes to low costs.

upvoted 4 times

  **JKevin778** 1 year, 1 month ago

**Selected Answer: B**

RDS is cheaper than Aurora.

upvoted 1 times

  **pentium75** 10 months, 1 week ago

RDS is cheaper than Aurora if you have a fixed instance size, but NOT if you have "unpredictable" usage patterns, then Aurora Serverless (!) is cheaper.

upvoted 5 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

Answer C, MySQL-compatible Amazon Aurora Serverless, would be the best solution to meet the company's requirements.

upvoted 2 times

  **MrAWSAssociate** 1 year, 4 months ago

**Selected Answer: C**

Since we have sporadic & unpredictable usage for DB, Aurora Serverless would be fit more cost-efficient for this case scenario than RDS MySQL. <https://www.techtarget.com/searchcloudcomputing/answer/When-should-I-use-Amazon-RDS-vs-Aurora-Serverless>

upvoted 1 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: C**

C for sure.  
upvoted 3 times

  **klayytech** 1 year, 7 months ago



**Selected Answer: B**

Amazon RDS for MySQL is a cost-effective database platform that will not require database modifications. It makes it easier to set up, operate, and scale MySQL deployments in the cloud. With Amazon RDS, you can deploy scalable MySQL servers in minutes with cost-efficient and resizable hardware capacity<sup>2</sup>.



Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. DynamoDB is a good choice for applications that require low-latency data access<sup>1</sup>.

MySQL-compatible Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora (MySQL-compatible edition), where the database will automatically start up, shut down, and scale capacity up or down based on your application's needs<sup>3</sup>.

So, Amazon RDS for MySQL is the best option for your requirements.  
upvoted 2 times

  **klayytech** 1 year, 7 months ago  
sorry i will change to C , because

Amazon RDS for MySQL is a fully-managed relational database service that makes it easy to set up, operate, and scale MySQL deployments in the cloud. Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora (MySQL-compatible edition), where the database will automatically start up, shut down, and scale capacity up or down based on your application’s needs. It is a simple, cost-effective option for infrequent, intermittent, or unpredictable workloads.  
upvoted 3 times

  **boxu03** 1 year, 7 months ago

**Selected Answer: C**

Amazon Aurora Serverless : a simple, cost-effective option for infrequent, intermittent, or unpredictable workloads  
upvoted 4 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: C**

cccccccccccccccccccc  
upvoted 2 times

An image-hosting company stores its objects in Amazon S3 buckets. The company wants to avoid accidental exposure of the objects in the S3 buckets to the public. All S3 objects in the entire AWS account need to remain private.

Which solution will meet these requirements?

- A. Use Amazon GuardDuty to monitor S3 bucket policies. Create an automatic remediation action rule that uses an AWS Lambda function to remediate any change that makes the objects public.
- B. Use AWS Trusted Advisor to find publicly accessible S3 buckets. Configure email notifications in Trusted Advisor when a change is detected. Manually change the S3 bucket policy if it allows public access.
- C. Use AWS Resource Access Manager to find publicly accessible S3 buckets. Use Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function when a change is detected. Deploy a Lambda function that programmatically remediates the change.
- D. Use the S3 Block Public Access feature on the account level. Use AWS Organizations to create a service control policy (SCP) that prevents IAM users from changing the setting. Apply the SCP to the account.

Correct Answer: D

Community vote distribution

D (96%)

4%

- Ruhi02

Highly Voted

2 years, 1 month ago

Answer is D ladies and gentlemen. While guard duty helps to monitor s3 for potential threats its a reactive action. We should always be proactive and not reactive in our solutions so D, block public access to avoid any possibility of the info becoming publicly accessible

upvoted 23 times
- Buruguduystunstugudunstuy

Highly Voted

2 years, 1 month ago

Selected Answer: D

Answer D is the correct solution that meets the requirements. The S3 Block Public Access feature allows you to restrict public access to S3 buckets and objects within the account. You can enable this feature at the account level to prevent any S3 bucket from being made public, regardless of the bucket policy settings. AWS Organizations can be used to apply a Service Control Policy (SCP) to the account to prevent IAM users from changing this setting, ensuring that all S3 objects remain private. This is a straightforward and effective solution that requires minimal operational overhead.

upvoted 10 times
- noircesar25

Most Recent

1 year, 2 months ago

its 1 aws account, how could D be the answer?

upvoted 1 times
- JA2018

5 months, 2 weeks ago

make no difference, 1 AWS account, 10 AWS accounts or > 100 AWS accounts, SCP enforces the same guard rails

upvoted 2 times
- TariqKipkemei

1 year, 6 months ago

Selected Answer: D

Use the S3 Block Public Access feature on the account level. Use AWS Organizations to create a service control policy (SCP) that prevents IAM users from changing the setting. Apply the SCP to the account

upvoted 2 times
- Guru4Cloud

1 year, 8 months ago

Selected Answer: D

Use the S3 Block Public Access feature on the account level. Use AWS Organizations to create a service control policy (SCP) that prevents IAM users from changing the setting. Apply the SCP to the account

upvoted 2 times
- MrAWSAssociate

1 year, 10 months ago

Selected Answer: A

A is correct!

upvoted 1 times
- pentium75

1 year, 4 months ago

No, first it would not remove any existing public access (only detect changes), second it would just detect and then remediate, but in the meantime someone could access the objects. It's clearly D.

upvoted 3 times
- Yadav\_Sanjay

1 year, 11 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-control-block-public-access.html>



upvoted 4 times

  **elearningtakai** 2 years, 1 month ago

**Selected Answer: D**

This is the most effective solution to meet the requirements.

upvoted 3 times

  **Bofi** 2 years, 1 month ago

**Selected Answer: D**

Option D provided real solution by using bucket policy to restrict public access. Other options were focus on detection which wasn't what was been asked

upvoted 3 times

An ecommerce company is experiencing an increase in user traffic. The company’s store is deployed on Amazon EC2 instances as a two-tier web application consisting of a web tier and a separate database tier. As traffic increases, the company notices that the architecture is causing significant delays in sending timely marketing and order confirmation email to users. The company wants to reduce the time it spends resolving complex email delivery issues and minimize operational overhead.

What should a solutions architect do to meet these requirements?

- A. Create a separate application tier using EC2 instances dedicated to email processing.
- B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES).
- C. Configure the web instance to send email through Amazon Simple Notification Service (Amazon SNS).
- D. Create a separate application tier using EC2 instances dedicated to email processing. Place the instances in an Auto Scaling group.

**Correct Answer: B**

*Community vote distribution*

B (100%)

- elearningtakai**

Highly Voted

1 year, 7 months ago

Selected Answer: B

Amazon SES is a cost-effective and scalable email service that enables businesses to send and receive email using their own email addresses and domains. Configuring the web instance to send email through Amazon SES is a simple and effective solution that can reduce the time spent resolving complex email delivery issues and minimize operational overhead.

upvoted 11 times
- Buruguduystunstugudunstuy**

Highly Voted

1 year, 7 months ago

Selected Answer: B

The best option for addressing the company's needs of minimizing operational overhead and reducing time spent resolving email delivery issues is to use Amazon Simple Email Service (Amazon SES).

Answer A of creating a separate application tier for email processing may add additional complexity to the architecture and require more operational overhead.

Answer C of using Amazon Simple Notification Service (Amazon SNS) is not an appropriate solution for sending marketing and order confirmation emails since Amazon SNS is a messaging service that is designed to send messages to subscribed endpoints or clients.

Answer D of creating a separate application tier using EC2 instances dedicated to email processing placed in an Auto Scaling group is a more complex solution than necessary and may result in additional operational overhead.

upvoted 6 times
- waldirlsantos**

Most Recent

6 months, 3 weeks ago

Selected Answer: B

B meet these requirements

upvoted 2 times
- TariqKipkemei**

1 year ago

Selected Answer: B

Amazon Simple Email Service (Amazon SES) lets you reach customers confidently without an on-premises Simple Mail Transfer Protocol (SMTP) email server using the Amazon SES API or SMTP interface.

upvoted 3 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: B

B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES)

upvoted 2 times
- nileshlg**

1 year, 7 months ago

Answer is B

upvoted 2 times
- Ruhi02**

1 year, 7 months ago

Answer B.. SES is meant for sending high volume e-mail efficiently and securely.

SNS is meant as a channel publisher/subscriber service

upvoted 5 times

  **[Removed]** 1 year, 7 months ago

**Selected Answer: B**

bbbbbbbbb

upvoted 2 times

A company has a business system that generates hundreds of reports each day. The business system saves the reports to a network share in CSV format. The company needs to store this data in the AWS Cloud in near-real time for analysis.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use AWS DataSync to transfer the files to Amazon S3. Create a scheduled task that runs at the end of each day.
- B. Create an Amazon S3 File Gateway. Update the business system to use a new network share from the S3 File Gateway.
- C. Use AWS DataSync to transfer the files to Amazon S3. Create an application that uses the DataSync API in the automation workflow.
- D. Deploy an AWS Transfer for SFTP endpoint. Create a script that checks for new files on the network share and uploads the new files by using SFTP.

**Correct Answer: B**

Community vote distribution

B (86%)

12%

- TariqKipkemei**

Highly Voted

1 year ago

Selected Answer: B

Both Amazon S3 File Gateway and AWS DataSync are suitable for this scenario. But there is a requirement for 'LEAST administrative overhead'. Option C involves the creation of an entirely new application to consume the DataSync API, this rules out this option.

upvoted 14 times
- channn**

Highly Voted

1 year, 7 months ago

Selected Answer: B

Key words:

  1. near-real-time (A is out)
  2. LEAST administrative (C n D is out)

upvoted 10 times
- zdi561**

Most Recent

3 months, 2 weeks ago

Selected Answer: D

Firstly the requirement is not about hybrid cloud e.g. use the data migrated to aws from on-premise, so B is gone. Secondly it is near real time, A is gone. Thirdly less admistratively, C is one since an application is a headache. D is the best , it only involves a script. SFTP is good for light file transfer. Here there are only hundreds of files in a day.

upvoted 1 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: C

This option has the least administrative overhead because:

Using DataSync avoids having to rewrite the business system to use a new file gateway or SFTP endpoint. Calling the DataSync API from an application allows automating the data transfer instead of running scheduled tasks or scripts. DataSync directly transfers files from the network share to S3 without needing an intermediate server

upvoted 2 times
- pentium75**

10 months, 1 week ago

"Create an application" hell no, the application must run somewhere etc., this is massive "administrative overhead".

upvoted 4 times
- antropaws**

1 year, 5 months ago

Selected Answer: B

B. Data Sync is better for one time migrations.

upvoted 3 times
- kruasan**

1 year, 6 months ago

Selected Answer: B

The correct solution here is:

B. Create an Amazon S3 File Gateway. Update the business system to use a new network share from the S3 File Gateway.

This option requires the least administrative overhead because:

  - It presents a simple network file share interface that the business system can write to, just like a standard network share. This requires minimal

changes to the business system.

- The S3 File Gateway automatically uploads all files written to the share to an S3 bucket in the background. This handles the transfer and upload to S3 without requiring any scheduled tasks, scripts or automation.

- All ongoing management like monitoring, scaling, patching etc. is handled by AWS for the S3 File Gateway.

upvoted 5 times

  **kruasan** 1 year, 6 months ago

The other options would require more ongoing administrative effort:

A) AWS DataSync would require creating and managing scheduled tasks and monitoring them.

C) Using the DataSync API would require developing an application and then managing and monitoring it.

D) The SFTP option would require creating scripts, managing SFTP access and keys, and monitoring the file transfer process.

So overall, the S3 File Gateway requires the least amount of ongoing management and administration as it presents a simple file share interface but handles the upload to S3 in a fully managed fashion. The business system can continue writing to a network share as is, while the files are transparently uploaded to S3.

The S3 File Gateway is the most hands-off, low-maintenance solution in this scenario.

upvoted 4 times

  **elearningtakai** 1 year, 7 months ago

**Selected Answer: B**

A - creating a scheduled task is not near-real time.

B - The S3 File Gateway caches frequently accessed data locally and automatically uploads it to Amazon S3, providing near-real-time access to the data.

C - creating an application that uses the DataSync API in the automation workflow may provide near-real-time data access, but it requires additional development effort.

D - it requires additional development effort.

upvoted 5 times

  **zooba72** 1 year, 7 months ago

**Selected Answer: B**

It's B. DataSync has a scheduler and it runs on hour intervals, it cannot be used real-time

upvoted 2 times

  **Buruguduystunstugudunstuy** 1 year, 7 months ago

**Selected Answer: C**

The correct answer is C. Use AWS DataSync to transfer the files to Amazon S3. Create an application that uses the DataSync API in the automation workflow.

To store the CSV reports generated by the business system in the AWS Cloud in near-real time for analysis, the best solution with the least administrative overhead would be to use AWS DataSync to transfer the files to Amazon S3 and create an application that uses the DataSync API in the automation workflow.

AWS DataSync is a fully managed service that makes it easy to automate and accelerate data transfer between on-premises storage systems and AWS Cloud storage, such as Amazon S3. With DataSync, you can quickly and securely transfer large amounts of data to the AWS Cloud, and you can automate the transfer process using the DataSync API.

upvoted 4 times



  **Buruguduystunstugudunstuy** 1 year, 7 months ago

Answer A, using AWS DataSync to transfer the files to Amazon S3 and creating a scheduled task that runs at the end of each day, is not the best solution because it does not meet the requirement of storing the CSV reports in near-real time for analysis.

Answer B, creating an Amazon S3 File Gateway and updating the business system to use a new network share from the S3 File Gateway, is not the best solution because it requires additional configuration and management overhead.

Answer D, deploying an AWS Transfer for the SFTP endpoint and creating a script to check for new files on the network share and upload the new files using SFTP, is not the best solution because it requires additional scripting and management overhead

upvoted 3 times



  **COTIT** 1 year, 7 months ago

**Selected Answer: B**

I think B is the better answer, "LEAST administrative overhead"

[https://aws.amazon.com/storagegateway/file/?nc1=h\\_ls](https://aws.amazon.com/storagegateway/file/?nc1=h_ls)



upvoted 5 times

  **andyto** 1 year, 7 months ago

B - S3 File Gateway.

C - this is wrong answer because data migration is scheduled (this is not continuous task), so condition "near-real time" is not fulfilled

upvoted 4 times

  **Thief** 1 year, 7 months ago

C is the best ans

upvoted 1 times

  **lizzard812** 1 year, 7 months ago

Why not A? There is no scheduled job?

upvoted 2 times

A company is storing petabytes of data in Amazon S3 Standard. The data is stored in multiple S3 buckets and is accessed with varying frequency. The company does not know access patterns for all the data. The company needs to implement a solution for each S3 bucket to optimize the cost of S3 usage.



Which solution will meet these requirements with the MOST operational efficiency?

- A. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Intelligent-Tiering.
- B. Use the S3 storage class analysis tool to determine the correct tier for each object in the S3 bucket. Move each object to the identified storage tier.
- C. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Glacier Instant Retrieval.
- D. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 One Zone-Infrequent Access (S3 One Zone-IA).

**Correct Answer: A**



*Community vote distribution*

A (100%)

  **TariqKipkemei** Highly Voted 11 months, 1 week ago



**Selected Answer: A**

Unknown access patterns for the data = S3 Intelligent-Tiering  
upvoted 8 times

  **channn** Highly Voted 1 year, 1 month ago



**Selected Answer: A**

Key words: 'The company does not know access patterns for all the data', so A.  
upvoted 5 times

  **Guru4Cloud** Most Recent 8 months, 1 week ago

**Selected Answer: A**

Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Intelligent-Tiering.  
upvoted 3 times



  **Buruguduystunstugudunstuy** 1 year, 1 month ago

**Selected Answer: A**

The correct answer is A.

Creating an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Intelligent-Tiering would be the most efficient solution to optimize the cost of S3 usage. S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers (frequent and infrequent) based on changing access patterns. It is a cost-effective solution that does not require any manual intervention to move data to different storage classes, unlike the other options.

upvoted 5 times



  **Buruguduystunstugudunstuy** 1 year, 1 month ago

Answer B, Using the S3 storage class analysis tool to determine the correct tier for each object and manually moving objects to the identified storage tier would be time-consuming and require more operational overhead.

Answer C, Transitioning objects to S3 Glacier Instant Retrieval would be appropriate for data that is accessed less frequently and does not require immediate access.

Answer D, S3 One Zone-IA would be appropriate for data that can be recreated if lost and does not require the durability of S3 Standard or S3 Standard-IA.

upvoted 3 times

  **COTIT** 1 year, 1 month ago

**Selected Answer: A**



For me is A. Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Intelligent-Tiering.

Why?

"S3 Intelligent-Tiering is the ideal storage class for data with unknown, changing, or unpredictable access patterns"

<https://aws.amazon.com/s3/storage-classes/intelligent-tiering/>

upvoted 3 times

  **Bofi** 1 year, 1 month ago

**Selected Answer: A**

Once the data traffic is unpredictable, Intelligent-Tiering is the best option  
upvoted 3 times

  **NIL8891** 1 year, 1 month ago

**Selected Answer: A**

Create an S3 Lifecycle configuration with a rule to transition the objects in the S3 bucket to S3 Intelligent-Tiering.  
upvoted 2 times

  **Maximus007** 1 year, 1 month ago

**Selected Answer: A**

A: as exact pattern is not clear  
upvoted 3 times

A rapidly growing global ecommerce company is hosting its web application on AWS. The web application includes static content and dynamic content. The website stores online transaction processing (OLTP) data in an Amazon RDS database The website’s users are experiencing slow page loads.

Which combination of actions should a solutions architect take to resolve this issue? (Choose two.)

- A. Configure an Amazon Redshift cluster.
- B. Set up an Amazon CloudFront distribution.
- C. Host the dynamic web content in Amazon S3.
- D. Create a read replica for the RDS DB instance.
- E. Configure a Multi-AZ deployment for the RDS DB instance.

Correct Answer: BD

Community vote distribution

BD (90%)

8%

  **Buruguduystunstugudunstuy** Highly Voted 2 years, 1 month ago

Selected Answer: BD

To resolve the issue of slow page loads for a rapidly growing e-commerce website hosted on AWS, a solutions architect can take the following two actions:

- 1. Set up an Amazon CloudFront distribution
- 2. Create a read replica for the RDS DB instance

Configuring an Amazon Redshift cluster is not relevant to this issue since Redshift is a data warehousing service and is typically used for the analytical processing of large amounts of data.

Hosting the dynamic web content in Amazon S3 may not necessarily improve performance since S3 is an object storage service, not a web application server. While S3 can be used to host static web content, it may not be suitable for hosting dynamic web content since S3 doesn't support server-side scripting or processing.

Configuring a Multi-AZ deployment for the RDS DB instance will improve high availability but may not necessarily improve performance.  
upvoted 16 times

  **acts268** Highly Voted 2 years, 1 month ago

Selected Answer: BD

Cloud Front and Read Replica  
upvoted 5 times

  **pentium75** Most Recent 1 year, 4 months ago

Selected Answer: BD

- A - Redshift is for OLAP, not OLTP
  - B - Caching, reduces page load time and server load
  - C - S3 can't host dynamic (!) content
  - D - Read Replica is meant for increasing DB performance
  - E - Multi-AZ is meant for HA (not asked here)
- upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

Selected Answer: BD

The two options that will best help resolve the slow page loads are:



B) Set up an Amazon CloudFront distribution

and

E) Configure a Multi-AZ deployment for the RDS DB instance

Explanation:

CloudFront can cache static content globally and improve latency for static content delivery.  
Multi-AZ RDS improves performance and availability of the database driving dynamic content.  
upvoted 4 times

  **MatAlves** 7 months, 3 weeks ago

Wrong. Multi-AZ is meant to "enhance availability by deploying a standby instance in a second AZ, and achieve fault tolerance in the event of an AZ or database instance failure."

It does not improve performance.  
upvoted 1 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: BD**

BD is correct.  
upvoted 4 times

  **TariqKipkemei** 1 year, 11 months ago



**Selected Answer: BD**

Resolve latency = Amazon CloudFront distribution and read replica for the RDS DB  
upvoted 5 times

  **SamDouk** 2 years, 1 month ago

**Selected Answer: BD**

B and D  
upvoted 3 times

  **klayytech** 2 years, 1 month ago

**Selected Answer: BD**

The website's users are experiencing slow page loads.


To resolve this issue, a solutions architect should take the following two actions:

Create a read replica for the RDS DB instance. This will help to offload read traffic from the primary database instance and improve performance.  
upvoted 3 times

  **zooba72** 2 years, 1 month ago

**Selected Answer: BD**

Question asked about performance improvements, not HA. Cloudfront & Read Replica  
upvoted 3 times

  **thaotnt** 2 years, 1 month ago


**Selected Answer: BD**

slow page loads. > > > D  
upvoted 3 times

  **andyto** 2 years, 1 month ago

**Selected Answer: BD**

Read Replica will speed up Reads on RDS DB.  
E is wrong. It brings HA but doesn't contribute to speed which is impacted in this case. Multi-AZ is Active-Standby solution.  
upvoted 1 times

  **COTIT** 2 years, 1 month ago

**Selected Answer: BE**

I agree with B & E.  
B. Set up an Amazon CloudFront distribution. (Amazon CloudFront is a content delivery network (CDN) service)  
E. Configure a Multi-AZ deployment for the RDS DB instance. (Good idea for loadbalance the DB workflow)  
upvoted 2 times

  **pentium75** 1 year, 4 months ago

Multi-AZ for HA, Read Replica for Scalability

[https://aws.amazon.com/rds/features/read-replicas/?nc1=h\\_ls](https://aws.amazon.com/rds/features/read-replicas/?nc1=h_ls)  
upvoted 2 times

  **Santosh43** 2 years, 1 month ago

B and E ( as there is nothing mention about read transactions)  
upvoted 1 times



  **awsgeek75** 1 year, 3 months ago

Why E? There is nothing mentioned about High Availability also. E is wrong because Multi AZ won't help with scaling  
upvoted 2 times

  **Akademik6** 2 years, 1 month ago

**Selected Answer: BD**



Cloudfront and Read Replica. We don't need HA here.  
upvoted 4 times

  **Bofi** 2 years, 1 month ago

**Selected Answer: BE**

Amazon CloudFront can handle both static and Dynamic contents hence there is not need for option C I.e hosting the static data on Amazon S3. RDS read replica will reduce the amount of reads on the RDS hence leading a better performance. Multi-AZ is for disaster Recovery , which means D is also out.

upvoted 1 times

  **Thief** 2 years, 1 month ago

**Selected Answer: BC**

CloudFont with S3

upvoted 1 times

  **pentium75** 1 year, 4 months ago

S3 can't host "dynamic content"

upvoted 3 times

  **NIL8891** 2 years, 1 month ago

**Selected Answer: BE**

B and E

upvoted 2 times

A company uses Amazon EC2 instances and AWS Lambda functions to run its application. The company has VPCs with public subnets and private subnets in its AWS account. The EC2 instances run in a private subnet in one of the VPCs. The Lambda functions need direct network access to the EC2 instances for the application to work.

The application will run for at least 1 year. The company expects the number of Lambda functions that the application uses to increase during that time. The company wants to maximize its savings on all application resources and to keep network latency between the services low.

Which solution will meet these requirements?

A. Purchase an EC2 Instance Savings Plan Optimize the Lambda functions' duration and memory usage and the number of invocations. Connect the Lambda functions to the private subnet that contains the EC2 instances.

B. Purchase an EC2 Instance Savings Plan Optimize the Lambda functions' duration and memory usage, the number of invocations, and the amount of data that is transferred. Connect the Lambda functions to a public subnet in the same VPC where the EC2 instances run.

C. Purchase a Compute Savings Plan. Optimize the Lambda functions' duration and memory usage, the number of invocations, and the amount of data that is transferred. Connect the Lambda functions to the private subnet that contains the EC2 instances.

D. Purchase a Compute Savings Plan. Optimize the Lambda functions' duration and memory usage, the number of invocations, and the amount of data that is transferred. Keep the Lambda functions in the Lambda service VPC.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Buruguduystunstugudunstuy** Highly Voted 2 years, 1 month ago

**Selected Answer: C**

Answer C is the best solution that meets the company's requirements.

By purchasing a Compute Savings Plan, the company can save on the costs of running both EC2 instances and Lambda functions. The Lambda functions can be connected to the private subnet that contains the EC2 instances through a VPC endpoint for AWS services or a VPC peering connection. This provides direct network access to the EC2 instances while keeping the traffic within the private network, which helps to minimize network latency.

Optimizing the Lambda functions' duration, memory usage, number of invocations, and amount of data transferred can help to further minimize costs and improve performance. Additionally, using a private subnet helps to ensure that the EC2 instances are not directly accessible from the public internet, which is a security best practice.

upvoted 18 times

  **Buruguduystunstugudunstuy** 2 years, 1 month ago

Answer A is not the best solution because connecting the Lambda functions directly to the private subnet that contains the EC2 instances may not be scalable as the number of Lambda functions increases. Additionally, using an EC2 Instance Savings Plan may not provide savings on the costs of running Lambda functions.

Answer B is not the best solution because connecting the Lambda functions to a public subnet may not be as secure as connecting them to a private subnet. Also, keeping the EC2 instances in a private subnet helps to ensure that they are not directly accessible from the public internet.

Answer D is not the best solution because keeping the Lambda functions in the Lambda service VPC may not provide direct network access to the EC2 instances, which may impact the performance of the application.

upvoted 9 times

  **TariqKipkemei** Highly Voted 1 year, 6 months ago

**Selected Answer: C**

Implement Compute Savings Plan because it applies to Lambda usage as well, then connect the Lambda functions to the private subnet that contains the EC2 instances

upvoted 7 times

  **MatAlves** Most Recent 7 months, 3 weeks ago

"Savings Plans are a flexible pricing model that offer low prices on Amazon EC2, AWS Lambda, and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term."

- That already excludes A and B.

The question requires to "keep network latency between the services low", which can be achieved by connecting the Lambda functions to the private subnet that contains the EC2 instances.

C is the answer.

upvoted 2 times

  **[Removed]** 10 months, 3 weeks ago

C

<https://aws.amazon.com/savingsplans/compute-pricing/>

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

A Compute Savings Plan covers both EC2 and Lambda and allows maximizing savings on all resources.

Optimizing Lambda configuration reduces costs.

Connecting the Lambda functions to the private subnet with the EC2 instances provides direct network access between them, keeping latency low.

The Lambda functions are isolated in the private subnet rather than public, improving security.

upvoted 4 times

  **jaehoon090** 1 year, 9 months ago

CCCCCCCCCCCCCCCCCCCC

upvoted 2 times

  **elearningtakai** 2 years, 1 month ago

**Selected Answer: C**

Connect Lambda to Private Subnet contains EC2



upvoted 2 times

  **zooba72** 2 years, 1 month ago

**Selected Answer: C**

Compute savings plan covers both EC2 & Lambda

upvoted 6 times

  **Zox42** 2 years, 1 month ago

C. I would go with C, because Compute savings plans cover Lambda as well.

upvoted 5 times

  **andyto** 2 years, 1 month ago

A. I would go with A. Saving and low network latency are required.

EC2 instance savings plans offer savings of up to 72%

Compute savings plans offer savings of up to 66%

Placing Lambda on the same private network with EC2 instances provides the lowest latency.

upvoted 1 times

  **abitwrong** 2 years, 1 month ago

EC2 Instance Savings Plans apply to EC2 usage only. Compute Savings Plans apply to usage across Amazon EC2, AWS Lambda, and AWS Fargate. (<https://aws.amazon.com/savingsplans/faq/>)

Lambda functions need direct network access to the EC2 instances for the application to work and these EC2 instances are in the private subnet. So the correct answer is C.

upvoted 3 times

A solutions architect needs to allow team members to access Amazon S3 buckets in two different AWS accounts: a development account and a production account. The team currently has access to S3 buckets in the development account by using unique IAM users that are assigned to an IAM group that has appropriate permissions in the account.

The solutions architect has created an IAM role in the production account. The role has a policy that grants access to an S3 bucket in the production account.

Which solution will meet these requirements while complying with the principle of least privilege?

- A. Attach the Administrator Access policy to the development account users.
- B. Add the development account as a principal in the trust policy of the role in the production account.
- C. Turn off the S3 Block Public Access feature on the S3 bucket in the production account.
- D. Create a user in the production account with unique credentials for each team member.

Correct Answer: B

Community vote distribution


B (100%)

  **kels1** Highly Voted 1 year ago

well, if you made it this far, it means you are persistent :) Good luck with your exam!  
upvoted 79 times

  **SkyZeroZx** 12 months ago

Thanks good luck for all  
upvoted 10 times

  **Kimnesh** 8 months, 2 weeks ago

thank you!  
upvoted 6 times

  **gpt\_test** Highly Voted 1 year, 1 month ago

**Selected Answer: B**

By adding the development account as a principal in the trust policy of the IAM role in the production account, you are allowing users from the development account to assume the role in the production account. This allows the team members to access the S3 bucket in the production account without granting them unnecessary privileges.  
upvoted 9 times

  **digital\_eye** Most Recent 1 month, 1 week ago

**Selected Answer: B**

Good luck everyone  
upvoted 1 times

  **TariqKipkemei** 6 months, 1 week ago

**Selected Answer: B**

Add the development account as a principal in the trust policy of the role in the production account  
upvoted 3 times

  **Guru4Cloud** 8 months, 1 week ago

**Selected Answer: B**

The best solution is B) Add the development account as a principal in the trust policy of the role in the production account.

This allows cross-account access to the S3 bucket in the production account by assuming the IAM role. The development account users can assume the role to gain temporary access to the production bucket.  
upvoted 5 times

  **nilandd44gg** 10 months ago

**Selected Answer: B**

<https://aws.amazon.com/blogs/security/how-to-use-trust-policies-with-iam-roles/>

An AWS account accesses another AWS account – This use case is commonly referred to as a cross-account role pattern. It allows human or machine IAM principals from one AWS account to assume this role and act on resources within a second AWS account. A role is assumed to enable this behavior when the resource in the target account doesn't have a resource-based policy that could be used to grant cross-account access.

upvoted 3 times

  **elearningtakai** 1 year, 1 month ago

**Selected Answer: B**

About Trust policy – The trust policy defines which principals can assume the role, and under which conditions. A trust policy is a specific type of resource-based policy for IAM roles.

Answer A: overhead permission Admin to development.

Answer C: Block public access is a security best practice and seems not relevant to this scenario.

Answer D: difficult to manage and scale

upvoted 3 times

  **Buruguduystunstugudunstuy** 1 year, 1 month ago

**Selected Answer: B**

Answer A, attaching the Administrator Access policy to development account users, provides too many permissions and violates the principle of least privilege. This would give users more access than they need, which could lead to security issues if their credentials are compromised.

Answer C, turning off the S3 Block Public Access feature, is not a recommended solution as it is a security best practice to enable S3 Block Public Access to prevent accidental public access to S3 buckets.

Answer D, creating a user in the production account with unique credentials for each team member, is also not a recommended solution as it can be difficult to manage and scale for large teams. It is also less secure, as individual user credentials can be more easily compromised.

upvoted 3 times

  **klayytech** 1 year, 1 month ago

**Selected Answer: B**

The solution that will meet these requirements while complying with the principle of least privilege is to add the development account as a principal in the trust policy of the role in the production account. This will allow team members to access Amazon S3 buckets in two different AWS accounts while complying with the principle of least privilege.

Option A is not recommended because it grants too much access to development account users. Option C is not relevant to this scenario. Option D is not recommended because it does not comply with the principle of least privilege.

upvoted 2 times

  **Akademik6** 1 year, 1 month ago

**Selected Answer: B**

B is the correct answer

upvoted 3 times

A company uses AWS Organizations with all features enabled and runs multiple Amazon EC2 workloads in the ap-southeast-2 Region. The company has a service control policy (SCP) that prevents any resources from being created in any other Region. A security policy requires the company to encrypt all data at rest.

An audit discovers that employees have created Amazon Elastic Block Store (Amazon EBS) volumes for EC2 instances without encrypting the volumes. The company wants any new EC2 instances that any IAM user or root user launches in ap-southeast-2 to use encrypted EBS volumes. The company wants a solution that will have minimal effect on employees who create EBS volumes.

Which combination of steps will meet these requirements? (Choose two.)

- A. In the Amazon EC2 console, select the EBS encryption account attribute and define a default encryption key.
- B. Create an IAM permission boundary. Attach the permission boundary to the root organizational unit (OU). Define the boundary to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
- C. Create an SCP. Attach the SCP to the root organizational unit (OU). Define the SCP to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
- D. Update the IAM policies for each account to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
- E. In the Organizations management account, specify the Default EBS volume encryption setting.

Correct Answer: CE

Community vote distribution

CE (76%)

13%

11%

 **Guru4Cloud** Highly Voted 1 year, 2 months ago

Selected Answer: CE

The correct answer is (C) and (E).

Option (C): Creating an SCP and attaching it to the root organizational unit (OU) will deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false. This means that any IAM user or root user in any account in the organization will not be able to create an EBS volume without encrypting it.

Option (E): Specifying the Default EBS volume encryption setting in the Organizations management account will ensure that all new EBS volumes created in any account in the organization are encrypted by default.

upvoted 11 times

 **AxauS** Highly Voted 1 year, 5 months ago

Selected Answer: CE

CE

Prevent future issues by creating a SCP and set a default encryption.

upvoted 10 times

 **network\_enthusiast** Most Recent 2 weeks ago

Selected Answer: AC

Answer will be A&C.

AWS does not currently provide a direct setting at the organization level to enforce default EBS encryption across all member accounts in AWS Organizations. However, you can achieve this goal by implementing default encryption at the account level and enforcing compliance using an SCP (Service Control Policy).


upvoted 1 times

 **Andreew883** 2 weeks, 5 days ago

Selected Answer: CE

C and E as shaded by guru4cloud

upvoted 1 times

 **Jazz888** 4 months, 3 weeks ago

The problem here is we don't know in which account the workload is on. The account in ap-xx-is that the management account or it's a member account?? That will decide to select either A or E. C is certainly correct

upvoted 2 times

 **NSA\_Poker** 5 months ago

Selected Answer: CE

(A) is incorrect bc absent of SCP or the Organizations management account, the scope of the EC2 console is too narrow to be applied to 'any IAM user or root user'.

upvoted 2 times

  **venutadi** 6 months, 1 week ago

**Selected Answer: AC**

<https://repost.aws/knowledge-center/ebs-automatic-encryption>

Newly created Amazon EBS volumes aren't encrypted by default. However, you can turn on default encryption for new EBS volumes and snapshot copies that are created within a specified Region. To turn on encryption by default, use the Amazon Elastic Compute Cloud (Amazon EC2) console.

upvoted 2 times

  **1rob** 10 months ago

**Selected Answer: AC**

A: will enforce automatic encryption in a account. This will have no effect on employees. Do this in every account.

B: permission boundary is not appropriate here.

C: an SCP will force employees to create encrypted volumes in every account.

D: This would work but is too much maintenance.

E: Setting EBS volume encryption in the Organizations management account will only have impact on volumes in that account, not on other accounts.

upvoted 2 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: AE**

The solution should "have minimal effect on employees who create EBS volumes". Thus new volumes should automatically be encrypted. Options B, C and D do NOT automatically encrypt volumes, they simply cause requests to create non-encrypted volumes to fail.

upvoted 3 times

  **dkw2342** 8 months ago

IMO the correct solution is AC:

In the Amazon EC2 console, select the EBS encryption account attribute and define a default encryption key.

-> This has to be done in every AWS account separately.

Create an SCP. Attach the SCP to the root organizational unit (OU). Define the SCP to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.

-> This will just act as a safeguard in case an admin would disable default encryption in the member account, so it should not have any effect on employees who create EBS volumes.

I think an updated question would offer options A and an updated C:

Create an SCP. Attach the SCP to the root organizational unit (OU). Define the SCP to deny the ec2:DisableEbsEncryptionByDefault action.

-> This will prevent disabling default encryption once it has been enabled.

upvoted 2 times

  **Valder21** 1 year, 1 month ago

Wondering if just C would be sufficient?

upvoted 1 times

  **bjexamprep** 1 year, 2 months ago

Seems many people selected E as part of the correct answer. But I didn't find so called Organization level EBS default setting in my Organization management account. I tried setting default EBS encryption setting in my Organization management account, and it didn't apply to the member account. If E cannot guarantee default encryption in all other account, E has no advantage over A. Anyone can explain why E is better than A?

upvoted 4 times

  **novelai\_me** 1 year, 4 months ago

**Selected Answer: AE**

Option A: By default, EBS encryption is not enabled for EC2 instances. However, you can set an EBS encryption by default in your AWS account in the Amazon EC2 console. This ensures that every new EBS volume that is created is encrypted.

Option E: With AWS Organizations, you can centrally set the default EBS encryption for your organization's accounts. This helps in enforcing a consistent encryption policy across your organization.

Option B, C and D are not correct because while you can use IAM policies or SCPs to restrict the creation of unencrypted EBS volumes, this could potentially impact employees' ability to create necessary resources if not properly configured. They might require additional permissions management, which is not mentioned in the requirements. By setting the EBS encryption by default at the account or organization level (Options A and E), you can ensure all new volumes are encrypted without affecting the ability of employees to create resources.

upvoted 3 times


  **Buruguduystunstugudunstuy** 1 year, 4 months ago

**Selected Answer: CE**

SCPs are a great way to enforce policies across an entire AWS Organization, preventing users from creating resources that do not comply with the set policies.

In AWS Management Console, one can go to EC2 dashboard -> Settings -> Data encryption -> Check "Always encrypt new EBS volumes" and choose a default KMS key. This ensures that every new EBS volume created will be encrypted by default, regardless of how it is created.

upvoted 3 times

  **PRASAD180** 1 year, 5 months ago

1000% CE crt

upvoted 1 times

  **[Removed]** 1 year, 5 months ago

Encryption by default allows you to ensure that all new EBS volumes created in your account are always encrypted, even if you don't specify encrypted=true request parameter.

<https://aws.amazon.com/blogs/compute/must-know-best-practices-for-amazon-ebs-encryption/>



upvoted 1 times

  **hiroohiroo** 1 year, 5 months ago

**Selected Answer: CE**

CとEが正しいと考える。

upvoted 3 times

  **Efren** 1 year, 5 months ago

**Selected Answer: CE**

CE for me as well

upvoted 2 times

A company wants to use an Amazon RDS for PostgreSQL DB cluster to simplify time-consuming database administrative tasks for production database workloads. The company wants to ensure that its database is highly available and will provide automatic failover support in most scenarios in less than 40 seconds. The company wants to offload reads off of the primary instance and keep costs as low as possible.

Which solution will meet these requirements?

- A. Use an Amazon RDS Multi-AZ DB instance deployment. Create one read replica and point the read workload to the read replica.
- B. Use an Amazon RDS Multi-AZ DB duster deployment Create two read replicas and point the read workload to the read replicas.
- C. Use an Amazon RDS Multi-AZ DB instance deployment. Point the read workload to the secondary instances in the Multi-AZ pair.
- D. Use an Amazon RDS Multi-AZ DB cluster deployment Point the read workload to the reader endpoint.

**Correct Answer: D**

*Community vote distribution*

D (86%)

13%

  **ogerber** Highly Voted 1 year, 11 months ago

**Selected Answer: D**

A - multi-az instance : failover takes between 60-120 sec

D - multi-az cluster: failover around 35 sec

upvoted 20 times

  **kelmryan1** 12 months ago

They want to keep cost low as possible, A is the right answer

upvoted 2 times

  **MatAlves** 7 months, 3 weeks ago

No, the question clearly says "failover support in most scenarios in less than 40 seconds."

D is the only possible answer.

upvoted 2 times

  **Buruguduystunstugudunstuy** Highly Voted 1 year, 10 months ago

**Selected Answer: D**

The correct answer is:

D. Use an Amazon RDS Multi-AZ DB cluster deployment. Point the read workload to the reader endpoint.

Explanation:

The company wants high availability, automatic failover support in less than 40 seconds, read offloading from the primary instance, and cost-effectiveness.

Answer D is the best choice for several reasons:

1. Amazon RDS Multi-AZ deployments provide high availability and automatic failover support.
2. In a Multi-AZ DB cluster, Amazon RDS automatically provisions and maintains a standby in a different Availability Zone. If a failure occurs, Amazon RDS performs an automatic failover to the standby, minimizing downtime.
3. The "Reader endpoint" for an Amazon RDS DB cluster provides load-balancing support for read-only connections to the DB cluster. Directing read traffic to the reader endpoint helps in offloading read operations from the primary instance.

upvoted 13 times

  **Kiki\_Pass** 1 year, 9 months ago

Sorry I'm a bit confused... I thought only Aurora DB Cluster has reader endpoint. Do you by any chance has the link to the doc for RDS Reader Endpoint?

upvoted 3 times

  **lemur88** 1 year, 8 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/multi-az-db-clusters-concepts-connection-management.html#multi-az-db-clusters-concepts-connection-management-endpoints-reader>

upvoted 6 times

  **zikou** Most Recent 8 months, 2 weeks ago



to offload read we use read replicas also there is no such thing as reader endpoint in rds, it is only on aurora

upvoted 2 times

  **rondelldell** 1 year ago

**Selected Answer: B**

<https://aws.amazon.com/rds/features/multi-az/>  
Amazon RDS Multi-AZ with two readable standbys  
upvoted 1 times

  **hro** 1 year, 1 month ago

I think the cluster is over-kill - but the company 'wants to use an Amazon RDS ... DB cluster'.  
upvoted 3 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: D**

A would be cheapest but "failover times are typically 60–120 seconds" which does not meet our requirements. We need Multi-AZ DB cluster (not instance). This has a reader endpoint by default, thus no need for additional read replicas (to "keep costs as low as possible").  
<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZSingleStandby.html>  
upvoted 7 times

  **master9** 1 year, 4 months ago

**Selected Answer: A**

in question, it has mentioned that "keep costs as low as possible"

In a Multi-AZ configuration, the DB instances and EBS storage volumes are deployed across two Availability Zones. It provides high availability and failover support for DB instances. This setup is primarily for disaster recovery. It involves a primary DB instance and a standby replica, which is a copy of the primary DB instance. The standby replica is not accessible directly; instead, it serves as a failover target in case the primary instance fails.  
upvoted 1 times

  **potomac** 1 year, 6 months ago

**Selected Answer: D**

It is D.  
A is not correct. Multi-AZ DB instance deployment, which creates a primary instance and a standby instance to provide failover support. However, the standby instance does not serve traffic.  
upvoted 2 times

  **maudsha** 1 year, 6 months ago

**Selected Answer: D**

<https://aws.amazon.com/blogs/database/choose-the-right-amazon-rds-deployment-option-single-az-instance-multi-az-instance-or-multi-az-database-cluster/#:~:text=Unlike%20Multi%2DAZ%20instance%20deployment,different%20AZs%20serving%20read%20traffic.>

According to this the answer is D  
"Unlike Multi-AZ instance deployment, where the secondary instance can't be accessed for read or writes, Multi-AZ DB cluster deployment consists of primary instance running in one AZ serving read-write traffic and two other standby running in two different AZs serving read traffic."

You don't have to create read replicas with cluster deployment so B is out.  
upvoted 2 times

  **kwang312** 1 year, 7 months ago

D  
Fail-over on Multi-AZ DB instance is 60-120s  
On Cluster, the time under 35s  
upvoted 5 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

D. Use an Amazon RDS Multi-AZ DB cluster deployment. Point the read workload to the reader endpoint  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

Use an Amazon RDS Multi-AZ DB cluster deployment Point the read workload to the reader endpoint.  
upvoted 2 times

  **Eminenza22** 1 year, 8 months ago

**Selected Answer: A**

The solutions architect should use an Amazon RDS Multi-AZ DB instance deployment. The company can create one read replica and point the read workload to the read replica. Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments.  
upvoted 1 times

  **Gooniegoogoo** 1 year, 10 months ago

and d..

Multi-AZ DB clusters typically have lower write latency when compared to Multi-AZ DB instance deployments. They also allow read-only workloads to run on reader DB instances.  
upvoted 2 times

🗨️ 👤 **TariqKipkemei** 1 year, 11 months ago

**Selected Answer: D**

This is as case where both option A and D can work, but option D gives 2 DB instances for read compared to only 1 given by option A. Costwise they are the same as both options use 3 DB instances.

upvoted 2 times

🗨️ 👤 **Henrytml** 1 year, 11 months ago

**Selected Answer: A**

lowest cost option, and effective with read replica

upvoted 3 times

🗨️ 👤 **antropaws** 1 year, 11 months ago

**Selected Answer: D**

It's D. Read well: "A company wants to use an Amazon RDS for PostgreSQL DB CLUSTER".

upvoted 4 times

A company runs a highly available SFTP service. The SFTP service uses two Amazon EC2 Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.

The company wants a serverless option that provides high IOPS performance and highly configurable security. The company also wants to maintain control over user permissions.

Which solution will meet these requirements?

- A. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses. Attach the EBS volume to the SFTP service endpoint. Grant users access to the SFTP service.
- B. Create an encrypted Amazon Elastic File System (Amazon EFS) volume. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing access. Attach a security group to the endpoint that allows only trusted IP addresses. Attach the EFS volume to the SFTP service endpoint. Grant users access to the SFTP service.
- C. Create an Amazon S3 bucket with default encryption enabled. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses. Attach the S3 bucket to the SFTP service endpoint. Grant users access to the SFTP service.
- D. Create an Amazon S3 bucket with default encryption enabled. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subnet. Attach a security group that allows only trusted IP addresses. Attach the S3 bucket to the SFTP service endpoint. Grant users access to the SFTP service.

**Correct Answer: B**

*Community vote distribution*

B (86%)

8%

  **pentium75** Highly Voted 1 year, 4 months ago

**Selected Answer: B**

Not A - Transfer Family can't use EBS

B - Possible and meets requirement

Not C - S3 doesn't guarantee "high IOPS performance"; also there is no "public endpoint that allows only trusted IP addresses" (you can assign a Security Group to a public endpoint but that is not mentioned here)

Not D - Endpoint would be in private subnet, not accessible from Internet at all

upvoted 9 times

  **alexandercamachop** Highly Voted 1 year, 11 months ago

**Selected Answer: B**

First Serverless - EFS

Second it says it is attached to the Linux instances at the same time, only EFS can do that.

upvoted 7 times

  **FlyingHawk** Most Recent 5 months ago

**Selected Answer: B**

D, the description "VPC endpoint that has internal access in a private subnet" is technically incorrect and doesn't make sense for AWS Transfer Family.

AWS Transfer Family actually supports:

Public endpoints (internet-facing)

VPC endpoints (private network access)

The VPC endpoint doesn't have "internal access" - it's about how the SFTP service is configured to interact with your VPC network resources.

upvoted 1 times

  **FlyingHawk** 5 months ago

A more accurate description for a VPC-based AWS Transfer Family SFTP service would be:

Create an AWS Transfer Family SFTP service with a VPC endpoint

Place the endpoint in a private subnet

Configure security groups to allow only trusted IP sources

Attach the S3 bucket to the SFTP service endpoint

So while the overall strategy in option D (using S3 with VPC endpoint and security group restrictions) could be valid, the specific wording about "internal access" is incorrect.

This technical inaccuracy in the description would make D an incorrect answer in a certification or technical assessment.

upvoted 1 times

  **FlyingHawk** 5 months ago

C public endpoint that allows only trusted IP addresses does not sound technical correct as only the security group can restrict the IP address. the correct description should " Create an AWS Transfer Family SFTP service in public subnet with security group allows only trusted IP address, create the vpc endpoint (gateway endpoint here for saving the cost) of S3 bucket to allow the SFTP service can access it privately, grant user access to the SFTP service via IAM policies and grant the access of S3 to SFTP service via IAM role

upvoted 1 times

  **FlyingHawk** 5 months ago

**Selected Answer: C**

Amazon S3 provides high throughput and performance suitable for many use cases, including those requiring high IOPS (Input/Output Operations Per Second). S3 is not a block storage solution (like EBS) or a file system (like EFS). While its performance is exceptional for its intended use case (object storage), it may not match the millisecond latency or consistent high IOPS required for transactional databases or other ultra-low latency applications.



For applications where sub-millisecond latency or extremely high random IOPS (e.g., 64,000 IOPS) is required, solutions like EBS or EFS would be better.

upvoted 1 times

  **FlyingHawk** 5 months ago

After read the description of C, I feel the description of "an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses" does not sound correct, the endpoint itself cannot restrict the IP address from internet, only the security group. the correct description should be: Create an AWS Transfer Family SFTP service with a public endpoint and security group that allows only trusted IP addresses. create the vpc endpoint (gateway endpoint here for saving the cost) of S3 bucket to allow the SFTP service can access it privately, grant user access to the SFTP service via IAM policies and grant the access of S3 to SFTP service via IAM role

upvoted 1 times

  **JA2018** 5 months, 1 week ago

**Selected Answer: B**

Actually AWS Transfer Family can use S3, so it's a toss-up between Options B & C but I tend to favour Option B for the following reasons (based on the keys in STEM):

company runs a highly available SFTP service. The SFTP service uses two Amazon EC2 Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.

1. Requires a serverless option that provides high IOPS performance and highly configurable security.

2. User also wants to maintain control over user permissions

upvoted 1 times

  **523db89** 8 months, 2 weeks ago

Option B best meets the company's requirements by leveraging AWS Transfer Family with an EFS volume, ensuring high availability, security, and performance.

upvoted 2 times

  **NickGordon** 1 year, 5 months ago

**Selected Answer: B**

A is incorrect as EBS is not an option

C is incorrect as when I select public accessible, I don't see an option I can set up trusted IP address

D is incorrect as it is internal.

B, followed the steps and I can set up a sftp in this way

upvoted 4 times



  **potomac** 1 year, 6 months ago

**Selected Answer: B**

B



EFS has lower latency and higher throughput than S3 when accessed from within the same availability zone.

upvoted 3 times

  **thanhnv142** 1 year, 6 months ago

C: Because it is server-less. definitely not A or B because it utilizes server.

upvoted 1 times

  **warp** 1 year, 6 months ago

Amazon Elastic File System - Serverless, fully elastic file storage:

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

upvoted 5 times

  **bsbs1234** 1 year, 7 months ago

B,

A), transfer family does not support EBS

C,D), S3 has lower IOPS than EFS

upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Create an encrypted Amazon Elastic File System (Amazon EFS) volume. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing access. Attach a security group to the endpoint that allows only trusted IP addresses. Attach the EFS volume to the SFTP service endpoint. Grant users access to the SFTP service.

upvoted 2 times

  **Axeashes** 1 year, 10 months ago

<https://aws.amazon.com/blogs/storage/use-ip-whitelisting-to-secure-your-aws-transfer-for-sftp-servers/>



upvoted 2 times

  **TariqKipkemei** 1 year, 11 months ago

**Selected Answer: B**

EFS is best to serve this purpose.

upvoted 2 times

  **envest** 1 year, 11 months ago

Answer C (from abylead.com)



Transfer Family offers fully managed serverless support for B2B file transfers via SFTP, AS2, FTPS, & FTP directly in & out of S3 or EFS. For a controlled internet access you can use internet-facing endpts with Transfer SFTP servers & restrict trusted internet sources with VPC's default Sgrp. In addition, S3 Access Points aliases allows you to use S3 bkt names for a unique access control plcy on shared S3 datasets. Transfer SFTP & S3: <https://aws.amazon.com/blogs/apn/how-to-use-aws-transfer-family-to-replace-and-scale-sftp-servers/>

A)Transfer SFTP doesn't support EBS, not for share data, & not serverless: infeasible.

B)EFS mounts via ENIs not endpts: infeasible.



D)pub endpt for internet access is missing: infeasible.

upvoted 4 times

  **omoakin** 1 year, 11 months ago

BBBBBBBBBBBBBBB



upvoted 1 times

  **vesen22** 1 year, 11 months ago

**Selected Answer: B**

EFS all day

upvoted 2 times

  **norris81** 1 year, 11 months ago

<https://aws.amazon.com/blogs/storage/use-ip-whitelisting-to-secure-your-aws-transfer-for-sftp-servers/> is worth a read

upvoted 2 times

A company is developing a new machine learning (ML) model solution on AWS. The models are developed as independent microservices that fetch approximately 1 GB of model data from Amazon S3 at startup and load the data into memory. Users access the models through an asynchronous API. Users can send a request or a batch of requests and specify where the results should be sent.

The company provides models to hundreds of users. The usage patterns for the models are irregular. Some models could be unused for days or weeks. Other models could receive batches of thousands of requests at a time.



Which design should a solutions architect recommend to meet these requirements?


- A. Direct the requests from the API to a Network Load Balancer (NLB). Deploy the models as AWS Lambda functions that are invoked by the NLB.
- B. Direct the requests from the API to an Application Load Balancer (ALB). Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from an Amazon Simple Queue Service (Amazon SQS) queue. Use AWS App Mesh to scale the instances of the ECS cluster based on the SQS queue size.
- C. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue. Deploy the models as AWS Lambda functions that are invoked by SQS events. Use AWS Auto Scaling to increase the number of vCPUs for the Lambda functions based on the SQS queue size.
- D. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue. Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from the queue. Enable AWS Auto Scaling on Amazon ECS for both the cluster and copies of the service based on the queue size.

**Correct Answer: D**



*Community vote distribution*



D (100%)


-   **examtopictempacc**

Highly Voted 

 1 year, 5 months ago

asynchronous=SQS, microservices=ECS.  
Use AWS Auto Scaling to adjust the number of ECS services.  
upvoted 17 times
-   **TariqKipkemei** 1 year, 4 months ago



good breakdown :)  
upvoted 4 times
-   **TariqKipkemei**


Highly Voted 

 1 year, 4 months ago

**Selected Answer: D**



For once examtopic answer is correct :) haha...

Batch requests/async = Amazon SQS  
Microservices = Amazon ECS  
Workload variations = AWS Auto Scaling on Amazon ECS  
upvoted 11 times
-   **wizcloudifa**



Most Recent 

 6 months, 3 weeks ago

**Selected Answer: D**

ALB is mentioned in other options to distract you, you dont need ALB for scaling here, we would need ECS autoscaling, they play with that idea in option B a bit however D gets it in a completely optimized way.... A and C both have lambda which for Machine learning models with workloads on heavy side, will not fly  
upvoted 2 times
-   **Guru4Cloud** 1 year, 2 months ago


**Selected Answer: D**

I go with everyone D.  
upvoted 3 times
-   **alexandercamachop** 1 year, 5 months ago

**Selected Answer: D**

D, no need for an App Load balancer like C says, no where in the text.  
SQS is needed to ensure all request gets routed properly in a Microservices architecture and also that it waits until its picked up.  
ECS with Autoscaling, will scale based on the unknown pattern of usage as mentioned.

upvoted 2 times

  **anibinaadi** 1 year, 5 months ago

It is D

Refer <https://aws.amazon.com/blogs/containers/amazon-elastic-container-service-ecs-auto-scaling-using-custom-metrics/> for additional information/knowledge.

upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: D**

because it is scalable, reliable, and efficient.

C does not scale the models automatically

upvoted 4 times

  **deechean** 1 year, 2 months ago

why C doesn't scale the model? Application Auto Scaling can apply to lambda.

upvoted 1 times

  **pentium75** 10 months, 1 week ago

How would you "use Auto Scaling (!) to increase the number of vCPUs (!) for the Lamba functions"?

upvoted 3 times

  **NSA\_Poker** 5 months ago

Auto Scaling doesn't apply to Lambda. As your functions receive more requests, Lambda automatically handles scaling the number of execution environments until you reach your account's concurrency limit.

upvoted 2 times

A solutions architect wants to use the following JSON text as an identity-based policy to grant specific permissions:

```
{
  "Statement": [
    {
      "Action": [
        "ssm:ListDocuments",
        "ssm:GetDocument"
      ],
      "Effect": "Allow",
      "Resource": "*",
      "Sid": ""
    }
  ],
  "Version": "2012-10-17"
}
```

Which IAM principals can the solutions architect attach this policy to? (Choose two.)

- A. Role
- B. Group
- C. Organization
- D. Amazon Elastic Container Service (Amazon ECS) resource
- E. Amazon EC2 resource

**Correct Answer: AB**

*Community vote distribution*

AB (100%)

- nosense**

Highly Voted

1 year, 5 months ago

Selected Answer: AB

identity-based policy used for role and group

upvoted 19 times
- pentium75**

Highly Voted

10 months, 1 week ago

Selected Answer: AB

Isn't the content of the policy completely irrelevant? IAM policies are applied to users, groups or roles ...

upvoted 7 times
- dkw2342**

Most Recent

8 months ago

AB is correct, but the question is misleading because, according to the AWS IAM documentation, groups are not considered principals: <https://docs.aws.amazon.com/IAM/latest/UserGuide/intro-structure.html#intro-structure-principal>.

upvoted 2 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: AB

A. Role

B. Group

upvoted 3 times
- TariqKipkemei**

1 year, 4 months ago

Selected Answer: AB

Role or group

upvoted 3 times

A company is running a custom application on Amazon EC2 On-Demand Instances. The application has frontend nodes that need to run 24 hours a day, 7 days a week and backend nodes that need to run only for a short time based on workload. The number of backend nodes varies during the day.

The company needs to scale out and scale in more instances based on workload.

Which solution will meet these requirements MOST cost-effectively?

- A. Use Reserved Instances for the frontend nodes. Use AWS Fargate for the backend nodes.
- B. Use Reserved Instances for the frontend nodes. Use Spot Instances for the backend nodes.
- C. Use Spot Instances for the frontend nodes. Use Reserved Instances for the backend nodes.
- D. Use Spot Instances for the frontend nodes. Use AWS Fargate for the backend nodes.

**Correct Answer: B**

*Community vote distribution*

B (62%)

A (38%)

- Ramdi1**

Highly Voted

1 year, 7 months ago

Selected Answer: A

Has to be A, It can scale down if required and you will be charged for what you use with fargate. Secondly they have not said the backend can have timeouts or can be down for a little period of time or something. So it has to rule out any spot instances even though they are cheaper.

upvoted 18 times
- awsgeek75**

1 year, 3 months ago

Fargate is serverless EKS so it cannot manage EC2 nodes

upvoted 7 times
- nosense**

Highly Voted

1 year, 11 months ago

Selected Answer: B

Reserved+ spot .  
Fargate for serverless

upvoted 15 times
- Andreew883**

Most Recent

2 weeks, 5 days ago

Selected Answer: B

The question is about instances so container and fargate is not applicable and we need to choose response B

upvoted 1 times
- zdi561**

2 months, 4 weeks ago

Selected Answer: A

I think in A it should be Fargate spot, B means standard spot. The former uses auto scaling and the latter does not

upvoted 1 times
- zdi561**

3 months, 2 weeks ago

Selected Answer: A

The requirement is to have an autoscaling , and fargate can do that. Spot instance only works when the work load is fault tolerant. Furthermore it is told to use on-demand instance.

upvoted 2 times
- Mayur\_B**

6 months, 3 weeks ago

I will go with A.  
Where does the question mention that the workload interruptions are accepted.

upvoted 3 times
- musha**

1 year, 1 month ago

Selected Answer: B

B) because fargate is container

upvoted 3 times
- noircesar25**

1 year, 2 months ago

so what we make up from this scenario is: the key word right here is "backend nodes" you can't use a serverless compute service with nodes and you need to use EC2s

so if we had ECS EC2 lunch type or on-demand EC2s as an options for the backend, they would be true?

upvoted 2 times

  **mwwt2022** 1 year, 3 months ago

**Selected Answer: B**

24-7 usage for fe -> reserved instance  
irregular workload for be -> spot instance  
upvoted 4 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: B**

Not A because Fargate runs containers, not EC2 instances. But we have no indication that the workload would be containerized; it runs "on EC2 instances".  
Not C and D because frontend must run 24/7, can't use Spot.

Thus B, yes, Spot instances are risky, but as they need to run "only for a short time" it seems acceptable.

Technically ideal option would be Reserved Instances for frontend nodes and On-demand instances for backend nodes, but that is not an option here.  
upvoted 8 times

  **[Removed]** 1 year, 5 months ago

**Selected Answer: B**

Not sure the application can be containerized  
upvoted 3 times

  **AwsZora** 1 year, 5 months ago

**Selected Answer: A**

it is safe  
upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

Fargate = containers  
A is wrong  
upvoted 2 times

  **meowruki** 1 year, 5 months ago

**Selected Answer: B**

Reserved Instances (RIs) for Frontend Nodes: Since the frontend nodes need to run continuously (24/7), using Reserved Instances for them makes sense. RIs provide significant cost savings compared to On-Demand Instances for steady-state workloads.

Spot Instances for Backend Nodes: Spot Instances are suitable for short-duration workloads and can be significantly cheaper than On-Demand Instances. Since the number of backend nodes varies during the day, Spot Instances can help you take advantage of spare capacity at a lower cost. Keep in mind that Spot Instances may be interrupted if the capacity is needed elsewhere, so they are best suited for stateless and fault-tolerant workloads.  
upvoted 2 times

  **meowruki** 1 year, 5 months ago

Option A (Use Reserved Instances for the frontend nodes. Use AWS Fargate for the backend nodes): AWS Fargate is a serverless compute engine for containers, and it may not be the best fit for the described backend workload, especially if the number of backend nodes varies during the day.  
upvoted 1 times

  **Goutham4981** 1 year, 5 months ago

**Selected Answer: B**

AWS Fargate is a serverless compute engine for containers that allows you to run containers without having to manage the underlying infrastructure. It simplifies the process of deploying and managing containerized applications by abstracting away the complexities of server management, scaling, and cluster orchestration.  
No containerized application requirements are mentioned in the question. Plain EC2 instances. So Fargate is not actually an option  
upvoted 3 times

  **thanhnv142** 1 year, 6 months ago

A is fargate, which is none sense. B seems more OK (though none-sense)  
upvoted 4 times

  **dilaaziz** 1 year, 7 months ago

**Selected Answer: A**

Fargate for backend node  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

Fargate is for containers not EC2 so A is wrong  
upvoted 2 times

  **Wayne23Fang** 1 year, 7 months ago

**Selected Answer: A**

(B) would take chance, though unlikely (A) is server-less auto-scaling. In case backend is idle, it might scale down, save money but no need to worry for interruption by Spot instance.

upvoted 3 times

A company uses high block storage capacity to runs its workloads on premises. The company's daily peak input and output transactions per second are not more than 15,000 IOPS. The company wants to migrate the workloads to Amazon EC2 and to provision disk performance independent of storage capacity.

Which Amazon Elastic Block Store (Amazon EBS) volume type will meet these requirements MOST cost-effectively?

- A. GP2 volume type
- B. io2 volume type
- C. GP3 volume type
- D. io1 volume type

**Correct Answer:** C

*Community vote distribution*

C (96%)4%

- nosense**

Highly Voted

1 year, 11 months ago

Selected Answer: C

Gp3 \$ 0.08 usd per gb  
Gp2 \$ 0.10 usd per gb  
upvoted 15 times
- Yadav\_Sanjay**

Highly Voted

1 year, 11 months ago

Selected Answer: C

Both GP2 and GP3 has max IOPS 16000 but GP3 is cost effective.  
<https://aws.amazon.com/blogs/storage/migrate-your-amazon-ebs-volumes-from-gp2-to-gp3-and-save-up-to-20-on-costs/>  
upvoted 12 times
- LeonSauveterre**

Most Recent

5 months ago

Selected Answer: C

OK, and is this something I should know to prepare for the exam...  
upvoted 1 times
- Guru4Cloud**

1 year, 8 months ago

Selected Answer: C

C. GP3 volume type  
upvoted 4 times
- james2033**

1 year, 9 months ago

Selected Answer: C

Quote "customers can scale up to 16,000 IOPS and" at <https://aws.amazon.com/about-aws/whats-new/2020/12/introducing-new-amazon-ebs-general-purpose-volumes-gp3/>  
upvoted 3 times
- alexandercamachop**

1 year, 11 months ago

Selected Answer: C

The GP3 (General Purpose SSD) volume type in Amazon Elastic Block Store (EBS) is the most cost-effective option for the given requirements. GP3 volumes offer a balance of price and performance and are suitable for a wide range of workloads, including those with moderate I/O needs.  
  
GP3 volumes allow you to provision performance independently from storage capacity, which means you can adjust the baseline performance (measured in IOPS) and throughput (measured in MiB/s) separately from the volume size. This flexibility allows you to optimize your costs while meeting the workload requirements.  
  
In this case, since the company's daily peak input and output transactions per second are not more than 15,000 IOPS, GP3 volumes provide a suitable and cost-effective option for their workloads.  
upvoted 2 times
- maver144**


1 year, 11 months ago

Selected Answer: B

It is not C pals. The company wants to migrate the workloads to Amazon EC2 and to provision disk performance independent of storage capacity. With GP3 we have to increase storage capacity to increase IOPS over baseline.  
  
You can only chose IOPS independetly with IO family and IO2 is in general better then IO1.  
upvoted 2 times



  **somsundar** 1 year, 9 months ago

@maver144 - That's the case with GP2 volumes. With GP3 we can define IOPS independent of storage capacity.  
upvoted 4 times

  **Joselucho38** 1 year, 11 months ago

**Selected Answer: C**

Therefore, the most suitable and cost-effective option in this scenario is the GP3 volume type (option C).  
upvoted 2 times

  **Efren** 1 year, 11 months ago

**Selected Answer: C**

GPS3 allows 16000 IOPS  
upvoted 4 times

A company needs to store data from its healthcare application. The application’s data frequently changes. A new regulation requires audit access at all levels of the stored data.

The company hosts the application on an on-premises infrastructure that is running out of storage capacity. A solutions architect must securely migrate the existing data to AWS while satisfying the new regulation.

Which solution will meet these requirements?

- A. Use AWS DataSync to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- B. Use AWS Snowcone to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.
- C. Use Amazon S3 Transfer Acceleration to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- D. Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.

**Correct Answer: A**

*Community vote distribution*

A (69%)

D (31%)

- thanhnv142**

Highly Voted

1 year, 6 months ago

A is better because:
  - Data sync is used for migrate. Storage gw is used to connect on-prem to AWS.
  - dataevents is to log for access, management events is for config or managementupvoted 14 times
- pentium75**

Highly Voted

1 year, 4 months ago

**Selected Answer: A**

We need to log "access at all levels" aka "data events", thus B and D are out (logging only "management events" like granting permissions or changing the access tier).  
C, S3 Transfer Acceleration is to increase upload performance from widespread sources or over unreliable networks, but it just provides an endpoint, it does not upload anything itself.

upvoted 11 times
- hpirnaj**

Most Recent

4 months ago

**Selected Answer: A**

data management shows management operations that are performed on resources in your AWS account so, D is wrong  
<https://repost.aws/knowledge-center/cloudtrail-data-management-events>

upvoted 2 times
- JA2018**

5 months, 1 week ago

**Selected Answer: A**

Keys in STEM:  
  
A company needs to store data from its healthcare application. The application’s data frequently changes. A new regulation requires audit access at all levels of the stored data.  
  
1. Must securely migrate the existing data to AWS while satisfying the new regulation.  
  
2. Logs access at all data levels

upvoted 2 times
- 1e22522**

9 months ago

**Selected Answer: A**

Typical DataSync scenario me thinks!

upvoted 3 times
- osmk**

1 year, 2 months ago

**Selected Answer: A**

Use AWS DataSync to migrate existing data to Amazon S3<https://aws.amazon.com/datasync/faqs/>



upvoted 4 times
- NayeraB**

1 year, 2 months ago

**Selected Answer: A**

It's DataSync for me

upvoted 3 times

  **frmrkc** 1 year, 3 months ago

**Selected Answer: D**



Storage Gateway integration with CloudTrail :

<https://docs.aws.amazon.com/filegateway/latest/filefsxw/logging-using-cloudtrail.html>

whereas DataSync can be monitored with Amazon CloudWatch:

<https://docs.aws.amazon.com/datasync/latest/userguide/monitor-datasync.html>

upvoted 2 times

  **frmrkc** 1 year, 3 months ago

and here are all Storage Gateway actions monitored by CloudTrail:

[https://docs.aws.amazon.com/storagegateway/latest/APIReference/API\\_Operations.html](https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_Operations.html)

upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

B and C don't solve the problem

A is extending the data and management events are for administrative actions only (tracking account creation, user security actions etc.).

C uses DataSync to move all the data and logs data events which include S3 file uploads and downloads.

Management events: User logs into an EC2 instance, creates an S3 IAM role

Data events: User uploads a file to S3

upvoted 4 times

  **benacert** 1 year, 3 months ago

A- DataSync secure fast data transfer

upvoted 2 times

  **ZZZ\_Sleep** 1 year, 4 months ago

**Selected Answer: D**

\*Keyword\* of this question = running out of storage capacity

AWS Storage Gateway = extend the on-premises storage

AWS DataSync = copy data between on-premises storage



So, the answer should be D (AWS Storage Gateway)

upvoted 6 times

  **pentium75** 1 year, 4 months ago

"Securely migrate the existing data to AWS" -> move data away from on-premises storage to AWS. Plus, D logs only management events, not "access at all levels".

upvoted 6 times

  **aws94** 1 year, 4 months ago

**Selected Answer: D**

AWS DataSync is designed for fast, simple, and secure data transfer, but it focuses more on data synchronization rather than on-premises migration.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

Thus it is wrong, but more because of the incorrect logging option in this answer.

upvoted 1 times

  **meowruki** 1 year, 5 months ago

**Selected Answer: A**

AWS DataSync is suitable for data transfer and synchronization

Option D (Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events): AWS Storage Gateway is typically used for hybrid cloud storage solutions and may introduce additional complexity for a one-time data migration task. It might not be as straightforward as using AWS Snowcone for this specific scenario.

upvoted 2 times

  **chikuwan** 1 year, 5 months ago

**Selected Answer: A**

both DataSync and Storage Gateway are fine to sync data...but to "audit access at all levels of the stored data" ...it should be data events(data plane operation)..management event is some account level things.

So answer should be A

upvoted 3 times

  **bogobob** 1 year, 5 months ago

**Selected Answer: D**

While both DataSync and Storage Gateway allow syncing of data between on-premise and cloud, DataSync is built for rapid shifting of data into a cloud environment, not specifically for continued use in on-premise servers.

upvoted 2 times

  **potomac** 1 year, 6 months ago

**Selected Answer: A**


AWS DataSync is an online data transfer service that simplifies, automates, and accelerates the process of copying large amounts of data to and from AWS storage services over the Internet or over AWS Direct Connect.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

What about logging?

upvoted 2 times

  **JA2018** 5 months, 1 week ago

This is where AWS CloudTrail comes into play.

upvoted 1 times

  **canonlycontainletters1** 1 year, 6 months ago

**Selected Answer: A**

A seems to be more convincing to me.

upvoted 2 times

A solutions architect is implementing a complex Java application with a MySQL database. The Java application must be deployed on Apache Tomcat and must be highly available.

What should the solutions architect do to meet these requirements?

- A. Deploy the application in AWS Lambda. Configure an Amazon API Gateway API to connect with the Lambda functions.
- B. Deploy the application by using AWS Elastic Beanstalk. Configure a load-balanced environment and a rolling deployment policy.
- C. Migrate the database to Amazon ElastiCache. Configure the ElastiCache security group to allow access from the application.
- D. Launch an Amazon EC2 instance. Install a MySQL server on the EC2 instance. Configure the application on the server. Create an AMI. Use the AMI to create a launch template with an Auto Scaling group.

**Correct Answer:** B

Community vote distribution

B (100%)

- cloudenthusiast

Highly Voted

1 year, 11 months ago

B

AWS Elastic Beanstalk provides an easy and quick way to deploy, manage, and scale applications. It supports a variety of platforms, including Java and Apache Tomcat. By using Elastic Beanstalk, the solutions architect can upload the Java application and configure the environment to run Apache Tomcat.

upvoted 11 times
- KennethNg923

Most Recent

10 months, 3 weeks ago

Selected Answer: B

Beanstalk for sure

upvoted 2 times
- zinabu

1 year ago

By using Elastic Beanstalk, the solutions architect can upload the Java application and configure the environment to run Apache Tomcat.

upvoted 3 times
- wizcloudifa

1 year ago

Selected Answer: B

The key word here from the question if you notice is "The Java application must be DEPLOYED..." hence Elastic Beanstalk, it is a serverless deployment service and supports a variety of platforms(apache Tomcat in our situation), and it will scale automatically with less operational overhead(unlike option D with a lot of operation overhead)

upvoted 3 times
- awsgeek75

1 year, 3 months ago

Selected Answer: B

<https://aws.amazon.com/elasticbeanstalk/details/>

upvoted 2 times
- Guru4Cloud

1 year, 8 months ago

Selected Answer: B

B. Deploy the application by using AWS Elastic Beanstalk. Configure a load-balanced environment and a rolling deployment policy.

upvoted 4 times
- james2033

1 year, 9 months ago

Selected Answer: B

Keyword "AWS Elastic Beanstalk" for re-architecture from Java web-app inside Apache Tomcat to AWS Cloud.

upvoted 3 times
- TariqKipkemei

1 year, 11 months ago

Selected Answer: B

Definitely B



upvoted 2 times
- antropaws

1 year, 11 months ago

Selected Answer: B

Clearly B.

upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: B**

Easy deploy, management and scale  
upvoted 3 times

  **greyrose** 1 year, 11 months ago

**Selected Answer: B**

[illegible]

A serverless application uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The Lambda function needs permissions to read and write to the DynamoDB table.



Which solution will give the Lambda function access to the DynamoDB table MOST securely?

- A. Create an IAM user with programmatic access to the Lambda function. Attach a policy to the user that allows read and write access to the DynamoDB table. Store the access\_key\_id and secret\_access\_key parameters as part of the Lambda environment variables. Ensure that other AWS users do not have read and write access to the Lambda function configuration.
- B. Create an IAM role that includes Lambda as a trusted service. Attach a policy to the role that allows read and write access to the DynamoDB table. Update the configuration of the Lambda function to use the new role as the execution role.
- C. Create an IAM user with programmatic access to the Lambda function. Attach a policy to the user that allows read and write access to the DynamoDB table. Store the access\_key\_id and secret\_access\_key parameters in AWS Systems Manager Parameter Store as secure string parameters. Update the Lambda function code to retrieve the secure string parameters before connecting to the DynamoDB table.
- D. Create an IAM role that includes DynamoDB as a trusted service. Attach a policy to the role that allows read and write access from the Lambda function. Update the code of the Lambda function to attach to the new role as an execution role.

**Correct Answer: B**



*Community vote distribution*

B (100%)

-   **awsgeek75** Highly Voted 1 year, 3 months ago



**Selected Answer: B**

DynamoDB needs to trust Lambda. NOT the other way around. So Lambda must be configured as a trusted service. Role for service which gives B and D options. D is setting up (somehow?) to allow Lambda to trust DynamoDB... or the wording makes no sense.

upvoted 8 times
-   **james2033** Highly Voted 1 year, 9 months ago



**Selected Answer: B**

Keyword B. " IAM role that includes Lambda as a trusted service", not "IAM role that includes DynamoDB as a trusted service" in D. It is IAM role, not IAM user.

upvoted 6 times
-   **KennethNg923** Most Recent 10 months, 3 weeks ago



**Selected Answer: B**

IAM Role for access to DynamoDB, not for access Lambda



upvoted 2 times
-   **antropaws** 1 year, 11 months ago

**Selected Answer: B**

B sounds better.



upvoted 3 times
-   **omoakin** 1 year, 11 months ago

BBBBBBBBBB

upvoted 2 times
-   **alvinnguyennexcel** 1 year, 11 months ago



**Selected Answer: B**

vote B

upvoted 2 times
-   **cloudenthusiast** 1 year, 11 months ago

B

Option B suggests creating an IAM role that includes Lambda as a trusted service, meaning the role is specifically designed for Lambda functions. The role should have a policy attached to it that grants the required read and write access to the DynamoDB table.

upvoted 4 times
-   **nosense** 1 year, 11 months ago

**Selected Answer: B**

B is right

Role key word and trusted service lambda

upvoted 5 times

The following IAM policy is attached to an IAM group. This is the only policy applied to the group.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "1",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    },
    {
      "Sid": "2",
      "Effect": "Deny",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "BoolIfExists": {"aws:MultiFactorAuthPresent": false}
      }
    }
  ]
}
```

What are the effective IAM permissions of this policy for group members?


- A. Group members are permitted any Amazon EC2 action within the us-east-1 Region. Statements after the Allow permission are not applied.
- B. Group members are denied any Amazon EC2 permissions in the us-east-1 Region unless they are logged in with multi-factor authentication (MFA).
- C. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for all Regions when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action.
- D. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for the us-east-1 Region only when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action within the us-east-1 Region.


**Correct Answer: D**


Community vote distribution

D (90%)

10%

 **jack79** Highly Voted 1 year, 10 months ago  
came in exam today  
upvoted 13 times

 **KennethNg923** Most Recent 10 months, 3 weeks ago  
**Selected Answer: D**  
for the us-east-1 Region only, not for all region  
upvoted 2 times

 **wizcloudifa** 1 year ago  
**Selected Answer: D**  
One of the few situations when actual answer is same as the most voted answer lol  
upvoted 2 times

 **pdragon1981** 1 year, 4 months ago  
**Selected Answer: C**

Not sure why everyone vote D, I think that the valid option as to be C as the second condition regarding MFA there is point that only refer to a specific region, so basically this means that is for all the regions

upvoted 2 times

  **pdragon1981** 1 year, 4 months ago

Ok ignore D is right as the first condition is what gives permission to make anything for EC2 but is restricted to us-east-1 region  
upvoted 6 times

  **yodelin** 1 year, 6 months ago

the json is describing a lot of things apparently, so I go with the longest answer lol  
upvoted 1 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

D. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for the us-east-1 Region only when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action within the us-east-1 Region  
upvoted 3 times

  **james2033** 1 year, 9 months ago

**Selected Answer: D**

- A. "Statements after the Allow permission are not applied." --> Wrong.
- B. "denied any Amazon EC2 permissions in the us-east-1 Region" --> Wrong. Just deny 2 items.
- C. "allowed the ec2:StopInstances and ec2:TerminateInstances permissions for all Regions" --> Wrong. Just region us-east-1.
- D. ok.  
upvoted 2 times

  **TariqKipkemei** 1 year, 11 months ago

**Selected Answer: D**

Only D makes sense  
upvoted 2 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: D**

D sounds about right.  
upvoted 2 times


  **alvinnguyennexcel** 1 year, 11 months ago

**Selected Answer: D**

D is correct  
upvoted 3 times

  **omoakin** 1 year, 11 months ago

D is correct  
upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: D**

D is right  
upvoted 3 times

A manufacturing company has machine sensors that upload .csv files to an Amazon S3 bucket. These .csv files must be converted into images and must be made available as soon as possible for the automatic generation of graphical reports.

The images become irrelevant after 1 month, but the .csv files must be kept to train machine learning (ML) models twice a year. The ML trainings and audits are planned weeks in advance.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Launch an Amazon EC2 Spot Instance that downloads the .csv files every hour, generates the image files, and uploads the images to the S3 bucket.
- B. Design an AWS Lambda function that converts the .csv files into images and stores the images in the S3 bucket. Invoke the Lambda function when a .csv file is uploaded.
- C. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket. Transition the .csv files from S3 Standard to S3 Glacier 1 day after they are uploaded. Expire the image files after 30 days.
- D. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket. Transition the .csv files from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) 1 day after they are uploaded. Expire the image files after 30 days.
- E. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket. Transition the .csv files from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) 1 day after they are uploaded. Keep the image files in Reduced Redundancy Storage (RRS).

**Correct Answer: BC**

*Community vote distribution*

BC (90%)

10%

- vesen22**

Highly Voted

1 year, 5 months ago

Selected Answer: BC

<https://docs.aws.amazon.com/amazonglacier/latest/dev/introduction.html>

upvoted 5 times
- awsgeek75**

Most Recent

10 months ago

Selected Answer: BC

B for processing the images via Lambda as it's more cost efficient than EC2 spot instances  
C for expiring images after 30 days and because the ML trainings are planned weeks in advance so S3 glacier is ideal for slow retrieval and cheap storage.

D and E uses S3 infrequent access which is more expensive than glacier

upvoted 3 times
- pentium75**

10 months, 1 week ago

Selected Answer: BC

Not A, we need the images "as soon as possible", A runs every hour  
"ML trainings and audits are planned weeks in advance" thus Glacier (C) is ok.

upvoted 3 times
- Xin123**

1 year, 1 month ago

Selected Answer: BC

Answer is B&C. For D, you must store data for 30 days in s3 standard before move to IA tiers, glacier is fine

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html#:~:text=Before%20you%20transition%20objects%20to%20S3%20Standard%2DIA%20or%20S3%20One%20Zone%2DIA%2C%20you%20must%20store%20them%20for%20at%20least%2030%20days%20in%20Amazon%20S3>

upvoted 4 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: BC

Definitely B & C

upvoted 3 times
- jayce5**

1 year, 3 months ago

Selected Answer: BC

A. Wrong, the .csv files must be processed asap.  
D and E are incorrect since Glacier is the most cost-effective option, and plans for using .csv files are known weeks in advance.  
upvoted 2 times

🗨️ 👤 **james2033** 1 year, 3 months ago

Why need "These .csv files must be converted into images"?  
upvoted 1 times

🗨️ 👤 **awsgeek75** 9 months, 2 weeks ago

Because they are being used in some graphical reports (probably fancy powerpoint presentations!)  
upvoted 2 times

🗨️ 👤 **smartegnine** 1 year, 4 months ago

**Selected Answer: BC**

the key word is Weeks in advance, even you save data in S3 Gracia will also OK to take couples days to retrieve the data  
upvoted 3 times

🗨️ 👤 **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: BC**

Definitely B & C  
upvoted 2 times

🗨️ 👤 **Abrar2022** 1 year, 5 months ago

**Selected Answer: BC**

A. Wrong because Lifecycle rule is not mentioned.

B. CORRECT

C. CORRECT

D. Why Store on S3 One Zone-Infrequent Access (S3 One Zone-IA) when the files are going to irrelevant after 1 month? (Availability 99.99% - consider cost)

E. again, Why use Reduced Redundancy Storage (RRS) when the files are irrelevant after 1 month? (Availability 99.99% - consider cost)  
upvoted 4 times

🗨️ 👤 **RoroJ** 1 year, 5 months ago

**Selected Answer: BE**

B: Serverless and fast responding  
E: will keep .csv file for a year, C and D expires the file after 30 days.  
upvoted 4 times

🗨️ 👤 **RoroJ** 1 year, 5 months ago

B&C, misread the question, expires the image files after 30 days.  
upvoted 3 times

🗨️ 👤 **hiroohiroo** 1 year, 5 months ago

**Selected Answer: BC**

<https://aws.amazon.com/jp/about-aws/whats-new/2021/11/amazon-s3-glacier-storage-class-amazon-s3-glacier-flexible-retrieval/>  
upvoted 3 times

🗨️ 👤 **nosense** 1 year, 5 months ago

**Selected Answer: BC**

B severless and cost effective  
C corrctl rule to store  
upvoted 3 times

A company has developed a new video game as a web application. The application is in a three-tier architecture in a VPC with Amazon RDS for MySQL in the database layer. Several players will compete concurrently online. The game’s developers want to display a top-10 scoreboard in near-real time and offer the ability to stop and restore the game while preserving the current scores.

What should a solutions architect do to meet these requirements?

- A. Set up an Amazon ElastiCache for Memcached cluster to cache the scores for the web application to display.
- B. Set up an Amazon ElastiCache for Redis cluster to compute and cache the scores for the web application to display.
- C. Place an Amazon CloudFront distribution in front of the web application to cache the scoreboard in a section of the application.
- D. Create a read replica on Amazon RDS for MySQL to run queries to compute the scoreboard and serve the read traffic to the web application.

**Correct Answer: B**

*Community vote distribution*

B (92%)

8%

- Guru4Cloud**

Highly Voted

1 year, 2 months ago

Selected Answer: B

Redis provides fast in-memory data storage and processing. It can compute the top 10 scores and update the cache in milliseconds. ElastiCache Redis supports sorting and ranking operations needed for the top 10 leaderboard. The cached leaderboard can be retrieved from Redis vs hitting the MySQL database for every read. This reduces load on the database. Redis supports persistence, so scores are preserved if the cache stops/restarts

upvoted 10 times
- TariqKipkemei**

Highly Voted

1 year ago

Selected Answer: B

Real-time gaming leaderboards are easy to create with Amazon ElastiCache for Redis. Just use the Redis Sorted Set data structure, which provides uniqueness of elements while maintaining the list sorted by their scores. Creating a real-time ranked list is as simple as updating a user's score each time it changes. You can also use Sorted Sets to handle time series data by using timestamps as the score.

<https://aws.amazon.com/elasticache/redis/#:~:text=ElastiCache%20for%20Redis,-Gaming,-Leaderboards>

upvoted 8 times
- potomac**

Most Recent

12 months ago

Selected Answer: B

ElastiCache for Redis sorts and ranks datasets

upvoted 5 times
- 5ab5e39**

1 year, 1 month ago

<https://aws.amazon.com/blogs/database/building-a-real-time-gaming-leaderboard-with-amazon-elasticache-for-redis/>

upvoted 4 times
- ukivanlamlpi**

1 year, 2 months ago

Selected Answer: A

concurrently = memcached

upvoted 3 times
- awsgeek75**

10 months ago

Amazon ElastiCache for Memcached is non-persistent so start/stop of game will lose the score." and offer the ability to stop and restore the game while preserving the current scores."

upvoted 3 times
- james2033**

1 year, 3 months ago

Selected Answer: B

See case study of leaderboard with Redis at <https://redis.io/docs/data-types/sorted-sets/> , it is feature "sorted sets". See comparison between Redis and Memcached at <https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/SelectEngine.html> , the different at feature "Sorted sets"

upvoted 4 times
- live\_reply\_developers**

1 year, 4 months ago



Selected Answer: B



advanced data structures, complex querying, pub/sub messaging, or persistence, Redis may be a better fit.



upvoted 2 times
- haoAWS**



1 year, 4 months ago



B is correct  
upvoted 2 times

  **jf\_topics** 1 year, 4 months ago  
B correct.  
upvoted 2 times

  **hiroohiroo** 1 year, 5 months ago  
**Selected Answer: B**  
<https://aws.amazon.com/jp/blogs/news/building-a-real-time-gaming-leaderboard-with-amazon-elasticache-for-redis/>  
upvoted 4 times

  **cloudenthusiast** 1 year, 5 months ago  
Amazon ElastiCache for Redis is a highly scalable and fully managed in-memory data store. It can be used to store and compute the scores in real time for the top-10 scoreboard. Redis supports sorted sets, which can be used to store the scores as well as perform efficient queries to retrieve the top scores. By utilizing ElastiCache for Redis, the web application can quickly retrieve the current scores without the need to perform complex and potentially resource-intensive database queries.  
upvoted 3 times

  **nosense** 1 year, 5 months ago  
**Selected Answer: B**  
B is right  
upvoted 2 times

  **Efren** 1 year, 5 months ago  
More questions!!!  
upvoted 5 times

An ecommerce company wants to use machine learning (ML) algorithms to build and train models. The company will use the models to visualize complex scenarios and to detect trends in customer data. The architecture team wants to integrate its ML models with a reporting platform to analyze the augmented data and use the data directly in its business intelligence dashboards.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Glue to create an ML transform to build and train models. Use Amazon OpenSearch Service to visualize the data.
- B. Use Amazon SageMaker to build and train models. Use Amazon QuickSight to visualize the data.
- C. Use a pre-built ML Amazon Machine Image (AMI) from the AWS Marketplace to build and train models. Use Amazon OpenSearch Service to visualize the data.
- D. Use Amazon QuickSight to build and train models by using calculated fields. Use Amazon QuickSight to visualize the data.

**Correct Answer: B**

*Community vote distribution*

B (100%)

  **vincent2023** 3 months, 2 weeks ago

**Selected Answer: B**

Sage Maker - ML  
upvoted 1 times

  **67db0ed** 9 months, 1 week ago

<https://docs.aws.amazon.com/quicksight/latest/user/sagemaker-integration.html>  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: B**

Machine Learning = Sage Maker so B for least operational overhead  
A and D are not right technologies.  
C is possible but with more overhead of using AMI even if you can get OpenSearch to visualize the data somehow which I don't think is possible without massive overhead  
upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Use Amazon SageMaker to build and train models. Use Amazon QuickSight to visualize the data.  
upvoted 3 times

  **james2033** 1 year, 9 months ago

**Selected Answer: B**

Question keyword "machine learning", answer keyword "Amazon SageMaker". Choose B. Use Amazon QuickSight for visualization. See "Gaining insights with machine learning (ML) in Amazon QuickSight" at <https://docs.aws.amazon.com/quicksight/latest/user/making-data-driven-decisions-with-ml-in-quicksight.html>  
upvoted 3 times

  **VellaDevil** 1 year, 10 months ago

**Selected Answer: B**

Sagemaker.  
upvoted 2 times

  **TariqKipkemei** 1 year, 11 months ago


**Selected Answer: B**

Business intelligence, visualiations = AmazonQuicksight  
ML = Amazon SageMaker  
upvoted 3 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: B**

Most likely B.  
upvoted 2 times

  **omoakin** 1 year, 11 months ago

Amazon SageMaker is a fully managed service that provides every developer and data scientist with the ability to build, train, and deploy ML models quickly.



upvoted 2 times

  **cloudenthusiast** 1 year, 11 months ago

Amazon SageMaker is a fully managed service that provides a complete set of tools and capabilities for building, training, and deploying ML models. It simplifies the end-to-end ML workflow and reduces operational overhead by handling infrastructure provisioning, model training, and deployment.

To visualize the data and integrate it into business intelligence dashboards, Amazon QuickSight can be used. QuickSight is a cloud-native business intelligence service that allows users to easily create interactive visualizations, reports, and dashboards from various data sources, including the augmented data generated by the ML models.



upvoted 4 times

  **Efren** 1 year, 11 months ago

**Selected Answer: B**

ML= = SageMaker

upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: B**

B sagemaker provide deploy ml models

upvoted 2 times

A company is running its production and nonproduction environment workloads in multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to design a solution that will prevent the modification of cost usage tags.

Which solution will meet these requirements?

- A. Create a custom AWS Config rule to prevent tag modification except by authorized principals.
- B. Create a custom trail in AWS CloudTrail to prevent tag modification.
- C. Create a service control policy (SCP) to prevent tag modification except by authorized principals.
- D. Create custom Amazon CloudWatch logs to prevent tag modification.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Guru4Cloud** Highly Voted  1 year, 2 months ago

**Selected Answer: C**

Tip: AWS Organziaton + service control policy (SCP) - This for any questions, you see both together. then you tell me C. Create a service control policy (SCP) to prevent tag modification except by authorized principals.

upvoted 7 times

  **awsgeek75** Most Recent  10 months ago

**Selected Answer: C**

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html)

AWS example for this question/use case:

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps\\_examples\\_tagging.html#example-require-restrict-tag-mods-to-admin](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.html#example-require-restrict-tag-mods-to-admin)

upvoted 3 times

  **james2033** 1 year, 3 months ago

**Selected Answer: C**

D "Amazon CloudWatch" just for logging, not for prevent tag modification

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_tag-policies-cwe.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_tag-policies-cwe.html)

Amazon Organziaton has "Service Control Policy (SCP)" with "tag policy"

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_tag-policies.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_tag-policies.html) . Choose C.

AWS Config for technical stuff, not for tag policies. Not A.

upvoted 4 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: C**

Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization.

upvoted 2 times

  **alexandercamachop** 1 year, 5 months ago

**Selected Answer: C**

Anytime we need to restrict anything in an AWS Organization, it is SCP Policies.

upvoted 3 times

  **Abrar2022** 1 year, 5 months ago

AWS Config is for tracking configuration changes

upvoted 2 times

  **Abrar2022** 1 year, 5 months ago

so it's wrong. Right asnwer is C

upvoted 3 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: C**

I'd say C.

upvoted 3 times

  **hiroohiroo** 1 year, 5 months ago

**Selected Answer: C**

[https://docs.aws.amazon.com/ja\\_jp/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps\\_examples\\_tagging.html](https://docs.aws.amazon.com/ja_jp/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.html)  
upvoted 4 times

  **nosense** 1 year, 5 months ago

**Selected Answer: C**

Denies tag: modify  
upvoted 3 times

A company hosts its application in the AWS Cloud. The application runs on Amazon EC2 instances behind an Elastic Load Balancer in an Auto Scaling group and with an Amazon DynamoDB table. The company wants to ensure the application can be made available in another AWS Region with minimal downtime.

What should a solutions architect do to meet these requirements with the LEAST amount of downtime?

- A. Create an Auto Scaling group and a load balancer in the disaster recovery Region. Configure the DynamoDB table as a global table. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- B. Create an AWS CloudFormation template to create EC2 instances, load balancers, and DynamoDB tables to be launched when needed. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- C. Create an AWS CloudFormation template to create EC2 instances and a load balancer to be launched when needed. Configure the DynamoDB table as a global table. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- D. Create an Auto Scaling group and load balancer in the disaster recovery Region. Configure the DynamoDB table as a global table. Create an Amazon CloudWatch alarm to trigger an AWS Lambda function that updates Amazon Route 53 pointing to the disaster recovery load balancer.

**Correct Answer: A**

*Community vote distribution*

A (57%)

C (38%)

5%

- lucdt4**

Highly Voted

1 year, 11 months ago

Selected Answer: A

A and D is correct.  
But Route 53 has a feature DNS failover when instances down so we dont need use Cloudwatch and lambda to trigger  
-> A correct  
upvoted 16 times
- Wablo**

1 year, 10 months ago

Yes it does but you configure it. Its not automated anymore. D is the best answer!  
upvoted 2 times
- JA2018**

5 months, 1 week ago

to be launched when needed => requires manual human intervention  
  
Certainly will not help with requirement for the least amount of downtime.  
upvoted 1 times
- Kp88**

1 year, 9 months ago

What are you talking about configuring ? Yes you have to configure everything at some point  
<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-configuring.html>  
upvoted 2 times
- smartegnine**

1 year, 10 months ago

Did not see Route 53 in this question right? So my opinion is D  
upvoted 1 times
- pentium75**

Highly Voted

1 year, 4 months ago

Selected Answer: C

They are not asking for automatic failover, they want to "ensure the application can (!) be made available in another AWS Region with minimal downtime". This works with C; they would just execute the template and it would be available in short time.  
  
A would create a DR environment that IS already available, which is not what the question asks for.  
D is like A, just abusing Lambda to update the DNS record (which doesn't make sense).  
B would create a separate, empty database  
upvoted 9 times
- bora4motion**

3 days, 22 hours ago

I am going with A. There's no downtime with A, maybe a blip because of DNS.  
upvoted 1 times
- [Removed]**

Most Recent

9 months, 1 week ago

ChatGPT:

Option C involves creating an AWS CloudFormation template to create EC2 instances and a load balancer only when needed, and configuring the DynamoDB table as a global table. This approach might introduce more downtime because the infrastructure in the disaster recovery region is not pre-deployed and ready to take over immediately. The process of launching instances and configuring the load balancer can take some time, leading to delays during the failover.

Option A, on the other hand, ensures that the necessary infrastructure (Auto Scaling group, load balancer, and DynamoDB global table) is already set up and running in the disaster recovery region. This pre-deployment reduces downtime since the failover can be handled quickly by updating DNS to point to the disaster recovery region's load balancer.

upvoted 2 times

  **anikolov** 1 year, 3 months ago

**Selected Answer: A**

With the LEAST amount of downtime = A  
Cost effective = C , but risky some of EC2 types/capacity not available in Region at the time, when need to switch to DR


upvoted 7 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: C**

There are 2 parts. DB and application. Dynamo DB recovery in another region is not possible without global table so option B is out.  
A will make the infra available in 2 regions which is not required. The question is about DR, not scaling.  
D Use Lambda to modify R53 to point to new region. This is going to cause delays but is possible and it will also be running a scaled EC2 instances in passive region.  
C Make a CF template which can launch the infra when needed. DB is global table so it will be available.

upvoted 4 times

  **bora4motion** 3 days, 22 hours ago

you need minimum downtime, nobody asks about what's required or not apart of minimal downtime. I'm going with a.

upvoted 1 times

  **LeonSauveterre** 5 months ago

Option C relies on launching EC2 instances and load balancers only when needed, which may increase downtime. On contrary, I think A is better because it ensures resources are pre-configured and ready-to-go.

It's true nobody mentioned anything about infra availability in 2 regions... but it couldn't hurt since we simply need the least amount of downtime.

upvoted 1 times

  **meowruki** 1 year, 5 months ago

**Selected Answer: C**

AWS CloudFormation Template: Use CloudFormation to define the infrastructure components (EC2 instances, load balancer, etc.) in a template. This allows for consistent and repeatable infrastructure deployment.

EC2 Instances and Load Balancer: Launch the EC2 instances and load balancer in the disaster recovery (DR) Region using the CloudFormation template. This enables the deployment of the application in the DR Region when needed.

DynamoDB Global Table: Configure the DynamoDB table as a global table. DynamoDB Global Tables provide automatic multi-region, multi-master replication, ensuring that the data is available in both the primary and DR Regions.

DNS Failover: Configure DNS failover to point to the new DR Region's load balancer. This allows for seamless failover of traffic to the DR Region when needed.

Option A is close, but it introduces an Auto Scaling group in the disaster recovery Region, which might introduce unnecessary complexity and potential scaling delays. Option D introduces a Lambda function triggered by CloudWatch alarms, which might add latency and complexity compared to the more direct approach in Option C.



upvoted 1 times

  **bogobob** 1 year, 5 months ago

**Selected Answer: A**

Assuming theyre using Route53 as a DNS then A <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

upvoted 3 times

  **EEK2k** 1 year, 5 months ago

**Selected Answer: C**

Only B and C take care of EC2 instants. But since B does not take care of Data in the Dynamo DB, C is the only correct Answer.


upvoted 1 times

  **potomac** 1 year, 6 months ago

**Selected Answer: A**

Route 53 haves a feature DNS failover when instances down

upvoted 1 times

  **thanhnv142** 1 year, 6 months ago

C is the best choice here

upvoted 1 times

  **Wayne23Fang** 1 year, 6 months ago



**Selected Answer: C**

I think CloudFormation is easier than manual provision of Auto Scaling group and load balancer in DR region.  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

Creating Auto Scaling group and load balancer in DR region allows fast launch of capacity when needed.  
Configuring DynamoDB as a global table provides continuous data replication.  
Using DNS failover via Route 53 to point to the DR region's load balancer enables rapid traffic shifting.  
upvoted 2 times



  **Wablo** 1 year, 10 months ago

Both Option A and Option D include the necessary steps of setting up an Auto Scaling group and load balancer in the disaster recovery Region, configuring the DynamoDB table as a global table, and updating DNS records. However, Option D provides a more detailed approach by explicitly mentioning the use of an Amazon CloudWatch alarm and AWS Lambda function to automate the DNS update process.

By leveraging an Amazon CloudWatch alarm, Option D allows for an automated failover mechanism. When triggered, the CloudWatch alarm can execute an AWS Lambda function, which in turn can update the DNS records in Amazon Route 53 to redirect traffic to the disaster recovery load balancer in the new Region. This automation helps reduce the potential for human error and further minimizes downtime.

Answer is D

upvoted 2 times

  **Kp88** 1 year, 9 months ago

Failover policy takes care of DNS record update so no need for cloud watch/lambda  
upvoted 1 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: C**

The company wants to ensure the application 'CAN' be made available in another AWS Region with minimal downtime. Meaning they want to be able to launch infra on need basis.

Best answer is C.

upvoted 2 times

  **Wablo** 1 year, 10 months ago

minimal downtme not minimal effort!

D

upvoted 1 times

  **dajform** 1 year, 10 months ago

B, C are not OK because "launching resources when needed", which will increase the time to recover "DR"


upvoted 2 times

  **AshishRocks** 1 year, 11 months ago

I feel it is A

Configure DNS failover: Use DNS failover to point the application's DNS record to the load balancer in the disaster recovery Region. DNS failover allows you to route traffic to the disaster recovery Region in case of a failure in the primary Region.

upvoted 3 times

  **Wablo** 1 year, 10 months ago

Once you configure manually the DNS , its no more automated like Lambda does.

upvoted 2 times

  **Yadav\_Sanjay** 1 year, 11 months ago

**Selected Answer: C**

C suits best

upvoted 3 times

  **hiroohiroo** 1 year, 11 months ago

**Selected Answer: A**

AがDNS フェイルオーバー

upvoted 1 times

A company needs to migrate a MySQL database from its on-premises data center to AWS within 2 weeks. The database is 20 TB in size. The company wants to complete the migration with minimal downtime.

Which solution will migrate the database MOST cost-effectively?

- A. Order an AWS Snowball Edge Storage Optimized device. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing changes. Send the Snowball Edge device to AWS to finish the migration and continue the ongoing replication.
- B. Order an AWS Snowmobile vehicle. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing changes. Send the Snowmobile vehicle back to AWS to finish the migration and continue the ongoing replication.
- C. Order an AWS Snowball Edge Compute Optimized with GPU device. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing changes. Send the Snowball device to AWS to finish the migration and continue the ongoing replication
- D. Order a 1 GB dedicated AWS Direct Connect connection to establish a connection with the data center. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing changes.

**Correct Answer: A**

*Community vote distribution*

A (88%)

12%

- nosense**

Highly Voted

1 year, 5 months ago

Selected Answer: A

A) 300 first 10 days. 150 shipping  
D) 750 for 2 weeks  
upvoted 12 times
- Efren**

1 year, 5 months ago

Thanks, i was checking the speed more than price. Thanks for the clarification  
upvoted 3 times
- Goutham4981**

Highly Voted

11 months, 3 weeks ago

Selected Answer: A

Direct Connect takes at least 1 month to setup - D is invalid  
AWS Snowmobile is used for transferring large amounts of data (petabytes) from remote locations where establishing a connection to the cloud is impossible - B is invalid  
AWS Snowball Edge Compute Optimized provides higher vCPU performance and lower storage as compared to Snowball storage optimized. As our need is solely data transfer, high vCPU performance is not required but high storage is - C is invalid  
upvoted 7 times
- zinabu**

Most Recent

6 months, 2 weeks ago

Selected Answer: A

But I didn't understand why we are using a schema conversion tool, because AWS have already a managed service engine for MySQL Db, (RDS for MySQL or Aurora for my SQL is on the table )  
upvoted 6 times
- EEK2k**

11 months, 3 weeks ago

Selected Answer: D

To calculate the time it would take to transfer 20TB of data over a 1 GB dedicated AWS Direct Connect, we can use the formula:  
  
time = data size / data transfer rate  
  
Here, the data size is 20TB, which is equivalent to 20,000 GB or 20,000,000 MB. The data transfer rate is 1 GB/s.  
  
Converting the data size to MB, we get:  
  
20,000,000 MB / 1 GB/s = 20,000 seconds  
  
Therefore, it would take approximately 20,000 seconds or 5.56 hours to transfer 20TB of data over a 1 GB dedicated AWS Direct Connect.  
upvoted 3 times
- Murtadhaceit**

11 months ago

It takes more way more than 2 weeks to setup Direct Connect. Therefore, D is not valid since we have to do the transfer within 2 weeks.  
upvoted 3 times

  **potomac** 12 months ago

**Selected Answer: A**

C is wrong, GPU is not needed  
upvoted 4 times

  **Ramdi1** 1 year, 1 month ago

**Selected Answer: A**

Has to be A. the option for D would only work if they said they have like 6 Months plus. It would take too long to set up.  
upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

I agreed with A.  
Why not D?  
When you initiate the process by requesting an AWS Direct Connect connection, it typically starts with the AWS Direct Connect provider. This provider may need to coordinate with AWS to allocate the necessary resources. This initial setup phase can take anywhere from a few days to a couple of weeks.  
Couple of weeks? No Good  
upvoted 4 times



  **Guru4Cloud** 1 year, 2 months ago

When you create a Snowball job in the AWS console, it will estimate the delivery date based on your location. Being near a facility shows 1-2 day estimated delivery.  
For extremely urgent requests, you can contact AWS Support and inquire about expedited Snowball delivery. If inventory is available, they may be able to ship same day or next day.  
upvoted 3 times

  **james2033** 1 year, 3 months ago

**Selected Answer: A**

Keyword "20 TB", choose "AWS Snowball", there are A or C. C has word "GPU" what is not related, therefore choose A.  
upvoted 3 times

  **Zox42** 1 year, 3 months ago

**Selected Answer: A**

Answer A  
upvoted 2 times

  **MrAWSAssociate** 1 year, 4 months ago

**Selected Answer: D**

D is correct  
upvoted 1 times



  **pentium75** 10 months, 1 week ago

No, takes months, not weeks  
upvoted 2 times

  **DrWatson** 1 year, 5 months ago

**Selected Answer: A**

[https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_LargeDBs.Process.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_LargeDBs.Process.html)  
upvoted 3 times

  **RoroJ** 1 year, 5 months ago



**Selected Answer: A**

D Direct Connection will need a long time to setup plus need to deal with Network and Security changes with existing environment. Ad then plus the Data trans time... No way can be done in 2 weeks.  
upvoted 5 times

  **Joselucho38** 1 year, 5 months ago



**Selected Answer: D**

Overall, option D combines the reliability and cost-effectiveness of AWS Direct Connect, AWS DMS, and AWS SCT to migrate the database efficiently and minimize downtime.  
upvoted 2 times

  **Abhineet9148232** 1 year, 5 months ago



**Selected Answer: A**

D - Direct Connect takes atleast a month to setup! Requirement is for within 2 weeks.  
upvoted 5 times



  **Rob1L** 1 year, 5 months ago

**Selected Answer: D**

AWS Snowball Edge Storage Optimized device is used for large-scale data transfers, but the lead time for delivery, data transfer, and return shipping would likely exceed the 2-week time frame. Also, ongoing database changes wouldn't be replicated while the device is in transit.  
upvoted 1 times



  **Rob1L** 1 year, 5 months ago

Change to A because "Most cost effective"  
upvoted 3 times

  **hiroohiroo** 1 year, 5 months ago

**Selected Answer: A**

[https://docs.aws.amazon.com/ja\\_jp/snowball/latest/developer-guide/device-differences.html#device-options](https://docs.aws.amazon.com/ja_jp/snowball/latest/developer-guide/device-differences.html#device-options)  
Aです。  
upvoted 4 times

  **norris81** 1 year, 5 months ago

**Selected Answer: A**

How long does direct connect take to provision ?  
upvoted 2 times

  **examtopictempacc** 1 year, 5 months ago

At least one month and expensive.  
upvoted 1 times

A company moved its on-premises PostgreSQL database to an Amazon RDS for PostgreSQL DB instance. The company successfully launched a new product. The workload on the database has increased. The company wants to accommodate the larger workload without adding infrastructure.

Which solution will meet these requirements MOST cost-effectively?

- A. Buy reserved DB instances for the total workload. Make the Amazon RDS for PostgreSQL DB instance larger.
- B. Make the Amazon RDS for PostgreSQL DB instance a Multi-AZ DB instance.
- C. Buy reserved DB instances for the total workload. Add another Amazon RDS for PostgreSQL DB instance.
- D. Make the Amazon RDS for PostgreSQL DB instance an on-demand DB instance.

**Correct Answer: A**

*Community vote distribution*

A (84%)

Other

- elmogy**

Highly Voted

1 year, 11 months ago

Selected Answer: A

A.  
"without adding infrastructure" means scaling vertically and choosing larger instance.  
"MOST cost-effectively" reserved instances  
upvoted 16 times
- wsdasdasdqwdaw**

1 year, 6 months ago

"MOST cost-effectively" doesn't mean reserved instances. Only in this case it is but not in general.  
upvoted 6 times
- zdi561**

Most Recent

3 months, 2 weeks ago

Selected Answer: B

without adding infrastructure means without adding more capacity per instance, multi-az adding read replicas solve the problem.  
upvoted 1 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: A

accommodate the larger workload without adding infrastructure. = Reserved DB instance  
upvoted 3 times
- awsgeek75**

1 year, 3 months ago

+ make the instance larger so most cost effective is to reserve a large instance suitable for workload which is A  
upvoted 3 times
- pentium75**

1 year, 4 months ago

Selected Answer: A

B - Multi-AZ is for HA, does not help 'accommodating the larger workload'  
C - Adding "another instance" will not help, we can't split the workload between two instances  
D - On-demand instance is a good choice for unknown workload, but here we know the workload, it's just higher than before  
upvoted 3 times
- Goutham4981**

1 year, 5 months ago

Selected Answer: A

Cannot add more infrastructure - C is invalid  
Multi AZ DB instance is for high availability and failure mitigation, does not increase performance, higher workload support - B is invalid  
On demand instances are costlier than Reserved instances - D is invalid  
upvoted 2 times
- bogobob**

1 year, 5 months ago

Selected Answer: D

Not A : "launched a new product", reserved instances are for known workloads, a new product doesn't have known workload.  
Not B : "accommodate the larger workload", while Multi-AZ can help with larger workloads, they are more for higher availability.  
Not C : "without adding infrastructure", adding a PostGresQL instance is new infrastructure.  
upvoted 3 times
- JA2018**

5 months, 1 week ago

Please refer to following Keys in Stem:

1. The company successfully launched a new product.
  2. The workload on the database has increased.
  3. The company wants to accommodate the larger workload without adding infrastructure.
- upvoted 1 times

  **pentium75** 1 year, 4 months ago

Question says nothing about unknown load. New product -> more total products -> load has increased.

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

B is the best approach in this scenario overall:

Making the RDS PostgreSQL instance Multi-AZ adds a standby replica to handle larger workloads and provides high availability. Even though it adds infrastructure, the cost is less than doubling the infrastructure with a separate DB instance. It provides better performance, availability, and disaster recovery than a single larger instance.

upvoted 2 times

  **BillyBlunts** 1 year, 7 months ago

Agreed the answer is B  
Multi-AZ deployments are cost-effective because they leverage the standby instance without incurring additional charges. You only pay for the primary instance's regular usage costs.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

Multi-AZ is for HA, does not add performance. Meaning, will not help 'accommodating the larger workload'.

upvoted 3 times

  **james2033** 1 year, 9 months ago

**Selected Answer: A**

Buy larger instance.

upvoted 2 times

  **james2033** 1 year, 9 months ago

**Selected Answer: A**

Keyword "Amazon RDS for PostgreSQL instance large" . See list of size of instance at <https://aws.amazon.com/rds/instance-types/>



upvoted 2 times

  **examtopictempacc** 1 year, 11 months ago

**Selected Answer: A**

A.  
Not C: without adding infrastructure

upvoted 3 times

  **EA100** 1 year, 11 months ago

Answer - C  
Option B, making the Amazon RDS for PostgreSQL DB instance a Multi-AZ DB instance, would provide high availability and fault tolerance but may not directly address the need for increased capacity to handle the larger workload.



Therefore, the recommended solution is Option C: Buy reserved DB instances for the workload and add another Amazon RDS for PostgreSQL DB instance to accommodate the increased workload in a cost-effective manner.

upvoted 1 times

  **cloudenthusiast** 1 year, 11 months ago



C  
Option C: buying reserved DB instances for the total workload and adding another Amazon RDS for PostgreSQL DB instance seems to be the most appropriate choice. It allows for workload distribution across multiple instances, providing scalability and potential performance improvements. Additionally, reserved instances can provide cost savings in the long term.

upvoted 1 times

  **nosense** 1 year, 11 months ago



A for me, because without adding additional infrastructure

upvoted 4 times

  **th3k33n** 1 year, 11 months ago

Should be C

upvoted 1 times

  **Efren** 1 year, 11 months ago



That would add more infrastructure. A would increase the size, keeping the number of instances, i think

upvoted 2 times

  **cloudenthusiast** 1 year, 11 months ago



Option A involves making the existing Amazon RDS for PostgreSQL DB instance larger. While this can improve performance, it may not be sufficient to handle a significantly increased workload. It also doesn't distribute the workload or provide scalability.

upvoted 1 times

  **nosense** 1 year, 11 months ago

The main not HA, cost-effectively and without adding infrastructure

upvoted 2 times

  **omoakin** 1 year, 11 months ago

A is the best

upvoted 2 times

A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users.

What should a solutions architect recommend?

- A. Deploy Amazon Inspector and associate it with the ALB.
- B. Deploy AWS WAF, associate it with the ALB, and configure a rate-limiting rule.
- C. Deploy rules to the network ACLs associated with the ALB to block the incoming traffic.
- D. Deploy Amazon GuardDuty and enable rate-limiting protection when configuring GuardDuty.

**Correct Answer: B**

*Community vote distribution*

B (96%)

2%

- samehpalass

Highly Voted

1 year, 4 months ago

Selected Answer: B

As no shield protect here so WAF rate limit  
upvoted 11 times
- hydro143

1 year ago

Where's your Shield Advanced now, in your hour of need he has abandoned you  
upvoted 11 times
- pentium75

Highly Voted

10 months, 1 week ago

Selected Answer: B

Best solution Shield Advanced, not listed here, thus second-best solution, WAF with rate limiting  
upvoted 8 times
- awsgeek75

Most Recent

10 months ago

Selected Answer: B

A. Amazon Inspector = Software vulnerabilities like OS patches etc. Not fit for purpose.  
C. Changing IP from DDoS so don't know the incoming traffic for configuration (even if it was possible)  
D. GardDuty is for workload and AWS account monitoring so it can't help with DDoS.  
  
B is correct as AWS WAF + ALB can configure rate limiting even if source IP changes.  
upvoted 6 times
- jAtlas7

10 months, 1 week ago

Selected Answer: B

according to some google searches... to protect against DDOS attack:  
\* AWS WAF(Web Application Firewall) provides protection on the application layer (I think Application Load Balancer belongs to this level)  
\* AWS Shield protects the infrastructure layers of the OSI mode (I think AWS Network Load Balancer belongs to this level)  
upvoted 3 times
- Guru4Cloud

1 year, 2 months ago

Selected Answer: A

This case is A  
upvoted 1 times
- vincent2023

3 months, 2 weeks ago

Inspector wouldn't help in this case as it's used to detect vulnerabilities.  
upvoted 1 times
- pentium75

10 months, 1 week ago

Inspector is for detecting vulnerabilities, has nothing to do with the requirement.  
upvoted 3 times
- james2033

1 year, 3 months ago

Selected Answer: B

AWS Web Application Firewall (WAF) + ALB (Application Load Balancer) See image at <https://aws.amazon.com/waf/> .  
<https://docs.aws.amazon.com/waf/latest/developerguide/ddos-responding.html> .

Question keyword "high request rate", answer keyword "rate-limiting rule" <https://docs.aws.amazon.com/waf/latest/developerguide/waf-rate-based-example-limit-login-page-keys.html>

Amazon GuardDuty for threat detection <https://aws.amazon.com/guardduty/> , not for DDoS.  
upvoted 3 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: B**

B in swahili 'ba' :)  
external systems, incoming requests = AWS WAF  
upvoted 2 times

  **Axeashes** 1 year, 4 months ago

**Selected Answer: B**

layer 7 DDoS protection with WAF  
<https://docs.aws.amazon.com/waf/latest/developerguide/ddos-get-started-web-acl-rbr.html>  
upvoted 2 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: B**

B no doubt.  
upvoted 1 times

  **Joselucho38** 1 year, 5 months ago



**Selected Answer: B**

AWS WAF (Web Application Firewall) is a service that provides protection for web applications against common web exploits. By associating AWS WAF with the Application Load Balancer (ALB), you can inspect incoming traffic and define rules to allow or block requests based on various criteria.  
upvoted 5 times

  **cloudenthusiast** 1 year, 5 months ago

B  
AWS Web Application Firewall (WAF) is a service that helps protect web applications from common web exploits and provides advanced security features. By deploying AWS WAF and associating it with the ALB, the company can set up rules to filter and block incoming requests based on specific criteria, such as IP addresses.

In this scenario, the company is facing performance issues due to a high request rate from illegitimate external systems with changing IP addresses. By configuring a rate-limiting rule in AWS WAF, the company can restrict the number of requests coming from each IP address, preventing excessive traffic from overwhelming the website. This will help mitigate the impact of potential DDoS attacks and ensure that legitimate users can access the site without interruption.  
upvoted 5 times

  **Efren** 1 year, 5 months ago



**Selected Answer: B**

If not AWS Shield, then WAF  
upvoted 4 times

  **nosense** 1 year, 5 months ago

**Selected Answer: B**

B obv for this  
upvoted 4 times



  **Efren** 1 year, 5 months ago

My mind slipped with AWS Shield. GuardDuty can be working along with WAF for DDOS attack, but ultimately would be WAF

<https://aws.amazon.com/blogs/security/how-to-use-amazon-guardduty-and-aws-web-application-firewall-to-automatically-block-suspicious-hosts/>  
upvoted 2 times

  **Mia2009687** 1 year, 3 months ago

Same here, I was looking for AWS Shield  
upvoted 1 times

  **Efren** 1 year, 5 months ago

**Selected Answer: D**

D, Guard Duty for me  
upvoted 1 times

  **pentium75** 10 months, 1 week ago

Guard Duty detects threats, has nothing to do with rate-limiting.  
upvoted 2 times



A company wants to share accounting data with an external auditor. The data is stored in an Amazon RDS DB instance that resides in a private subnet. The auditor has its own AWS account and requires its own copy of the database.

What is the MOST secure way for the company to share the database with the auditor?

- A. Create a read replica of the database. Configure IAM standard database authentication to grant the auditor access.
- B. Export the database contents to text files. Store the files in an Amazon S3 bucket. Create a new IAM user for the auditor. Grant the user access to the S3 bucket.
- C. Copy a snapshot of the database to an Amazon S3 bucket. Create an IAM user. Share the user's keys with the auditor to grant access to the object in the S3 bucket.
- D. Create an encrypted snapshot of the database. Share the snapshot with the auditor. Allow access to the AWS Key Management Service (AWS KMS) encryption key.

**Correct Answer: D**

*Community vote distribution*

D (100%)

  **alexandercamachop** Highly Voted 1 year, 11 months ago

**Selected Answer: D**

The most secure way for the company to share the database with the auditor is option D: Create an encrypted snapshot of the database, share the snapshot with the auditor, and allow access to the AWS Key Management Service (AWS KMS) encryption key.

By creating an encrypted snapshot, the company ensures that the database data is protected at rest. Sharing the encrypted snapshot with the auditor allows them to have their own copy of the database securely.

In addition, granting access to the AWS KMS encryption key ensures that the auditor has the necessary permissions to decrypt and access the encrypted snapshot. This allows the auditor to restore the snapshot and access the data securely.

This approach provides both data protection and access control, ensuring that the database is securely shared with the auditor while maintaining the confidentiality and integrity of the data.

upvoted 23 times

  **TariqKipkemei** 1 year, 10 months ago

best explanation ever

upvoted 4 times

  **Guru4Cloud** Highly Voted 1 year, 8 months ago

**Selected Answer: D**

Key word: "Secure way"

The snapshot contents are encrypted using KMS keys for data security.

Sharing the snapshot directly removes risks of extracting/transferring data.

The auditor can restore the snapshot into their own RDS instance.

Access is controlled through sharing the encrypted snapshot and KMS key.

upvoted 5 times

  **24b2e9e** Most Recent 10 months, 2 weeks ago

why not A ?

upvoted 5 times

  **LeonSauveterre** 5 months ago

I think it's simply because it introduces more risks than sharing just a snapshot. External entities don't need that many privileges, so it's not the MOST secure way


upvoted 1 times

  **KennethNg923** 10 months, 3 weeks ago

**Selected Answer: D**

Encrypted snapshot must be most secure compare others

upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: D**

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_ShareSnapshot.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html)

With AWS RDS, you can share snapshots across accounts so no need to go through S3 or replication. Option D allows more secure way by using encryption and sharing encryption key.  
upvoted 3 times

  **potomac** 1 year, 6 months ago

**Selected Answer: D**

MOST secure way  
upvoted 3 times


  **antropaws** 1 year, 11 months ago

**Selected Answer: D**

Most likely D.  
upvoted 4 times

  **cloudenthusiast** 1 year, 11 months ago

Option D (Creating an encrypted snapshot of the database, sharing the snapshot, and allowing access to the AWS Key Management Service encryption key) is generally considered a better option for sharing the database with the auditor in terms of security and control.  
upvoted 3 times

  **nosense** 1 year, 11 months ago

**Selected Answer: D**

D for me  
upvoted 3 times

A solutions architect configured a VPC that has a small range of IP addresses. The number of Amazon EC2 instances that are in the VPC is increasing, and there is an insufficient number of IP addresses for future workloads.

Which solution resolves this issue with the LEAST operational overhead?

- A. Add an additional IPv4 CIDR block to increase the number of IP addresses and create additional subnets in the VPC. Create new resources in the new subnets by using the new CIDR.
- B. Create a second VPC with additional subnets. Use a peering connection to connect the second VPC with the first VPC. Update the routes and create new resources in the subnets of the second VPC.
- C. Use AWS Transit Gateway to add a transit gateway and connect a second VPC with the first VPC. Update the routes of the transit gateway and VPCs. Create new resources in the subnets of the second VPC.
- D. Create a second VPC. Create a Site-to-Site VPN connection between the first VPC and the second VPC by using a VPN-hosted solution on Amazon EC2 and a virtual private gateway. Update the route between VPCs to the traffic through the VPN. Create new resources in the subnets of the second VPC.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **antropaws** Highly Voted 1 year, 11 months ago

**Selected Answer: A**

A is correct: You assign a single CIDR IP address range as the primary CIDR block when you create a VPC and can add up to four secondary CIDR blocks after creation of the VPC.

upvoted 8 times

  **Guru4Cloud** Highly Voted 1 year, 8 months ago

**Selected Answer: A**

the architect just needs to:

Add the CIDR using the AWS console or CLI  
Create new subnets in the VPC using the new CIDR  
Launch resources in the new subnets


upvoted 5 times

  **Tendu** Most Recent 2 months, 1 week ago

**Selected Answer: A**

A - is the correct answer

upvoted 1 times

  **rosanna** 4 months, 4 weeks ago

**Selected Answer: A**

The CIDR block of a VPC can NOT be changed but can be extended.  
Thus, for the LEAST operational overhead; extend the existing one.

upvoted 1 times

  **f2e2419** 1 year, 3 months ago

**Selected Answer: A**

best option

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

A: LEAST operational overhead is by creating a new CIDR block in existing VPC.  
All other options require additional overhead of gateway or second VPC

upvoted 4 times

  **potomac** 1 year, 6 months ago

**Selected Answer: A**

After you've created your VPC, you can associate additional IPv4 CIDR blocks with the VPC

upvoted 4 times

  **TariqKipkemei** 1 year, 10 months ago



**Selected Answer: A**

A is best  
upvoted 2 times

  **Yadav\_Sanjay** 1 year, 11 months ago



**Selected Answer: A**

Add additional CIDR of bigger range  
upvoted 3 times

  **Efren** 1 year, 11 months ago

**Selected Answer: A**

Add new bigger subnets  
upvoted 3 times

  **nosense** 1 year, 11 months ago

**Selected Answer: A**

A valid  
upvoted 2 times

A company used an Amazon RDS for MySQL DB instance during application testing. Before terminating the DB instance at the end of the test cycle, a solutions architect created two backups. The solutions architect created the first backup by using the mysqldump utility to create a database dump. The solutions architect created the second backup by enabling the final DB snapshot option on RDS termination.

The company is now planning for a new test cycle and wants to create a new DB instance from the most recent backup. The company has chosen a MySQL-compatible edition of Amazon Aurora to host the DB instance.

Which solutions will create the new DB instance? (Choose two.)

- A. Import the RDS snapshot directly into Aurora.
- B. Upload the RDS snapshot to Amazon S3. Then import the RDS snapshot into Aurora.
- C. Upload the database dump to Amazon S3. Then import the database dump into Aurora.
- D. Use AWS Database Migration Service (AWS DMS) to import the RDS snapshot into Aurora.
- E. Upload the database dump to Amazon S3. Then use AWS Database Migration Service (AWS DMS) to import the database dump into Aurora.

**Correct Answer:** AC

*Community vote distribution*

AC (76%)

CE (15%)

4%

- Axaus**

Highly Voted

1 year, 11 months ago

Selected Answer: AC

A,C  
A because the snapshot is already stored in AWS.  
C because you dont need a migration tool going from MySQL to MySQL. You would use the MySQL utility.  
upvoted 13 times
- oras2023**

Highly Voted

1 year, 11 months ago

Selected Answer: AC

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.RDSMySQL.Import.html>  
upvoted 8 times
- bora4motion**

Most Recent

3 days, 22 hours ago

Selected Answer: CE

c + e is correct  
upvoted 1 times
- tch**

1 month, 1 week ago

Selected Answer: AC

Answer B is out, why upload to S3?  
This is to call the most recent backup, answer D & E are out to use AWS Database Migration Service (AWS DMS).  
upvoted 1 times
- rpmaws**

7 months, 3 weeks ago

Selected Answer: CE

AWS DMS does not support migrating data directly from an RDS snapshot. DMS can migrate data from a live RDS instance or from a database dump, but not from a snapshot.  
Also dump can be migrated to aurora using sql client.  
upvoted 3 times
- Mayur\_B**

6 months, 3 weeks ago

Correct.  
Homogeneous migration in AWS Database Migration Service (DMS) refers to migrating databases between the same type of database engine.  
Here's an example:  
  
Example:  
Source Database: Amazon RDS for MySQL  
Target Database: Amazon Aurora MySQL  
upvoted 1 times
- JackyCCK**

1 year ago

If you can use option A - "Import the RDS snapshot directly into Aurora", why go S3 in option C ?  
Non sense, A and C cannot co-exist

upvoted 1 times



  **pentium75** 1 year, 4 months ago

**Selected Answer: AC**

A per <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.RDSMySQL.Snapshot.html>

C per <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.ExtMySQL.html>

upvoted 7 times

  **aws94** 1 year, 4 months ago

**Selected Answer: AB**

A and B

upvoted 1 times

  **meowruki** 1 year, 5 months ago

**Selected Answer: AC**

Similar : <https://repost.aws/knowledge-center/aurora-postgresql-migrate-from-rds>

upvoted 3 times

  **potomac** 1 year, 6 months ago

**Selected Answer: AD**

A and C


upvoted 1 times

  **TariqKipkemei** 1 year, 6 months ago

**Selected Answer: AC**

Either import the RDS snapshot directly into Aurora or upload the database dump to Amazon S3, then import the database dump into Aurora.

upvoted 3 times

  **thanhnv142** 1 year, 6 months ago

AC:

- store dump in s3 then upload to aurora
- no need to store snapshot in s3 because is in AWS already

upvoted 5 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: CE**

C and E are the solutions that can restore the backups into Amazon Aurora.

The RDS DB snapshot contains backup data in a proprietary format that cannot be directly imported into Aurora.

The mysqldump database dump contains SQL statements that can be imported into Aurora after uploading to S3.

AWS DMS can migrate the dump file from S3 into Aurora.

upvoted 4 times

  **james2033** 1 year, 9 months ago

**Selected Answer: AC**

Amazon RDS for MySQL --> Amazon Aurora MySQL-compatible.

\* mysqldump, database dump --> (C) Upload to Amazon S3, Import dump to Aurora.

\* DB snapshot --> (A) Import RDS Snapshot directly Aurora. The correct word should be "migration". "Use console to migrate the DB snapshot and create an Aurora MySQL DB cluster with the same databases as the original MySQL DB instance."

Exclude B, because no need upload DB snapshot to Amazon S3. Exclude D, because no need Migration service. Exclude E, because no need Migration service. Use exclusion method is more easy for this question.

Related links:

- Amazon RDS create database snapshot [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_CreateSnapshot.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html)
- <https://aws.amazon.com/rds/aurora/>

upvoted 3 times

  **marufxplorer** 1 year, 10 months ago

CE

Since the backup created by the solutions architect was a database dump using the mysqldump utility, it cannot be directly imported into Aurora using RDS snapshots. Amazon Aurora has its own specific backup format that is different from RDS snapshots

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

C and E are the solutions that can restore the backups into Amazon Aurora.

The RDS DB snapshot contains backup data in a proprietary format that cannot be directly imported into Aurora.

The mysqldump database dump contains SQL statements that can be imported into Aurora after uploading to S3.

AWS DMS can migrate the dump file from S3 into Aurora.

upvoted 3 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: AC**

Migrating data from MySQL by using an Amazon S3 bucket

You can copy the full and incremental backup files from your source MySQL version 5.7 database to an Amazon S3 bucket, and then restore an Amazon Aurora MySQL DB cluster from those files.

This option can be considerably faster than migrating data using mysqldump, because using mysqldump replays all of the commands to recreate the schema and data from your source database in your new Aurora MySQL DB cluster.

By copying your source MySQL data files, Aurora MySQL can immediately use those files as the data for an Aurora MySQL DB cluster.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.ExtMySQL.html>

upvoted 4 times

  **omoakin** 1 year, 11 months ago

BE

Upload the RDS snapshot to Amazon S3. Then import the RDS snapshot into Aurora.

Upload the database dump to Amazon S3. Then use AWS Database Migration Service (AWS DMS) to import the database dump into Aurora

upvoted 1 times

A company hosts a multi-tier web application on Amazon Linux Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company observes that the Auto Scaling group launches more On-Demand Instances when the application's end users access high volumes of static web content. The company wants to optimize cost.

What should a solutions architect do to redesign the application MOST cost-effectively?

- A. Update the Auto Scaling group to use Reserved Instances instead of On-Demand Instances.
- B. Update the Auto Scaling group to scale by launching Spot Instances instead of On-Demand Instances.
- C. Create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket.
- D. Create an AWS Lambda function behind an Amazon API Gateway API to host the static website contents.

**Correct Answer:** C

*Community vote distribution*

C (100%)

- awsgeek75** 10 months ago

Selected Answer: C

C: Cost effective static content scaling = CloudFront  
A and B scale instances so not the best use of money for static content  
D Probably most expensive way of service static content at scale as you'll be charged for Lambda execution also  
upvoted 4 times
- mwwt2022** 10 months ago

Selected Answer: C

static content -> CloudFront  
upvoted 3 times
- Guru4Cloud** 1 year, 2 months ago

Selected Answer: C

implementing CloudFront to serve static content is the most cost-optimal architectural change for this use case.  
upvoted 4 times
- james2033** 1 year, 3 months ago

Selected Answer: C

Keyword "Amazon CloudFront", "high volumes of static web content", choose C.  
upvoted 3 times
- TariqKipkemei** 1 year, 4 months ago

Selected Answer: C

static web content = Amazon CloudFront  
upvoted 2 times
- alexandercamachop** 1 year, 5 months ago

Selected Answer: C

Static Web Content = S3 Always.  
CloudFront = Closer to the users locations since it will cache in the Edge nodes.  
upvoted 3 times
- cloudenthusiast** 1 year, 5 months ago

By leveraging Amazon CloudFront, you can cache and serve the static web content from edge locations worldwide, reducing the load on your EC2 instances. This can help lower the number of On-Demand Instances required to handle high volumes of static web content requests. Storing the static content in an Amazon S3 bucket and using CloudFront as a content delivery network (CDN) improves performance and reduces costs by reducing the load on your EC2 instances.  
upvoted 4 times
- Efren** 1 year, 5 months ago

Selected Answer: C

Static content, cloudFront plus S3  
upvoted 3 times
- nosense** 1 year, 5 months ago

Selected Answer: C

c for me  
upvoted 2 times

Question #442

Topic 1

A company stores several petabytes of data across multiple AWS accounts. The company uses AWS Lake Formation to manage its data lake. The company's data science team wants to securely share selective data from its accounts with the company's engineering team for analytical purposes.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Copy the required data to a common account. Create an IAM access role in that account. Grant access by specifying a permission policy that includes users from the engineering team accounts as trusted entities.
- B. Use the Lake Formation permissions Grant command in each account where the data is stored to allow the required engineering team users to access the data.
- C. Use AWS Data Exchange to privately publish the required data to the required engineering team accounts.
- D. Use Lake Formation tag-based access control to authorize and grant cross-account permissions for the required data to the engineering team accounts.

Correct Answer: D

Community vote distribution

D (100%)

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

Selected Answer: D

By utilizing Lake Formation's tag-based access control, you can define tags and tag-based policies to grant selective access to the required data for the engineering team accounts. This approach allows you to control access at a granular level without the need to copy or move the data to a common account or manage permissions individually in each account. It provides a centralized and scalable solution for securely sharing data across accounts with minimal operational overhead.

upvoted 18 times

  **NSA\_Poker** Most Recent  5 months ago

Selected Answer: D

(B) uses the CLI command that has many options: principal, TableName, ColumnNames, LFTag etc providing a way to manage granular access permissions for different users at the table and column level. That way you don't give full access to the all the data. The problem with (B) is to implement this in each account has a lot more operational overhead than (D).

upvoted 2 times

  **awsgeek75** 10 months ago

Selected Answer: D

D: Selective data = tagging  
A and B gives full access to all the data  
C is possible but with complex operational overhead as you have to publish your data to the Data Exchange. (this is based on my limited knowledge so happy to be corrected)

upvoted 4 times

  **Guru4Cloud** 1 year, 2 months ago



Selected Answer: D

D is the correct option with the least operational overhead.

Using Lake Formation tag-based access control allows granting cross-account permissions to access data in other accounts based on tags, without having to copy data or configure individual permissions in each account.

This provides a centralized, tag-based way to share selective data across accounts to authorized users with least operational overhead.

upvoted 4 times

  **luisgu** 1 year, 5 months ago

Selected Answer: D

<https://aws.amazon.com/blogs/big-data/securely-share-your-data-across-aws-accounts-using-aws-lake-formation/>

upvoted 4 times

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.

**Correct Answer: A**

*Community vote distribution*

A (76%)C (24%)

- pentium75**

Highly Voted

1 year, 4 months ago

**Selected Answer: A**

The question asks for "a cost-effective solution [ONLY TO] to minimize upload and download latency and maximize performance", not for the actual application. And the 'cost-effective solution to minimize upload and download latency and maximize performance' is S3 Transfer Acceleration. Obviously there is more required to host the app, but that is not asked for.

upvoted 14 times
- chris0975**

Highly Voted

1 year, 6 months ago

**Selected Answer: A**

The question is focused on large downloads and uploads. S3 Transfer Acceleration is what fits. CloudFront is for caching which cannot be used when the data is unique. They aren't as concerned with regular web traffic.

Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects.

upvoted 7 times
- bora4motion**

Most Recent

3 days, 22 hours ago

**Selected Answer: A**

voting for A.

upvoted 1 times
- Ben\_88**

6 months, 1 week ago

**Selected Answer: C**

S3 can't "host an application" , so even tho S3 Transfer acceleration would optimize the performances of the download/upload , still it doesn't host the app. the question is poorly written i guess

upvoted 1 times
- vincent2023**

3 months, 2 weeks ago

I agree with your point !!

upvoted 1 times
- ChymKuBoy**

8 months, 2 weeks ago

**Selected Answer: A**

A for sure

upvoted 3 times
- awsgeek75**

1 year, 3 months ago

**Selected Answer: A**

Not C, D No requirements to scale the application itself so EC2 is not applicable.  
B is for caching so not sure how/if that helps the upload speed for global users  
A is correct as Transfer Accelerator is best for uploading and downloading unique items near the user's region/location

upvoted 6 times
- tosuccess**

1 year, 4 months ago

**Selected Answer: A**

for datas greater tham 1 GB, s3 transfer acceleration is the best

upvoted 3 times

🗨️ 👤 **Cyberkayu** 1 year, 4 months ago

**Selected Answer: A**

Application users will be able to download and upload UNIQUE data up to gigabytes in size

Thus all caching related solution dont work.

upvoted 5 times

🗨️ 👤 **Goutham4981** 1 year, 5 months ago

**Selected Answer: A**

Downloading data upto gigabytes in size - Cloudfront is a content delivery service that acts as an edge caching layer for images and other data. Not a service that minimizes upload and download latency.

upvoted 2 times

🗨️ 👤 **potomac** 1 year, 6 months ago

**Selected Answer: A**

The question is focused on large downloads and uploads. S3 Transfer Acceleration is what fits. CloudFront is for caching which cannot be used when the data is unique. They aren't as concerned with regular web traffic.

C didn't mention S3. Where the data is stored?

upvoted 4 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

A doesn't mention EC2 or EKS or ECS or Elastic Beanstalk or Lambda. Where does the "scalable web application" run?

upvoted 2 times

🗨️ 👤 **JA2018** 5 months, 1 week ago

hmm... could be a red herring?

upvoted 1 times

🗨️ 👤 **beast2091** 1 year, 6 months ago

It is A.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html>

upvoted 2 times

🗨️ 👤 **danielmakita** 1 year, 6 months ago

It is A as the Transfer Acceleration will minimize upload and download latency.

If you choose C, where would the files be stored? There is no mention of any S3. Will it be stored inside the EC2? That's why I didn't go for C

upvoted 5 times

🗨️ 👤 **Sindokuhlep** 1 year, 6 months ago

**Selected Answer: C**

Amazon S3 with Transfer Acceleration (option A) is designed for speeding up uploads to Amazon S3, and it's not used for hosting scalable web applications. It doesn't mention using EC2 instances for hosting the application.

upvoted 4 times

🗨️ 👤 **canonlycontainletters1** 1 year, 6 months ago

**Selected Answer: C**

My answer is C

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 6 months ago

C because A is for upload data to S3, not for web app

upvoted 1 times

🗨️ 👤 **DamyanG** 1 year, 6 months ago

**Selected Answer: C**

The correct answer is C!!! It is not A, because

- Amazon S3 with Transfer Acceleration (option A) is designed for speeding up uploads to Amazon S3, and it's not used for hosting scalable web applications. It doesn't mention using EC2 instances for hosting the application.

upvoted 3 times

🗨️ 👤 **Victory007** 1 year, 6 months ago

**Selected Answer: C**

Amazon CloudFront is a global content delivery network (CDN) that delivers web content to users with low latency and high transfer speeds. It does this by caching content at edge locations around the world, which are closer to the users than the origin server.

By using Amazon EC2 with Auto Scaling and Amazon CloudFront, the company can create a scalable and high-performance web application that is accessible to users from different geographic regions of the world.

upvoted 1 times

A company has hired a solutions architect to design a reliable architecture for its application. The application consists of one Amazon RDS DB instance and two manually provisioned Amazon EC2 instances that run web servers. The EC2 instances are located in a single Availability Zone.

An employee recently deleted the DB instance, and the application was unavailable for 24 hours as a result. The company is concerned with the overall reliability of its environment.

What should the solutions architect do to maximize reliability of the application's infrastructure?

- A. Delete one EC2 instance and enable termination protection on the other EC2 instance. Update the DB instance to be Multi-AZ, and enable deletion protection.
- B. Update the DB instance to be Multi-AZ, and enable deletion protection. Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling group across multiple Availability Zones.
- C. Create an additional DB instance along with an Amazon API Gateway and an AWS Lambda function. Configure the application to invoke the Lambda function through API Gateway. Have the Lambda function write the data to the two DB instances.
- D. Place the EC2 instances in an EC2 Auto Scaling group that has multiple subnets located in multiple Availability Zones. Use Spot Instances instead of On-Demand Instances. Set up Amazon CloudWatch alarms to monitor the health of the instances Update the DB instance to be Multi-AZ, and enable deletion protection.

Correct Answer: B

Community vote distribution

B (100%)

  **awsgeek75** Highly Voted 10 months ago

Option E: Sack the employee who did this :)  
upvoted 10 times

  **nosense** Highly Voted 1 year, 5 months ago

**Selected Answer: B**

B is correct. HA ensured by DB in Mutli-AZ and EC2 in AG  
upvoted 5 times

  **wizcloudifa** Most Recent 6 months, 2 weeks ago

**Selected Answer: B**

A: delete one instance, why?. Although takes care of reliability of DB instance however not EC2.  
B. seems perfect as takes care of reliability of both EC2 as well as DB  
C. DB instance's reliability is not taken care of  
D. seems to be trying to address cost alongside reliability of EC2 and DB.  
upvoted 3 times

  **awsgeek75** 10 months ago

**Selected Answer: B**

A: Deleting one EC2 instance makes no sense. Why would you do that?  
C: API Gateway, Lambda etc are all nice but they don't solve the problem of DB instance deletion  
D: EC2 subnet blah blah, what? The problem is reliability, not networking!

B is correct as it solves the DB deletion issue and increases reliability by Multi AZ scaling of EC2 instances  
upvoted 5 times

  **Guru4Cloud** 1 year, 2 months ago


**Selected Answer: B**

The key points:  
° RDS Multi-AZ and deletion protection provide high availability for the database.  
° The load balancer and Auto Scaling group across AZs give high availability for EC2.  
° Options A, C, D have limitations that would reduce reliability vs option B.  
upvoted 3 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: B**

Update the DB instance to be Multi-AZ, and enable deletion protection. Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling group across multiple Availability Zones  
upvoted 2 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: B**

B for sure.

upvoted 2 times

  **alexandercamachop** 1 year, 5 months ago

**Selected Answer: B**

It is the only one with High Availability.

Amazon RDS with Multi AZ

EC2 with Auto Scaling Group in Multi Az

upvoted 2 times

  **omoakin** 1 year, 5 months ago

same question from

<https://www.examttopics.com/exams/amazon/aws-certified-solutions-architect-associate-saa-c02/>

long time ago and still same option B

upvoted 3 times

A company is storing 700 terabytes of data on a large network-attached storage (NAS) system in its corporate data center. The company has a hybrid environment with a 10 Gbps AWS Direct Connect connection.

After an audit from a regulator, the company has 90 days to move the data to the cloud. The company needs to move the data efficiently and without disruption. The company still needs to be able to access and update the data during the transfer window.

Which solution will meet these requirements?

- A. Create an AWS DataSync agent in the corporate data center. Create a data transfer task Start the transfer to an Amazon S3 bucket.
- B. Back up the data to AWS Snowball Edge Storage Optimized devices. Ship the devices to an AWS data center. Mount a target Amazon S3 bucket on the on-premises file system.
- C. Use rsync to copy the data directly from local storage to a designated Amazon S3 bucket over the Direct Connect connection.
- D. Back up the data on tapes. Ship the tapes to an AWS data center. Mount a target Amazon S3 bucket on the on-premises file system.

Correct Answer: A

Community vote distribution

A (100%)

  **wRhIH** Highly Voted 1 year, 10 months ago

For those who wonders why not B. Snowball Edge Storage Optimized device for data transfer is up to 100TB  
<https://docs.aws.amazon.com/snowball/latest/developer-guide/device-differences.html>  
upvoted 12 times

  **Maru86** 1 year, 1 month ago


The question explicitly mentioned "devices", also Snowball Edge Storage Optimized is 80TB HDD. So it is possible, but the answer is A because we can transfer with DataSync in 6.5 days.  
upvoted 4 times

  **smartegnine** 1 year, 10 months ago

10GBs \* 24\*60\*60 =864,000 GB estimate around 864 TB a day, 2 days will transfer all data. But for snowball at least 4 days for delivery to the data center.  
upvoted 1 times

  **siGma182** 1 year, 9 months ago

This account is wrong but I get your point. It is wrong cause 10Gb/s is not the same as 10GB/s (Gigabits vs Gigabytes). However, the correct count is 864Tb/8 = 108TB per day. In one week you should've transferred all the data.  
upvoted 8 times

  **JA2018** 5 months, 1 week ago

folks, you had to consider the typical TCP overheads for such over-the-wire data transfers.... it could be as high as 70% of the total available bandwidth due to a wide variety of factors  
upvoted 1 times

  **Maru86** 1 year, 1 month ago

That's right, 1 GB = 8 Gb. Essentially we have a speed of 1.25GB/s.  
upvoted 2 times

  **hsinchang** Highly Voted 1 year, 9 months ago



**Selected Answer: A**

Access during the transfer window -> DataSync  
upvoted 9 times



  **surajkrishnamurthy** Most Recent 3 months ago

**Selected Answer: A**

700TB with a Bandwidth of 10Gbps takes approximately 7 days, 3 Hours, 2 Minutes, 7 Seconds  
As per the below link>>  
<https://expedient.com/knowledgebase/tools-and-calculators/file-transfer-time-calculator/>  
So Option A is Correct & using AWS Datasync facilitates further access to data still present in the OnPrem to access and S3 provides with unlimited data storage capacity  
upvoted 1 times

  **Danilus** 5 months, 4 weeks ago

A- Porque La compañía todavía necesita poder acceder y actualizar los datos durante la ventana de transferencia  
B-Aunque esta podría ser una opción porque dice de enviar los dispositivos al centro de datos que serían más o menos entre 9 y 10 dispositivos ya que el almacenamiento de storage optimized es de 80 TB el problema está en que con Snowball Edge Storage Optimized no se puede acceder y actualizar a los datos durante la transferencia  
upvoted 1 times

  **JA2018** 5 months, 1 week ago  
From Google Translate:

A- Because The company still needs to be able to access and update the data during the transfer window

B-Although this could be an option because it says that sending the devices to the data center would be more or less between 9 and 10 devices since the storage optimized storage is 80 TB, the problem is that with Snowball Edge Storage Optimized it does not Data can be accessed and updated during transfer  
upvoted 3 times

  **Danilus** 5 months, 4 weeks ago

**Selected Answer: A**

A- Porque La compañía todavía necesita poder acceder y actualizar los datos durante la ventana de transferencia  
B-Aunque esta podría ser una opción porque dice de enviar los dispositivos al centro de datos que serían más o menos entre 9 y 10 dispositivos ya que el almacenamiento de storage optimized es de 80 TB el problema está en que con Snowball Edge Storage Optimized no se puede acceder y actualizar a los datos durante la transferencia  
upvoted 1 times

  **MatAlves** 7 months, 3 weeks ago



**Selected Answer: A**

Finally, a company with good bandwidth.  
upvoted 4 times

  **NSA\_Poker** 11 months, 1 week ago


**Selected Answer: A**

(B) is incorrect bc although Mountpoint for S3 is possible for on-premises NAS, this is not as efficient as AWS DataSync. Data updates made during the transfer window would have to be resolved later.  
upvoted 2 times


  **Burrito69** 1 year, 1 month ago

I will put simple calculation for you guys to just store in your head to quickly answer:  
for 10GBPS it's 1.25GBPS because it's bits to bytes.  
for one minute it's 75GBPS  
for one hour it's 4500 GBPS  
for one day it's 10.8 TB

so you can calculate easily if you just store these numbers in your head. let's say if question is 1GBPS DirectConnect that means everything above should be divided by 8. cool  
upvoted 1 times

  **JA2018** 5 months, 1 week ago

folks, you had to consider the typical TCP overheads for such over-the-wire data transfers.... it could be as high as 70% of the total available bandwidth due to a wide variety of factors  
  
in a worst case scenario, we could be looking @ 3.24 TB/ day  
upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

Critical requirement: "The company needs to move the data efficiently and without disruption."  
B: Causes disruption  
C: I don't think that is possible without a gateway kind of thing  
D: Tape backups? "Mount a target Amazon S3 bucket on the on-premises file system"? This requires some gateway which is not mentioned  
  
A is the answer as DataSync allows transfer without disruption and with 10Gbps, it can be done in 90 days.  
upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago


**Selected Answer: A**

AWS DataSync can efficiently transfer large datasets from on-premises NAS to Amazon S3 over Direct Connect.  
  
DataSync allows accessing and updating the data continuously during the transfer process.  
upvoted 5 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: A**

AWS DataSync is a secure, online service that automates and accelerates moving data between on-premises and AWS Storage services.  
upvoted 3 times

  **omoakin** 1 year, 11 months ago

A

<https://www.examttopics.com/discussions/amazon/view/46492-exam-aws-certified-solutions-architect-associate-saa-c02/#:~:text=Exam%20question%20from,Question%20%23%3A%20385>

upvoted 2 times

  **cloudenthusiast** 1 year, 11 months ago

**Selected Answer: A**

By leveraging AWS DataSync in combination with AWS Direct Connect, the company can efficiently and securely transfer its 700 terabytes of data to an Amazon S3 bucket without disruption. The solution allows continued access and updates to the data during the transfer window, ensuring business continuity throughout the migration process.

upvoted 4 times

  **nosense** 1 year, 11 months ago

**Selected Answer: A**

A for me, bcs egde storage up to 100tb

upvoted 5 times

A company stores data in PDF format in an Amazon S3 bucket. The company must follow a legal requirement to retain all new and existing data in Amazon S3 for 7 years.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Turn on the S3 Versioning feature for the S3 bucket. Configure S3 Lifecycle to delete the data after 7 years. Configure multi-factor authentication (MFA) delete for all S3 objects.
- B. Turn on S3 Object Lock with governance retention mode for the S3 bucket. Set the retention period to expire after 7 years. Recopy all existing objects to bring the existing data into compliance.
- C. Turn on S3 Object Lock with compliance retention mode for the S3 bucket. Set the retention period to expire after 7 years. Recopy all existing objects to bring the existing data into compliance.
- D. Turn on S3 Object Lock with compliance retention mode for the S3 bucket. Set the retention period to expire after 7 years. Use S3 Batch Operations to bring the existing data into compliance.

Correct Answer: D

Community vote distribution

D (86%)

8%

- omarshaban

Highly Voted

1 year, 3 months ago

THIS WAS IN MY EXAM

upvoted 11 times
- awsgeek75

1 year, 3 months ago

Did you pass?

upvoted 4 times
- DrWatson

Highly Voted

1 year, 11 months ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-retention-date.html>

upvoted 6 times
- Lin878

Most Recent

10 months, 3 weeks ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html>

upvoted 3 times
- iapps369

1 year, 3 months ago

D

as S3 batch operations reduce risk and manual copy/paste overhead.

upvoted 6 times
- awsgeek75

1 year, 3 months ago

Selected Answer: D

A: Versioning, not relevant

B: Governance, it won't enforce object lock

C: Recopy existing objects may work but lots of operational overhead (see link)

D: Compliance on existing objects with batch operations is least operational overhead

<https://repost.aws/questions/QUGKrl8XRLTEeulzUHq0Ikew/s3-object-lock-on-existing-s3-objects>

upvoted 5 times
- awsgeek75

1 year, 3 months ago

With option C, you have to copy the object for it to be complaint and then delete the original as only the new copy will be compliant. So D is the only option

upvoted 2 times
- mr123dd

1 year, 4 months ago

Selected Answer: A

To enable Object Lock on an Amazon S3 bucket, you must first enable versioning on that bucket. other 3 option did not enable versioning first

upvoted 2 times
- fb4afde

1 year, 4 months ago

**Selected Answer: D**

Recopying offers more control but requires users to manage the process. S3 Batch Operations automates the process at scale but with less granular control - LEAST operational overhead

upvoted 3 times

  **moonster** 1 year, 5 months ago

Its C because you only need to recopy all existing objects one time, so why use S3 batch operations if new datas going to be in compliance retention mode? I can see why its C although my initial gut answer was D.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

What if I don't have the original files anymore? Where should I copy them from?

upvoted 3 times

  **kwang312** 1 year, 7 months ago

You can only enable Object Lock for new buckets. If you want to turn on Object Lock for an existing bucket, contact AWS Support.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

You need a token from AWS Support, but you CAN enable Object Lock for an existing bucket.

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

Turn on S3 Object Lock with compliance retention mode for the S3 bucket. Set the retention period to expire after 7 years. Use S3 Batch Operations to bring the existing data into compliance.

upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: D**

To replicate existing object/data in S3 Bucket to bring them to compliance, optionally we use "S3 Batch Replication", so option D is the most appropriate, especially if we have big data in S3.

upvoted 2 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: D**

For minimum ops D is best

upvoted 2 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: C**

Batch operations will add operational overhead.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

And gathering all the files for copying them again does not?

upvoted 2 times

  **Abrar2022** 1 year, 11 months ago

Use Object Lock in Compliance mode. Then Use Batch operation.

WRONG>>manual work and not automated>>>Recopy all existing objects to bring the existing data into compliance.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

Batch IS automated. You just need to create the batch which is a one-time operation.

"Recopy all existing objects" is not operational overhead?

upvoted 2 times

  **omoakin** 1 year, 11 months ago

C

When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period.

upvoted 3 times

  **omoakin** 1 year, 11 months ago

error i meant to type D

i wont do recopy

upvoted 2 times

  **lucdt4** 1 year, 11 months ago

No, D for me because the requirement is LEAST operational overhead

So RECOPy ..... is the manual operation -> C is wrong

D is correct



upvoted 3 times

  **cloudenthusiast** 1 year, 11 months ago

Recopying vs. S3 Batch Operations: In Option C, the recommendation is to recopy all existing objects to ensure they have the appropriate retention settings. This can be done using simple S3 copy operations. On the other hand, Option D suggests using S3 Batch Operations, which is a more advanced feature and may require additional configuration and management. S3 Batch Operations can be beneficial if you have a massive number of objects and need to perform complex operations, but it might introduce more overhead for this specific use case.

Operational complexity: Option C has a straightforward process of recopying existing objects. It is a well-known operation in S3 and doesn't require additional setup or management. Option D introduces the need to set up and configure S3 Batch Operations, which can involve creating job definitions, specifying job parameters, and monitoring the progress of batch operations. This additional complexity may increase the operational overhead.

upvoted 2 times

  **Efren** 1 year, 11 months ago

**Selected Answer: D**

You need AWS Batch to re-apply certain config to files that were already in S3, like encryption

upvoted 5 times

A company has a stateless web application that runs on AWS Lambda functions that are invoked by Amazon API Gateway. The company wants to deploy the application across multiple AWS Regions to provide Regional failover capabilities.

What should a solutions architect do to route traffic to multiple Regions?

- A. Create Amazon Route 53 health checks for each Region. Use an active-active failover configuration.
- B. Create an Amazon CloudFront distribution with an origin for each Region. Use CloudFront health checks to route traffic.
- C. Create a transit gateway. Attach the transit gateway to the API Gateway endpoint in each Region. Configure the transit gateway to route requests.
- D. Create an Application Load Balancer in the primary Region. Set the target group to point to the API Gateway endpoint hostnames in each Region.

**Correct Answer: A**

*Community vote distribution*

A (89%)

11%

- TariqKipkemei**

Highly Voted

1 year, 10 months ago

Selected Answer: A

Global, Reduce latency, health checks, no failover = Amazon CloudFront  
Global ,Reduce latency, health checks, failover, Route traffic = Amazon Route 53  
option A has more weight.  
upvoted 32 times
- ManikRoy**

12 months ago

Cloud front does have failover capabilities.  
[https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high\\_availability\\_origin\\_failover.html#:~:text=the%20secondary%20origin.,Note,Choose%20Create%20origin%20group](https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html#:~:text=the%20secondary%20origin.,Note,Choose%20Create%20origin%20group).  
upvoted 3 times
- Anmol\_1010**

1 year, 6 months ago

nicley explained  
upvoted 3 times
- examtopictempacc**

Highly Voted

1 year, 11 months ago

Selected Answer: A

A. I'm not an expert in this area, but I still want to express my opinion. After carefully reviewing the question and thinking about it for a long time, I actually don't know the reason. As I mentioned at the beginning, I'm not an expert in this field.  
upvoted 19 times
- awsgeek75**

1 year, 3 months ago

All the explanation you need for this question and option A is in this article:  
<https://aws.amazon.com/blogs/compute/building-a-multi-region-serverless-application-with-amazon-api-gateway-and-aws-lambda/>  
upvoted 3 times
- MatAlves**

Most Recent

7 months, 3 weeks ago

Selected Answer: A

Correct me if I'm wrong but CloudFront DOES NOT have health check capabilities out of the box. Route 53 and Global Accelerator do.  
upvoted 3 times
- ChymKuBoy**

8 months, 1 week ago

Selected Answer: A

A for sure  
upvoted 2 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: A

B: Caching solution. Not ideal for failover although it will work. Would have been a correct answer if A wasn't an option  
C: Transit gateway is for VPC connectivity not AWS API or Lambda  
D: Even if it was possible, there is a primary region dependency of ALB  
A: correct because R53 health checks can failover across regions  
  
Good explanation here:  
<https://aws.amazon.com/blogs/compute/building-a-multi-region-serverless-application-with-amazon-api-gateway-and-aws-lambda/>

upvoted 4 times

  **awsgeek75** 1 year, 3 months ago

The article also explains why you cannot use a CloudFront distribution for API Gateway, Lambda for failover

upvoted 2 times

  **tosuccess** 1 year, 4 months ago

**Selected Answer: B**

we can set primary and secondry regions in cloud front for failover.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: A**

Application is serverless, it doesn't matter where it runs, so can be active-active setup and run wherever the request comes in. Route 53 with health checks will route to a healthy region.

B, could work too, but CloudFront is for caching which does not seem to help with an API. The goal here is "failover capabilities", not caching/performance/latency etc.

upvoted 3 times

  **Goutham4981** 1 year, 5 months ago

**Selected Answer: A**

In activ active failover config, route53 continuously monitors its endpoints and if one of them is unhealthy, it excludes the region/endpoint from its valid traffic route - Only Sensible option

Cloudfront is a content delivery network - not used to route traffic

Transit gateway for traffic routing - aws devs will hit us with a stick on hearing this option

You cant use a load balancer for cross region load balancing - invalid

upvoted 2 times

  **potomac** 1 year, 6 months ago

**Selected Answer: A**

Global ,Reduce latency, health checks, failover, Route traffic = Amazon Route 53

upvoted 2 times

  **yodelin** 1 year, 6 months ago

"What the?" yeah I know right

upvoted 1 times

  **jrestrepob** 1 year, 8 months ago

**Selected Answer: B**

"Stateless applications provide one service or function and use content delivery network (CDN), web, or print servers to process these short-term requests.

<https://docs.aws.amazon.com/architecture-diagrams/latest/multi-region-api-gateway-with-cloudfront/multi-region-api-gateway-with-cloudfront.html>

upvoted 1 times

  **deechean** 1 year, 8 months ago

its not static content, actually they deployed a API Gateway backed by lambda

upvoted 3 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: A**

A option does make sense.

upvoted 2 times

  **Sangsation** 1 year, 10 months ago

**Selected Answer: B**

By creating an Amazon CloudFront distribution with origins in each AWS Region where the application is deployed, you can leverage CloudFront's global edge network to route traffic to the closest available Region. CloudFront will automatically route the traffic based on the client's location and the health of the origins using CloudFront health checks.

Option A (creating Amazon Route 53 health checks with an active-active failover configuration) is not suitable for this scenario as it is primarily used for failover between different endpoints within the same Region, rather than routing traffic to different Regions.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

Option A does not speak of Route 53 failover routing policies.

upvoted 1 times

  **Axeashes** 1 year, 10 months ago

**Selected Answer: A**

<https://aws.amazon.com/blogs/compute/building-a-multi-region-serverless-application-with-amazon-api-gateway-and-aws-lambda/>

upvoted 4 times

🗨️ 👤 **Gooniegoogoo** 1 year, 10 months ago  
that is from 2017.. i wonder if it is still relevant..  
upvoted 1 times

🗨️ 👤 **DrWatson** 1 year, 11 months ago  
**Selected Answer: A**  
<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>  
upvoted 2 times

🗨️ 👤 **antropaws** 1 year, 11 months ago  
**Selected Answer: A**  
I understand that you can use Route 53 to provide regional failover.  
upvoted 1 times

🗨️ 👤 **alexandercamachop** 1 year, 11 months ago  
**Selected Answer: A**  
To route traffic to multiple AWS Regions and provide regional failover capabilities for a stateless web application running on AWS Lambda functions invoked by Amazon API Gateway, you can use Amazon Route 53 with an active-active failover configuration.  
  
By creating Amazon Route 53 health checks for each Region and configuring an active-active failover configuration, Route 53 can monitor the health of the endpoints in each Region and route traffic to healthy endpoints. In the event of a failure in one Region, Route 53 automatically routes traffic to the healthy endpoints in other Regions.  
  
This setup ensures high availability and failover capabilities for your web application across multiple AWS Regions.  
upvoted 3 times

A company has two VPCs named Management and Production. The Management VPC uses VPNs through a customer gateway to connect to a single device in the data center. The Production VPC uses a virtual private gateway with two attached AWS Direct Connect connections. The Management and Production VPCs both use a single VPC peering connection to allow communication between the applications.

What should a solutions architect do to mitigate any single point of failure in this architecture?

- A. Add a set of VPNs between the Management and Production VPCs.
- B. Add a second virtual private gateway and attach it to the Management VPC.
- C. Add a second set of VPNs to the Management VPC from a second customer gateway device.
- D. Add a second VPC peering connection between the Management VPC and the Production VPC.

**Correct Answer: C**

*Community vote distribution*

C (95%)5%

**Guru4Cloud**

Highly Voted

8 months, 1 week ago

Selected Answer: C

C is the correct option to mitigate the single point of failure.

The Management VPC currently has a single VPN connection through one customer gateway device. This is a single point of failure.

Adding a second set of VPN connections from the Management VPC to a second customer gateway device provides redundancy and eliminates this single point of failure.

upvoted 7 times

**Guru4Cloud**

8 months, 1 week ago

As @Abrar2022 explains

(production) VPN 1-----> cgw 1

(management) VPN 2-----> cgw

upvoted 4 times

**Abrar2022**

Highly Voted

11 months ago

(production) VPN 1-----> cgw 1

(management) VPN 2-----> cgw 2

upvoted 6 times

**Abrar2022**

11 months ago

ANSWER IS C

upvoted 2 times

**Dantecito**

Most Recent

2 months, 3 weeks ago

Selected Answer: D

but what if the VPC peering connection fails? how The Management and Production vpc communicate?

upvoted 1 times

**bsbs1234**

7 months ago

C,

(production) --PrivateGateway----->Direct Connect Gateway 1 ---> cgw 1 ---> DataCenter

(production) -- PrivateGateway -----> Direct Connect Gateway 2 --->cgw 2 --> DataCenter

(Management) -- > VPN ----> (Direct Connect Gateway 1?) --->cgw1 ---> dataCenter---> device in dataCenter

upvoted 2 times

**omoakin**

11 months, 3 weeks ago

I agree to C

upvoted 2 times

**cloudenthusiast**

11 months, 3 weeks ago

Selected Answer: C

option D is not a valid solution for mitigating single points of failure in the architecture. I apologize for the confusion caused by the incorrect information.

To mitigate single points of failure in the architecture, you can consider implementing option C: adding a second set of VPNs to the Management

VPC from a second customer gateway device. This will introduce redundancy at the VPN connection level for the Management VPC, ensuring that if one customer gateway or VPN connection fails, the other connection can still provide connectivity to the data center.


upvoted 5 times

  **Efren** 11 months, 3 weeks ago

**Selected Answer: C**

Redundant VPN connections: Instead of relying on a single device in the data center, the Management VPC should have redundant VPN connections established through multiple customer gateways. This will ensure high availability and fault tolerance in case one of the VPN connections or customer gateways fails.

upvoted 5 times

  **nosense** 11 months, 3 weeks ago

**Selected Answer: C**

<https://www.examttopics.com/discussions/amazon/view/53908-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 3 times

A company runs its application on an Oracle database. The company plans to quickly migrate to AWS because of limited resources for the database, backup administration, and data center maintenance. The application uses third-party database features that require privileged access.

Which solution will help the company migrate the database to AWS MOST cost-effectively?

- A. Migrate the database to Amazon RDS for Oracle. Replace third-party features with cloud services.
- B. Migrate the database to Amazon RDS Custom for Oracle. Customize the database settings to support third-party features.
- C. Migrate the database to an Amazon EC2 Amazon Machine Image (AMI) for Oracle. Customize the database settings to support third-party features.
- D. Migrate the database to Amazon RDS for PostgreSQL by rewriting the application code to remove dependency on Oracle APEX.

**Correct Answer: B**

*Community vote distribution*

B (97%)

- AbRAR2022**

Highly Voted

1 year, 5 months ago

**Selected Answer: B**

RDS Custom since it's related to 3rd vendor  
RDS Custom since it's related to 3rd vendor  
RDS Custom since it's related to 3rd vendor  
upvoted 6 times
- awsgeek75**

Most Recent

10 months ago

**Selected Answer: B**

Key constraints: Limited resources for DB admin and cost. 3rd party db features with privileged access.  
A: Won't work due to 3rd party features  
C: AMI with Oracle may work but again overhead of backed, maintenance etc  
D: Too much overhead in rewrite  
B: Actually supports Oracle 3rd party features  
Caution: If this is only about APEX as suggested in option D, then A is also a possible answer:  
<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.Oracle.Options.APEX.html>  
upvoted 3 times
- awsgeek75**

10 months ago

Action ignore the last line of my previous comment, A is not a valid option in any case as it suggest replacing 3rd party features with cloud services which is not possible without more details.  
upvoted 2 times
- pentium75**

10 months, 1 week ago

**Selected Answer: B**

"Amazon RDS Custom is a managed database service for applications that require customization of the underlying operating system and database environment. Benefits of RDS automation with the access needed for legacy, packaged, and custom applications."  
  
That should allow the "privileged access".  
upvoted 2 times
- Guru4Cloud**

1 year, 2 months ago

**Selected Answer: B**

Migrate the database to Amazon RDS Custom for Oracle. Customize the database settings to support third-party features.  
upvoted 4 times
- TariqKipkemei**

1 year, 4 months ago

**Selected Answer: B**

Custom database features = Amazon RDS Custom for Oracle  
upvoted 4 times
- antropaws**

1 year, 5 months ago

**Selected Answer: B**

Most likely B.  
upvoted 2 times
- omoakin**

1 year, 5 months ago

CCCCCCCCCCCCCCCCCCCC

upvoted 1 times

  **aqmdla2002** 1 year, 5 months ago

**Selected Answer: B**

<https://aws.amazon.com/about-aws/whats-new/2021/10/amazon-rds-custom-oracle/>

upvoted 3 times

  **hiroohiroo** 1 year, 5 months ago

**Selected Answer: B**

[https://docs.aws.amazon.com/ja\\_jp/AmazonRDS/latest/UserGuide/Oracle.Resources.html](https://docs.aws.amazon.com/ja_jp/AmazonRDS/latest/UserGuide/Oracle.Resources.html)

upvoted 2 times

  **karbob** 1 year, 5 months ago

Amazon RDS Custom for Oracle, which is not an actual service. !!!!

upvoted 2 times

  **nosense** 1 year, 5 months ago

Option C is also a valid solution, but it is not as cost-effective as option B.  
Option C requires the company to manage its own database infrastructure, which can be expensive and time-consuming. Additionally, the company will need to purchase and maintain Oracle licenses.

upvoted 3 times

  **y0** 1 year, 5 months ago

RDS Custom enables the capability to access the underlying database and OS so as to configure additional settings to support 3rd party. This feature is applicable only for Oracle and Postgresql

upvoted 2 times

  **y0** 1 year, 5 months ago

Sorry Oracle and sql server (not posstgresql)

upvoted 2 times

  **omoakin** 1 year, 5 months ago

I will say C cos of this  
"application uses third-party "

upvoted 1 times

  **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: C**



Should not it be since for Ec2, the company will have full control over the database and this is the reason that they are moving to AWS in the first place "The company plans to quickly migrate to AWS because of limited resources for the database, backup administration, and data center maintenance?"

upvoted 1 times

  **pentium75** 10 months, 1 week ago

"Amazon RDS Custom (B) is a managed database service for applications that require customization of the underlying operating system and database environment. Benefits of RDS automation with the access needed for legacy, packaged, and custom applications."

upvoted 2 times

  **Efren** 1 year, 5 months ago

**Selected Answer: B**

RDS Custom when is something related to 3rd vendor, for me

upvoted 2 times

  **nosense** 1 year, 5 months ago

not sure, but b probably

upvoted 3 times

A company has a three-tier web application that is in a single server. The company wants to migrate the application to the AWS Cloud. The company also wants the application to align with the AWS Well-Architected Framework and to be consistent with AWS recommended best practices for security, scalability, and resiliency.

Which combination of solutions will meet these requirements? (Choose three.)

- A. Create a VPC across two Availability Zones with the application's existing architecture. Host the application with existing architecture on an Amazon EC2 instance in a private subnet in each Availability Zone with EC2 Auto Scaling groups. Secure the EC2 instance with security groups and network access control lists (network ACLs).
- B. Set up security groups and network access control lists (network ACLs) to control access to the database layer. Set up a single Amazon RDS database in a private subnet.
- C. Create a VPC across two Availability Zones. Refactor the application to host the web tier, application tier, and database tier. Host each tier on its own private subnet with Auto Scaling groups for the web tier and application tier.
- D. Use a single Amazon RDS database. Allow database access only from the application tier security group.
- E. Use Elastic Load Balancers in front of the web tier. Control access by using security groups containing references to each layer's security groups.
- F. Use an Amazon RDS database Multi-AZ cluster deployment in private subnets. Allow database access only from application tier security groups.

**Correct Answer:** CEF

Community vote distribution

CEF (96%)

4%

  **awsgeek75** Highly Voted 10 months ago

**Selected Answer:** CEF

The wording on this question makes things ambiguous for C. But, remember well-architected so:

A: Not ideal as it is suggesting using existing architecture but with autoscaling EC2. Doesn't leave room for improvement on scaling or reliability on each tier.

B: Single RDS, not well-architected

D: Again, single RDS

E,F are good options and C is only remaining good one.

upvoted 8 times

  **awsgeek75** 10 months ago

C is badly worded IMHO because of this part " Refactor the application to host the web tier, application tier, and database tier." The database tier tier just makes it confusing when you don't read E and F.

upvoted 2 times

  **Abrar2022** Highly Voted 1 year, 5 months ago

**Selected Answer:** CEF

C-scalable and resilient

E-high availability of the application

F-Multi-AZ configuration provides high availability

upvoted 6 times

  **4919c76** Most Recent 2 weeks ago

**Selected Answer:** EF

The web tier can only be in a public subnet so C is wrong

upvoted 1 times

  **ChhatwaniB** 3 weeks ago

in Option C , it saying all layer should be in it's own private subnet . Isn't the web layer should be in public subnet ?

upvoted 1 times

  **ChhatwaniB** 3 weeks ago

in Option C , it saying all layer should be in it's own private subnet . Isn't the web layer should be in public subnet ?

upvoted 1 times

  **Burrito69** 7 months, 1 week ago

remove singles and remove network ACLs

upvoted 3 times

  **jjcode** 9 months, 2 weeks ago

i would flag this on the test and do it last.

upvoted 5 times

  **argl1995** 1 year, 3 months ago

option A cannot be the answer as Security group is at instance level whereas a NACL is at the subnet level. Having said that option C is the right one as the VPC cannot span across the regions and here it is mentioned two AZs for which I am guessing it is a default VPC which is created in each region with a subnet in each AZ.

upvoted 2 times

  **argl1995** 1 year, 3 months ago

So, CEF is the right answer

upvoted 2 times

  **Gooniegoogoo** 1 year, 4 months ago

How can you create a VPC across 2 AZ? i only see EF here.. if they mean 2 separate VPC then that is different but a VPC cannot span two AZ..

upvoted 1 times

  **lemur88** 1 year, 2 months ago

A VPC most definitely can span across 2 AZ. You may be thinking of subnets.


upvoted 3 times

  **marufexplorer** 1 year, 4 months ago

I also agree with CEF but chatGPT answer is ACE. A and C is the similar

Another Logic F is not True because in the question not mentioned about DB

upvoted 1 times

  **awsgeek75** 9 months, 2 weeks ago

ChatGPT is a language parser. It is not an AWS solution architect!

upvoted 3 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: CEF**

CEF is best

upvoted 2 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: CEF**

It's clearly CEF.

upvoted 2 times



  **omoakin** 1 year, 5 months ago

B- to control access to database

C-scalable and resilient

E-high availability of the application

upvoted 1 times

  **lucdt4** 1 year, 5 months ago

**Selected Answer: CEF**

CEF

A: application's existing architecture is wrong (single AZ)

B: single AZ

D: Single AZ

upvoted 2 times

  **cloudenthusiast** 1 year, 5 months ago

C.

This solution follows the recommended architecture pattern of separating the web, application, and database tiers into different subnets. It provides better security, scalability, and fault tolerance.



E.By using Elastic Load Balancers (ELBs), you can distribute traffic to multiple instances of the web tier, increasing scalability and availability.

Controlling access through security groups allows for fine-grained control and ensures only authorized traffic reaches each layer.

F.

Deploying an Amazon RDS database in a Multi-AZ configuration provides high availability and automatic failover. Placing the database in private subnets enhances security. Allowing database access only from the application tier security groups limits exposure and follows the principle of least privilege.

upvoted 5 times

  **mwwt2022** 10 months ago

good explanation

upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: CEF**

Only this valid for best practices and well architected  
upvoted 5 times

A company is migrating its applications and databases to the AWS Cloud. The company will use Amazon Elastic Container Service (Amazon ECS), AWS Direct Connect, and Amazon RDS.



Which activities will be managed by the company's operational team? (Choose three.)


- A. Management of the Amazon RDS infrastructure layer, operating system, and platforms
- B. Creation of an Amazon RDS DB instance and configuring the scheduled maintenance window
- C. Configuration of additional software components on Amazon ECS for monitoring, patch management, log management, and host intrusion detection
- D. Installation of patches for all minor and major database versions for Amazon RDS
- E. Ensure the physical security of the Amazon RDS infrastructure in the data center
- F. Encryption of the data that moves in transit through Direct Connect

**Correct Answer:** BCF

*Community vote distribution*



BCF (98%)


-   **pentium75**

Highly Voted 

 10 months, 1 week ago



Selected Answer: BCF


ADE = AWS responsibility  
upvoted 10 times
-   **awsgeek75**

Highly Voted 

 9 months, 2 weeks ago

Selected Answer: BCF

Just to clarify on F. Direct Connect is an ISP and AWS offering, I consider it as a physical connection just like you get from your ISP at home. There is not security on it until you build security on the connection. AWS provides Direct Connect but it does not provide encryption level security on data movement through it by default. It's the customer's responsibility.  
upvoted 9 times
-   **Guru4Cloud**



Most Recent 

 1 year, 2 months ago

Selected Answer: BCF

B: Creating an RDS instance and configuring the maintenance window is done by the customer.

C: Adding monitoring, logging, etc on ECS is managed by the customer.



F: Encrypting Direct Connect traffic is handled by the customer.  
upvoted 5 times
-   **james2033** 1 year, 3 months ago



Selected Answer: BCF

In question has 3 keyword "Amazon ECS", "AWS Direct Connect", "Amazon RDS". With per Amazon services, choose 1 according answer. Has 6 items, need pick 3 items.



ECS --> choose C.

Direct Connect --> choose F.

RDS --> Exclude A (by keyword "infrastructure layer"), Choose B. Exclusive D (by keyword "patches for all minor and major database versions for Amazon RDS"). Exclusive E (by keyword "Ensure the physical security of the Amazon RDS"). Easy question.  
upvoted 5 times
-   **kapit** 1 year, 4 months ago

BC & F ( no automatic encryption with direct connect  
upvoted 2 times
-   **TariqKipkemei** 1 year, 4 months ago

Selected Answer: BF

Amazon ECS is a fully managed service, the ops team only focus on building their applications, not the environment. Only option B and F makes sense.  
upvoted 1 times
-   **pentium75** 10 months, 1 week ago

Plus C (we were asked for three). Configuration (!) of components for monitoring, log management etc.; those services exist from AWS but you need to configure them (which logs do you want to store for how long etc.).



upvoted 2 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: BCF**

100% BCF.

upvoted 2 times

  **lucdt4** 1 year, 5 months ago

**Selected Answer: BCF**

BCF

B: Mentioned RDS

C: Mentioned ECS

F: Mentioned Direct connect

upvoted 5 times

  **hiroohiroo** 1 year, 5 months ago

**Selected Answer: BCF**


Yes BCF

upvoted 2 times

  **omoakin** 1 year, 5 months ago

I agree BCF

upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: BCF**

Bcf for me

upvoted 3 times

A company runs a Java-based job on an Amazon EC2 instance. The job runs every hour and takes 10 seconds to run. The job runs on a scheduled interval and consumes 1 GB of memory. The CPU utilization of the instance is low except for short surges during which the job uses the maximum CPU available. The company wants to optimize the costs to run the job.

Which solution will meet these requirements?

- A. Use AWS App2Container (A2C) to containerize the job. Run the job as an Amazon Elastic Container Service (Amazon ECS) task on AWS Fargate with 0.5 virtual CPU (vCPU) and 1 GB of memory.
- B. Copy the code into an AWS Lambda function that has 1 GB of memory. Create an Amazon EventBridge scheduled rule to run the code each hour.
- C. Use AWS App2Container (A2C) to containerize the job. Install the container in the existing Amazon Machine Image (AMI). Ensure that the schedule stops the container when the task finishes.
- D. Configure the existing schedule to stop the EC2 instance at the completion of the job and restart the EC2 instance when the next job starts.

**Correct Answer: B**

*Community vote distribution*

B (100%)

  **omarshaban** Highly Voted 1 year, 3 months ago

THIS WAS IN MY EXAM

upvoted 9 times

  **awsgeek75** Highly Voted 1 year, 3 months ago

**Selected Answer: B**

Never done it myself but apparently you can run Java in Lambda all the way to latest version  
<https://docs.aws.amazon.com/lambda/latest/dg/lambda-java.html>

upvoted 6 times

  **Danilus** Most Recent 5 months, 4 weeks ago

**Selected Answer: B**

key -The company wants to optimize the costs to run the job.

the answer is B because lambda provides the most cost-effective,easy to mange and scalable solution.it minimizes the need to manage infrastructure because you only pay for the compute time used during job execution

upvoted 2 times

  **noircesar25** 1 year, 2 months ago

can someone explain what makes A wrong, im aware that C hasnt covered all the requirements but A seems good with fargate serverless and autoscaling functionalities, plus AWS app2container is for .NET and JAVA

upvoted 4 times

  **bujuman** 1 year, 1 month ago

Statement: A company runs a Java-based job on an Amazon EC2 instance

Requirement: The company wants to optimize the costs to run the job

Regarding option A: App2Container is more likely for migratiing legacy application to conatiner based application.

Which is not the main purpose of this use case.

We are asked to reduce cost on a application that is already running under EC2 instance.

So option B has the hight weight, cause lambda could perfectly do the job with the minimal cost

upvoted 4 times

  **JA2018** 5 months, 1 week ago

you only need to pay for the compute resources consumed as and when each Lambda function is actually running (for the running duration)

upvoted 1 times

  **Murtadhaceit** 1 year, 4 months ago

**Selected Answer: B**

This question is intended for Lambda. Just searched for Lambda with Event bridge. I


upvoted 3 times

  **potomac** 1 year, 6 months ago

**Selected Answer: B**

Lambda allows you to allocate memory for your functions in increments of 1 MB, ranging from a minimum of 128 MB to a maximum of 10,240 MB (10 GB).



upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Remember - AWS Lambda function can go up to 10 GB of memory, instead of free tier only allow 512MB.

upvoted 6 times

  **james2033** 1 year, 9 months ago

**Selected Answer: B**

"AWS Batch jobs as EventBridge targets" at <https://docs.aws.amazon.com/batch/latest/userguide/batch-cwe-target.html>

AWS Batch + Amazon EventBridge <https://docs.aws.amazon.com/batch/latest/userguide/batch-cwe-target.html> .

AWS Lambda just for a point of time per period. Choose B.

upvoted 2 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: B**

10 seconds to run, optimize the costs, consumes 1 GB of memory = AWS Lambda function.

upvoted 3 times

  **alexandercamachop** 1 year, 11 months ago

**Selected Answer: B**

AWS Lambda automatically scales resources to handle the workload, so you don't have to worry about managing the underlying infrastructure. It provisions the necessary compute resources based on the configured memory size (1 GB in this case) and executes the job in a serverless environment.

By using Amazon EventBridge, you can create a scheduled rule to trigger the Lambda function every hour, ensuring that the job runs on the desired interval.

upvoted 2 times

  **Yadav\_Sanjay** 1 year, 11 months ago

**Selected Answer: B**



B - Within 10 sec and 1 GB Memory (Lambda Memory 128MB to 10GB)

upvoted 3 times

  **Yadav\_Sanjay** 1 year, 11 months ago

<https://docs.aws.amazon.com/lambda/latest/operatorguide/computing-power.html>

upvoted 2 times

  **Efren** 1 year, 11 months ago

**Selected Answer: B**

Agreed, B Lambda

upvoted 3 times

A company wants to implement a backup strategy for Amazon EC2 data and multiple Amazon S3 buckets. Because of regulatory requirements, the company must retain backup files for a specific time period. The company must not alter the files for the duration of the retention period.

Which solution will meet these requirements?

- A. Use AWS Backup to create a backup vault that has a vault lock in governance mode. Create the required backup plan.
- B. Use Amazon Data Lifecycle Manager to create the required automated snapshot policy.
- C. Use Amazon S3 File Gateway to create the backup. Configure the appropriate S3 Lifecycle management.
- D. Use AWS Backup to create a backup vault that has a vault lock in compliance mode. Create the required backup plan.

Correct Answer: D

Community vote distribution

D (100%)

- Efren

Highly Voted

1 year, 11 months ago

D, Governance is like the government, they can do things you cannot , like delete files or backups :D Compliance, nobody can!

upvoted 39 times
- cmbt

1 year, 10 months ago

Finally I understood!

upvoted 4 times
- joshnort

1 year, 10 months ago

Great analogy

upvoted 8 times
- nosense

Highly Voted

1 year, 11 months ago

Selected Answer: D

D bcs in governance we can delete backup

upvoted 5 times
- Danilus

Most Recent

5 months, 4 weeks ago

Selected Answer: D

key-The company must not alter the files for the duration of the retention period  
the answer is D because backup with vault lock in compliance mode make sure that the files when are locked they cannot be deleted or modified  
for the specified retention period it makes the files immutable  
why is not A? because the governance mode... with this mode backups allows flexibility to alter or delete backups

upvoted 2 times
- f2e2419

1 year, 3 months ago

Selected Answer: D

D. Use AWS Backup to create a backup vault that has a vault lock in compliance mode. Create the required backup plan

upvoted 3 times
- Guru4Cloud

1 year, 8 months ago

Selected Answer: D

D. Use AWS Backup to create a backup vault that has a vault lock in compliance mode. Create the required backup plan

upvoted 3 times
- Guru4Cloud

1 year, 8 months ago

Selected Answer: D

Use AWS Backup to create a backup vault that has a vault lock in compliance mode. Create the required backup plan

upvoted 3 times
- ccat91

1 year, 9 months ago

Selected Answer: D

Compliance mode

upvoted 2 times
- TariqKipkemei

1 year, 10 months ago

Selected Answer: D

Must not alter the files for the duration of the retention period = Compliance Mode



upvoted 2 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: D**

D for sure.

upvoted 2 times

  **dydzah** 1 year, 11 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/aws-backup/latest/devguide/vault-lock.html>

upvoted 3 times

  **cloudenthusiast** 1 year, 11 months ago

**Selected Answer: D**

compliance mode

upvoted 4 times

A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers a previous employee did not provide details about the resources inventory. The solutions architect needs to build and map the relationship details of the various workloads across all accounts.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report.
- B. Use AWS Step Functions to collect workload details. Build architecture diagrams of the workloads manually.
- C. Use Workload Discovery on AWS to generate architecture diagrams of the workloads.
- D. Use AWS X-Ray to view the workload details. Build architecture diagrams with relationships.

**Correct Answer: C**

*Community vote distribution*

C (97%)

- Danilus** 5 months, 4 weeks ago

Key - Build and map the relationship details of the various workloads.  
Key - MOST operationally efficient way.  
The answer is C because Workload Discovery is a tool that allows you to find automatically the workloads on the cloud, giving you a visual representation of the infrastructure and the relationship between the resources through different accounts.

upvoted 2 times
- zinabu** 1 year ago

workload discovery=architecture diagram

upvoted 3 times
- osmk** 1 year, 3 months ago

<https://docs.aws.amazon.com/solutions/latest/workload-discovery-on-aws/solution-overview.html>Workload Discovery on AWS is a visualization tool that automatically generates architecture diagrams of your workload on AWS. You can use this solution to build, customize, and share detailed workload visualizations based on live data from AWS

upvoted 3 times
- awsgeek75** 1 year, 3 months ago

**Selected Answer: C**

A: Systems Manager Inventory -> Metadata  
B: Not possible (correct me if I'm wrong)  
D: X-Ray is for application debugging  
C: Workload Discovery is purpose built tool for this type of usage

upvoted 4 times
- NayeraB** 1 year, 2 months ago

Even if B is possible, it has "manually" in it which we won't do because we're lazy in this question

upvoted 3 times
- potomac** 1 year, 6 months ago

**Selected Answer: C**

Workload Discovery on AWS (formerly called AWS Perspective) is a tool to visualize AWS Cloud workloads. Use Workload Discovery on AWS to build, customize, and share detailed architecture diagrams of your workloads based on live data from AWS.

upvoted 3 times
- TariqKipkemei** 1 year, 6 months ago

**Selected Answer: C**

use Workload Discovery on AWS

upvoted 3 times
- Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

Workload Discovery is purpose-built to automatically generate visual mappings of architectures across accounts and Regions. This makes it the most operationally efficient way to meet the requirements.

upvoted 4 times
- MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: C**

Option A: AWS SSM offers "Software inventory": Collect software catalog and configuration for your instances.  
Option C: Workload Discovery on AWS: is a tool for maintaining an inventory of the AWS resources across your accounts and various Regions and mapping relationships between them, and displaying them in a web UI.  
upvoted 4 times

  **DrWatson** 1 year, 11 months ago

**Selected Answer: A**

<https://aws.amazon.com/blogs/mt/visualizing-resources-with-workload-discovery-on-aws/>  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

That is C  
upvoted 2 times

  **Abrar2022** 1 year, 11 months ago

**Selected Answer: C**

AWS Workload Discovery - create diagram, map and visualise AWS resources across AWS accounts and Regions  
upvoted 3 times

  **Abrar2022** 1 year, 11 months ago

Workload Discovery on AWS can map AWS resources across AWS accounts and Regions and visualize them in a UI provided on the website.  
upvoted 2 times

  **hiroohiroo** 1 year, 11 months ago

**Selected Answer: C**

[https://aws.amazon.com/jp/builders-flash/202209/workload-discovery-on-aws/?awsf.filter-name=\\*all](https://aws.amazon.com/jp/builders-flash/202209/workload-discovery-on-aws/?awsf.filter-name=*all)  
upvoted 3 times

  **omoakin** 1 year, 11 months ago

Only C makes sense  
upvoted 3 times

  **cloudenthusiast** 1 year, 11 months ago

**Selected Answer: C**



Workload Discovery on AWS is a service that helps visualize and understand the architecture of your workloads across multiple AWS accounts and Regions. It automatically discovers and maps the relationships between resources, providing an accurate representation of the architecture.  
upvoted 3 times

  **Efren** 1 year, 11 months ago

Not sure here tbh

To efficiently build and map the relationship details of various workloads across multiple AWS Regions and accounts, you can use the AWS Systems Manager Inventory feature in combination with AWS Resource Groups. Here's a solution that can help you achieve this:

AWS Systems Manager Inventory:  
upvoted 1 times

  **nosense** 1 year, 11 months ago

**Selected Answer: C**

only c mapping relationships  
upvoted 2 times

A company uses AWS Organizations. The company wants to operate some of its AWS accounts with different budgets. The company wants to receive alerts and automatically prevent provisioning of additional resources on AWS accounts when the allocated budget threshold is met during a specific period.

Which combination of solutions will meet these requirements? (Choose three.)

- A. Use AWS Budgets to create a budget. Set the budget amount under the Cost and Usage Reports section of the required AWS accounts.
- B. Use AWS Budgets to create a budget. Set the budget amount under the Billing dashboards of the required AWS accounts.
- C. Create an IAM user for AWS Budgets to run budget actions with the required permissions.
- D. Create an IAM role for AWS Budgets to run budget actions with the required permissions.
- E. Add an alert to notify the company when each account meets its budget threshold. Add a budget action that selects the IAM identity created with the appropriate config rule to prevent provisioning of additional resources.
- F. Add an alert to notify the company when each account meets its budget threshold. Add a budget action that selects the IAM identity created with the appropriate service control policy (SCP) to prevent provisioning of additional resources.

**Correct Answer:** BDF

*Community vote distribution*

BDF (71%)

ADF (15%)

12%

- vesen22**

Highly Voted

1 year, 11 months ago

Selected Answer: BDF

I don't see why adf has the most voted when almost everyone has chosen bdf, smh  
[https://acloudguru.com/videos/acg-fundamentals/how-to-set-up-an-aws-billing-and-budget-alert?utm\\_source=google&utm\\_medium=paid-search&utm\\_campaign=cloud-transformation&utm\\_term=ssi-global-acg-core-dsa&utm\\_content=free-trial&gclid=Cj0KCQjwmtGjBhDhARIsAEqfDEcDfXdLul2NxgSMxKracIITZimWOtDBRpsJPpx8IS9T4NndKhbUqPlaAlzhEALw\\_wcB](https://acloudguru.com/videos/acg-fundamentals/how-to-set-up-an-aws-billing-and-budget-alert?utm_source=google&utm_medium=paid-search&utm_campaign=cloud-transformation&utm_term=ssi-global-acg-core-dsa&utm_content=free-trial&gclid=Cj0KCQjwmtGjBhDhARIsAEqfDEcDfXdLul2NxgSMxKracIITZimWOtDBRpsJPpx8IS9T4NndKhbUqPlaAlzhEALw_wcB)  
upvoted 14 times
- omarshaban**

Highly Voted

1 year, 3 months ago

IN MY EXAM

upvoted 9 times
- FlyingHawk**

Most Recent

4 months, 4 weeks ago

Selected Answer: B

AWS budget is under the Budget and Planning section in Billing and Cost Management cosole  
upvoted 1 times
- FlyingHawk**

4 months, 4 weeks ago

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/view-billing-dashboard.html>  
upvoted 1 times
- TariqKipkemei**

1 year, 6 months ago

Selected Answer: DF

Its 11/Nov/2023. Options D&F are definitely required.  
As for the budget, right from the aws console, the only place to set this up is:  
AWS Billing>Cost Management>Budgets.  
upvoted 5 times
- Guru4Cloud**

1 year, 8 months ago

Selected Answer: BDF

How to create a budget:  
Billing console > budget > create budget!  
upvoted 6 times
- Chris22usa**

1 year, 10 months ago

ACF:  
Option B is incorrect because the budget amount should be set under the Cost and Usage Reports section, not the Billing dashboards.  
upvoted 1 times
- pentium75**

1 year, 4 months ago

"Create an AWS Budget: Go to the AWS Billing Dashboard"

<https://awslabs.github.io/scale-out-computing-on-aws/workshops/TKO-Scale-Out-Computing/modules/071-budgets/>  
upvoted 1 times

  **Abrar2022** 1 year, 11 months ago

**Selected Answer: BDF**

How to create a budget:  
Billing console > budget > create budget!  
upvoted 3 times

  **udo2020** 1 year, 11 months ago

It is BDF because there is actually a Billing Dashboard available.  
upvoted 8 times

  **hiroohiroo** 1 year, 11 months ago

**Selected Answer: BDF**

[https://docs.aws.amazon.com/ja\\_jp/awsaccountbilling/latest/aboutv2/view-billing-dashboard.html](https://docs.aws.amazon.com/ja_jp/awsaccountbilling/latest/aboutv2/view-billing-dashboard.html)  
upvoted 6 times

  **y0** 1 year, 11 months ago

BDF - Budgets can be set from the billing dashboard in AWS console  
upvoted 3 times


  **cloudenthusiast** 1 year, 11 months ago

**Selected Answer: ADF**

Currently, AWS does not have a specific feature called "AWS Billing Dashboards."  
upvoted 6 times

  **[Removed]** 1 year, 11 months ago

<https://awslabs.github.io/scale-out-computing-on-aws/workshops/TKO-Scale-Out-Computing/modules/071-budgets/>  
upvoted 4 times

  **Efren** 1 year, 11 months ago

if im not wrong, those are correct  
upvoted 3 times

A company runs applications on Amazon EC2 instances in one AWS Region. The company wants to back up the EC2 instances to a second Region. The company also wants to provision EC2 resources in the second Region and manage the EC2 instances centrally from one AWS account.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a disaster recovery (DR) plan that has a similar number of EC2 instances in the second Region. Configure data replication.
- B. Create point-in-time Amazon Elastic Block Store (Amazon EBS) snapshots of the EC2 instances. Copy the snapshots to the second Region periodically.
- C. Create a backup plan by using AWS Backup. Configure cross-Region backup to the second Region for the EC2 instances.
- D. Deploy a similar number of EC2 instances in the second Region. Use AWS DataSync to transfer the data from the source Region to the second Region.



**Correct Answer: C**


*Community vote distribution*

C (63%)

B (25%)

13%



-   **cloudenthusiast**


Highly Voted 

 1 year, 11 months ago

Selected Answer: C

Using AWS Backup, you can create backup plans that automate the backup process for your EC2 instances. By configuring cross-Region backup, you can ensure that backups are replicated to the second Region, providing a disaster recovery capability. This solution is cost-effective as it leverages AWS Backup's built-in features and eliminates the need for manual snapshot management or deploying and managing additional EC2 instances in the second Region.

upvoted 8 times
-   **Guru4Cloud**

Highly Voted 



 1 year, 8 months ago

Selected Answer: C



C is the most cost-effective solution that meets all the requirements.


AWS Backup provides automated backups across Regions for EC2 instances. This handles the backup requirement.

AWS Backup is more cost-effective for cross-Region EC2 backups than using EBS snapshots manually or DataSync.

upvoted 5 times
-   **wizcloudifa** 1 year ago

Well Datasync is to transfer data from on-prem to AWS rather than AWS R1 to R2, and it is use for a one-time migration rather than repetitive data migration activities like backups hence its out(on so many levels)



upvoted 2 times
-   **elmyth**

Most Recent 

 5 months, 4 weeks ago



Selected Answer: B

We have implemented solution similar to B, it's really small and simple script, so we don't have much manual work here, and it's really the most cost-effective.



upvoted 3 times
-   **Certibanksman** 8 months, 1 week ago

Selected Answer: B

Option B (EBS snapshots with cross-Region copy) is the most cost-effective solution for backing up EC2 instances to a second Region while allowing for centralized management and easy recovery when needed.



upvoted 3 times
-   **Daminaij** 8 months, 2 weeks ago

The answer is c

upvoted 1 times
-   **bogobob** 1 year, 5 months ago

Selected Answer: D

How does AWS Backup address that "The company also wants to provision EC2 resources in the second Region"?

upvoted 3 times
-   **NSA\_Poker** 11 months, 1 week ago

With AWS Backup you can 'back up the EC2 instances to a second Region' by implementing cross-Region backup & 'provision EC2 resources in the second Region' by restoring the backup using the AWS Backup console.

<https://docs.aws.amazon.com/aws-backup/latest/devguide/restore-resource.html>  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

How do A, B or D address that? They want to "provision EC2 resources", nobody says that this should be copies of the existing servers. And if it should be copies of the existing servers, wouldn't we need the same (not "a similar") number of servers? We have no idea how many applications on how many servers they have.  
upvoted 3 times



  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: C**



AWS backup  
upvoted 2 times

  **omoakin** 1 year, 11 months ago

CCCCC  
. Create a backup plan by using AWS Backup. Configure cross-Region backup to the second Region for the EC2 instances.  
upvoted 1 times

  **Blingy** 1 year, 11 months ago

CCCCCCC  
upvoted 2 times

  **Efren** 1 year, 11 months ago

C, i would say same, always AWS Backup  
upvoted 1 times

A company that uses AWS is building an application to transfer data to a product manufacturer. The company has its own identity provider (IdP). The company wants the IdP to authenticate application users while the users use the application to transfer data. The company must use Applicability Statement 2 (AS2) protocol.

Which solution will meet these requirements?

- A. Use AWS DataSync to transfer the data. Create an AWS Lambda function for IdP authentication.
- B. Use Amazon AppFlow flows to transfer the data. Create an Amazon Elastic Container Service (Amazon ECS) task for IdP authentication.
- C. Use AWS Transfer Family to transfer the data. Create an AWS Lambda function for IdP authentication.
- D. Use AWS Storage Gateway to transfer the data. Create an Amazon Cognito identity pool for IdP authentication.

Correct Answer: C

Community vote distribution

C (91%) 9%

  **TariqKipkemei** Highly Voted 1 year, 10 months ago

Selected Answer: C

Option C stands out stronger because AWS Transfer Family securely scales your recurring business-to-business file transfers to AWS Storage services using SFTP, FTPS, FTP, and AS2 protocols.

And AWS Lambda can be used to authenticate users with the company's IdP.

upvoted 10 times

  **baba365** 1 year, 10 months ago

Ans : C

To authenticate your users, you can use your existing identity provider with AWS Transfer Family. You integrate your identity provider using an AWS Lambda function, which authenticates and authorizes your users for access to Amazon S3 or Amazon Elastic File System (Amazon EFS).

<https://docs.aws.amazon.com/transfer/latest/userguide/custom-identity-provider-users.html>

upvoted 6 times

  **Danilus** Most Recent 5 months, 4 weeks ago

Selected Answer: C

key-The company must use Applicability Statement 2 (AS2) protocol

The only service in AWS that accepts this protocol is AWS Transfer Family, which provides a managed solution for secure file transfers using AS2, along with essential features such as encryption, digital signatures, and message disposition notifications (MDN)

upvoted 2 times

  **zinabu** 1 year ago

aws transfer family for data transfer and lamda function for idp authentication

upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

Selected Answer: C

<https://aws.amazon.com/about-aws/whats-new/2022/07/aws-transfer-family-support-applicability-statement-2-as2/>

upvoted 4 times

  **potomac** 1 year, 6 months ago

Selected Answer: C

To authenticate your users, you can use your existing identity provider with AWS Transfer Family. You integrate your identity provider using an AWS Lambda function, which authenticates and authorizes your users for access to Amazon S3 or Amazon Elastic File System (Amazon EFS).

upvoted 3 times

  **potomac** 1 year, 6 months ago

Selected Answer: C

Applicability Statement 2 (AS2) is a business-to-business (B2B) messaging protocol used to exchange Electronic Data Interchange (EDI) documents. With AWS Transfer Family's AS2 capabilities, you can securely exchange AS2 messages at scale while maintaining compliance and interoperability with your trading partners.

upvoted 4 times

  **thanhnv142** 1 year, 6 months ago

D is ok

upvoted 1 times

- hsinchang

1 year, 9 months ago

its own IdP -> Lambda  
upvoted 3 times
- dydzah

1 year, 11 months ago

Selected Answer: C

<https://docs.aws.amazon.com/transfer/latest/userguide/custom-identity-provider-users.html>  
upvoted 2 times
- examtopictempacc

1 year, 11 months ago

Selected Answer: C

C is correct. AWS Transfer Family supports the AS2 protocol, which is required by the company. Also, AWS Lambda can be used to authenticate users with the company's IdP, which meets the company's requirement.  
upvoted 3 times

EA100

1 year, 11 months ago

Answer - D  
AS2 is a widely used protocol for secure and reliable data transfer. In this scenario, the company wants to transfer data using the AS2 protocol and authenticate application users using their own identity provider (IdP). AWS Storage Gateway provides a hybrid cloud storage solution that enables data transfer between on-premises environments and AWS.  
  
By using AWS Storage Gateway, you can set up a gateway that supports the AS2 protocol for data transfer. Additionally, you can configure authentication using an Amazon Cognito identity pool. Amazon Cognito provides a comprehensive authentication and user management service that integrates with various identity providers, including your own IdP.  
  
Therefore, Option D is the correct solution as it leverages AWS Storage Gateway for AS2 data transfer and allows authentication using an Amazon Cognito identity pool integrated with the company's IdP.  
upvoted 2 times

deechean

1 year, 8 months ago

AWS Transfer Family also support AS2  
upvoted 2 times

hiroohiroo

1 year, 11 months ago

Selected Answer: C

<https://repost.aws/articles/ARo2ihKKThT2Cue5j6yVUgsQ/articles/ARo2ihKKThT2Cue5j6yVUgsQ/aws-transfer-family-announces-support-for-sending-as2-messages-over-https?>  
upvoted 1 times

omoakin

1 year, 11 months ago

C is correct  
upvoted 1 times

omoakin

1 year, 11 months ago

This is a new Qtn n AS2 is newly supported by AWS Transfer family.....good timing to know ur stuffs.  
upvoted 2 times

nosense

1 year, 11 months ago

Option D looks the better option because it is more secure, scalable, cost-effective, and easy to use than option C.  
upvoted 1 times

cloudenthusiast

1 year, 11 months ago

Selected Answer: D

AWS Storage Gateway supports the AS2 protocol for transferring data. By using AWS Storage Gateway, the company can integrate its own IdP authentication by creating an Amazon Cognito identity pool. Amazon Cognito provides user authentication and authorization capabilities, allowing the company to authenticate application users using its own IdP.  
  
AWS Transfer Family does not currently support the AS2 protocol. AS2 is a specific protocol used for secure and reliable data transfer, often used in business-to-business (B2B) scenarios. In this case, option C, which suggests using AWS Transfer Family, would not meet the requirement of using the AS2 protocol.  
upvoted 3 times

omoakin

1 year, 11 months ago

AWS Transfer Family now supports the Applicability Statement 2 (AS2) protocol, complementing existing protocol support for SFTP, FTPS, and FTP  
upvoted 2 times

y0

1 year, 11 months ago

This is not a case for storage gateway which is more used for a hybrid like environment. Here, to transfer data, we can think or Datasync or Transfer family and considering AS2 protocol, transfer family looks good  
upvoted 3 times

Efren

1 year, 11 months ago



ChatGP

To meet the requirements of using an identity provider (IdP) for user authentication and the AS2 protocol for data transfer, you can implement the following solution:

AWS Transfer Family: Use AWS Transfer Family, specifically AWS Transfer for SFTP or FTPS, to handle the data transfer using the AS2 protocol. AWS Transfer for SFTP and FTPS provide fully managed, highly available SFTP and FTPS servers in the AWS Cloud.

Not sure about Lambda tho

upvoted 3 times

  **Efren** 1 year, 11 months ago

Maybe yes

The Lambda authorizer authenticates the token with the third-party identity provider.

upvoted 2 times

  **cloudenthusiast** 1 year, 11 months ago

Also from ChatGPT

AWS Transfer Family supports multiple protocols, including AS2, and can be used for data transfer. By utilizing AWS Transfer Family, the company can integrate its own IdP authentication by creating an AWS Lambda function.

Both options D and C are valid solutions for the given requirements. The choice between them would depend on additional factors such as specific preferences, existing infrastructure, and overall architectural considerations.

upvoted 2 times

A solutions architect is designing a RESTAPI in Amazon API Gateway for a cash payback service. The application requires 1 GB of memory and 2 GB of storage for its computation resources. The application will require that the data is in a relational format.

Which additional combination of AWS services will meet these requirements with the LEAST administrative effort? (Choose two.)

- A. Amazon EC2
- B. AWS Lambda
- C. Amazon RDS
- D. Amazon DynamoDB
- E. Amazon Elastic Kubernetes Services (Amazon EKS)

**Correct Answer:** *BC*

*Community vote distribution*

BC (89%)11%

- cloudenthusiast**

Highly Voted

1 year, 11 months ago

Selected Answer: BC

"The application will require that the data is in a relational format" so DynamoDB is out. RDS is the choice. Lambda is severless.

upvoted 17 times
- smurfing2k17**

Most Recent

1 year ago

Why cant it be AC? We don't know the time of job runs right?

upvoted 2 times
- JA2018**

5 months, 1 week ago

requirement/ key not stated in stem

upvoted 1 times
- TariqKipkemei**

1 year, 10 months ago

Selected Answer: BC

AWS Lambda and Amazon RDS

upvoted 3 times
- handsonlabsaws**

1 year, 11 months ago

Selected Answer: AC

"2 GB of storage for its COMPUTATION resources" the maximum for Lambda is 512MB.

upvoted 3 times
- PLN6302**

1 year, 8 months ago

Lambda now supports upto 10GB of memory

upvoted 8 times
- Kp88**

1 year, 9 months ago

I thought the same but seems like you can go all the way to 10gb. 512mb is the free tier

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-function-common.html#configuration-ephemeral-storage>

upvoted 3 times
- r3mo**

1 year, 10 months ago

At first I was thinking the same. But the computation memery for the lambda function is 1gb not 2gb. Hence. if you go to basic settings when you create the lambda function you can select a in the memori settings the 1024 MB (1Gb) and that solve the problem.

upvoted 2 times
- Efren**

1 year, 11 months ago

Selected Answer: BC

Relational Data RDS and computing for Lambda

upvoted 4 times
- nosense**

1 year, 11 months ago

bc for me

upvoted 3 times

A company uses AWS Organizations to run workloads within multiple AWS accounts. A tagging policy adds department tags to AWS resources when the company creates tags.

An accounting team needs to determine spending on Amazon EC2 consumption. The accounting team must determine which departments are responsible for the costs regardless of AWS account. The accounting team has access to AWS Cost Explorer for all AWS accounts within the organization and needs to access all reports from Cost Explorer.

Which solution meets these requirements in the MOST operationally efficient way?

- A. From the Organizations management account billing console, activate a user-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.
- B. From the Organizations management account billing console, activate an AWS-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.
- C. From the Organizations member account billing console, activate a user-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by the tag name, and filter by EC2.
- D. From the Organizations member account billing console, activate an AWS-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.

Correct Answer: A

Community vote distribution

A (100%)

- cloudenthusiast

Highly Voted

1 year, 5 months ago

Selected Answer: A

By activating a user-defined cost allocation tag named "department" and creating a cost report in Cost Explorer that groups by the tag name and filters by EC2, the accounting team will be able to track and attribute costs to specific departments across all AWS accounts within the organization. This approach allows for consistent cost allocation and reporting regardless of the AWS account structure.

upvoted 9 times
- luisgu

Highly Voted

1 year, 5 months ago

Selected Answer: A

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html>

upvoted 6 times
- awsgeek75

Most Recent

10 months ago

Selected Answer: A

Management not user.

upvoted 5 times
- Guru4Cloud

1 year, 2 months ago

Selected Answer: A

From the Organizations management account billing console, activate a user-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.

upvoted 5 times
- TariqKipkemei

1 year, 4 months ago

Selected Answer: A

From the Organizations management account billing console, activate a user-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.

upvoted 3 times
- hiroohiroo

1 year, 5 months ago

Selected Answer: A

[https://docs.aws.amazon.com/ja\\_jp/awsaccountbilling/latest/aboutv2/activating-tags.html](https://docs.aws.amazon.com/ja_jp/awsaccountbilling/latest/aboutv2/activating-tags.html)

upvoted 4 times
- nosense

1 year, 5 months ago

Selected Answer: A

a for me

upvoted 3 times





A company wants to securely exchange data between its software as a service (SaaS) application Salesforce account and Amazon S3. The company must encrypt the data at rest by using AWS Key Management Service (AWS KMS) customer managed keys (CMKs). The company must also encrypt the data in transit. The company has enabled API access for the Salesforce account.


- A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3.
- B. Create an AWS Step Functions workflow. Define the task to transfer the data securely from Salesforce to Amazon S3.
- C. Create Amazon AppFlow flows to transfer the data securely from Salesforce to Amazon S3.
- D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3.

**Correct Answer:** C

*Community vote distribution*

C (100%)



-   **cloudenthusiast**


Highly Voted 

 1 year, 11 months ago

Selected Answer: C

Amazon AppFlow is a fully managed integration service that allows you to securely transfer data between different SaaS applications and AWS services. It provides built-in encryption options and supports encryption in transit using SSL/TLS protocols. With AppFlow, you can configure the data transfer flow from Salesforce to Amazon S3, ensuring data encryption at rest by utilizing AWS KMS CMKs.



upvoted 15 times
-   **Guru4Cloud**


Highly Voted 

 1 year, 8 months ago

Selected Answer: C

  - ° Amazon AppFlow can securely transfer data between Salesforce and Amazon S3.
  - ° AppFlow supports encrypting data at rest in S3 using KMS CMKs.
  - ° AppFlow supports encrypting data in transit using HTTPS/TLS.
  - ° AppFlow provides built-in support and templates for Salesforce and S3, requiring less custom configuration than solutions like Lambda, Step Functions, or custom connectors.
  - ° So Amazon AppFlow is the easiest way to meet all the requirements of securely transferring data between Salesforce and S3 with encryption at rest and in transit.



upvoted 9 times
-   **1e22522**

Most Recent 



 9 months ago

Selected Answer: C

i do like myself some appflow flow



upvoted 2 times
-   **zinabu** 1 year ago

SAAS=aws appflow



upvoted 4 times
-   **cvoiceip** 1 year, 3 months ago

Ans : C

Salesforce -----> Amazon AppFlow -----> S3



upvoted 3 times
-   **hsinchang** 1 year, 9 months ago

securely transfer data between Software-as-a-Service (SaaS) applications and AWS -> AppFlow

upvoted 3 times
-   **TariqKipkemei** 1 year, 10 months ago



Selected Answer: C

With Amazon AppFlow automate bi-directional data flows between SaaS applications and AWS services in just a few clicks



upvoted 2 times
-   **DrWatson** 1 year, 11 months ago

Selected Answer: C

<https://docs.aws.amazon.com/appflow/latest/userguide/what-is-appflow.html>

upvoted 3 times
-   **Abrar2022** 1 year, 11 months ago



All you need to know is that AWS AppFlow securely transfers data between different SaaS applications and AWS services

upvoted 3 times
-   **hiroohiroo** 1 year, 11 months ago

**Selected Answer: C**

<https://docs.aws.amazon.com/appflow/latest/userguide/salesforce.html>

upvoted 4 times

  **Efren** 1 year, 11 months ago

**Selected Answer: C**

Saas with another service, AppFlow

upvoted 2 times

A company is developing a mobile gaming app in a single AWS Region. The app runs on multiple Amazon EC2 instances in an Auto Scaling group. The company stores the app data in Amazon DynamoDB. The app communicates by using TCP traffic and UDP traffic between the users and the servers. The application will be used globally. The company wants to ensure the lowest possible latency for all users.

Which solution will meet these requirements?

- A. Use AWS Global Accelerator to create an accelerator. Create an Application Load Balancer (ALB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB.
- B. Use AWS Global Accelerator to create an accelerator. Create a Network Load Balancer (NLB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB.
- C. Create an Amazon CloudFront content delivery network (CDN) endpoint. Create a Network Load Balancer (NLB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB. Update CloudFront to use the NLB as the origin.
- D. Create an Amazon CloudFront content delivery network (CDN) endpoint. Create an Application Load Balancer (ALB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB. Update CloudFront to use the ALB as the origin.

**Correct Answer:** B

Community vote distribution

B (100%)

- zinabu**

Highly Voted

6 months, 4 weeks ago

TCP/UDP/IP based communication with server =NLB  
for global low latency communication if IP/udp/tcp based = aws global accelerator  
upvoted 5 times
- sandordini**

Most Recent

6 months, 2 weeks ago

Selected Answer: B

Mobile gaming, UDP > AWS Global Accelerator + NLB  
upvoted 3 times
- Mikado211**

10 months, 3 weeks ago

Selected Answer: B

UDP == NLB  
NLB can't be used with Cloudfront, so we have to play with AWS Global accelerator  
upvoted 4 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: B

Use AWS Global Accelerator to create an accelerator. Create a Network Load Balancer (NLB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB  
upvoted 3 times
- TariqKipkemei**

1 year, 4 months ago

Selected Answer: B

TCP and UDP = global accelerator and Network Load Balancer  
upvoted 3 times
- antropaws**

1 year, 5 months ago

Selected Answer: B

Clearly B.  
upvoted 2 times
- eddie5049**

1 year, 5 months ago

Selected Answer: B



NLB + Accelerator  
upvoted 4 times
- hiroohiroo**

1 year, 5 months ago

Selected Answer: B

AWS Global Accelerator+NLB

upvoted 4 times

  **Efren** 1 year, 5 months ago

**Selected Answer: B**

UDP, Global Accelerator plus NLB

upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: B**

AWS Global Accelerator is a better solution for the mobile gaming app than CloudFront

upvoted 4 times

A company has an application that processes customer orders. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high the workload does not process orders fast enough.

What should a solutions architect do to write the orders reliably to the database as quickly as possible?

- A. Increase the instance size of the EC2 instance when traffic is high. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
- B. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.
- C. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- D. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limits. Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.

**Correct Answer:** B

Community vote distribution

B (100%)

- cloudenthusiast

Highly Voted

1 year, 5 months ago

Selected Answer: B

By decoupling the write operation from the processing operation using SQS, you ensure that the orders are reliably stored in the queue, regardless of the processing capacity of the EC2 instances. This allows the processing to be performed at a scalable rate based on the available EC2 instances, improving the overall reliability and speed of order processing.

upvoted 12 times
- omarshaban

Highly Voted

9 months, 3 weeks ago

IN MY EXAM

upvoted 9 times
- Guru4Cloud

Most Recent

1 year, 2 months ago

Selected Answer: B

Decoupling the order processing from the application using Amazon SQS and leveraging Auto Scaling to handle the processing of orders based on the workload in the SQS queue is indeed the most efficient and scalable approach. This architecture addresses both reliability and performance concerns during traffic spikes.

upvoted 4 times
- TariqKipkemei

1 year, 4 months ago

Selected Answer: B

Write orders to an Amazon Simple Queue Service (Amazon SQS) queue. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.

upvoted 3 times
- antropaws

1 year, 5 months ago

Selected Answer: B

100% B.

upvoted 3 times
- omoakin

1 year, 5 months ago

BBBBBBBBBB

upvoted 3 times

An IoT company is releasing a mattress that has sensors to collect data about a user's sleep. The sensors will send data to an Amazon S3 bucket. The sensors collect approximately 2 MB of data every night for each mattress. The company must process and summarize the data for each mattress. The results need to be available as soon as possible. Data processing will require 1 GB of memory and will finish within 30 seconds.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Glue with a Scala job
- B. Use Amazon EMR with an Apache Spark script
- C. Use AWS Lambda with a Python script
- D. Use AWS Glue with a PySpark job

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: C**

AWS Lambda charges you based on the number of invocations and the execution time of your function. Since the data processing job is relatively small (2 MB of data), Lambda is a cost-effective choice. You only pay for the actual usage without the need to provision and maintain infrastructure.

upvoted 7 times

  **joechen2023** 1 year, 4 months ago

but the question states "Data processing will require 1 GB of memory and will finish within 30 seconds." so it can't be C as Lambda support maximum 512M

upvoted 2 times

  **BillaRanga** 8 months, 4 weeks ago

Lambda can support upto 10 GB, But 512M is under free tier

upvoted 3 times

  **nilandd44gg** 1 year, 3 months ago

C is valid.

Lambda quotas:

Memory - 128 MB to 10,240 MB, in 1-MB increments.

Note: Lambda allocates CPU power in proportion to the amount of memory configured. You can increase or decrease the memory and CPU power allocated to your function using the Memory (MB) setting. At 1,769 MB, a function has the equivalent of one vCPU.

Function timeout 900 seconds (15 minutes)

4 KB, for all environment variables associated with the function, in aggregate

<https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html>

upvoted 3 times

  **Guru4Cloud** Highly Voted 1 year, 2 months ago

**Selected Answer: C**

The data processing is lightweight, only requiring 1 GB memory and finishing in under 30 seconds. Lambda is designed for short, transient workloads like this.

Lambda scales automatically, invoking the function as needed when new data arrives. No servers to manage.

Lambda has a very low cost. You only pay for the compute time used to run the function, billed in 100ms increments. Much cheaper than provisioning EMR or Glue.

Processing can begin as soon as new data hits the S3 bucket by triggering the Lambda function. Provides low latency.

upvoted 5 times

  **Chiquitabandita** Most Recent 1 year ago

I understand C is a common answer "throw Lambda" seems to be a common theme for questions that need processing under 15 minutes for the test. but in reality, can the other solutions be viable options as well?

upvoted 4 times

  **Mikado211** 10 months, 4 weeks ago

That's the point here, technically all the options are good and will work, but since we are on a small amount of data Lambda will be the cheapest one, usually Glue or EMR will be kept for a big amount of data.

Here is a topic where people did a comparison in comments :

[https://www.reddit.com/r/aws/comments/9umxv1/aws\\_glue\\_vs\\_lambda\\_costbenefit/](https://www.reddit.com/r/aws/comments/9umxv1/aws_glue_vs_lambda_costbenefit/)

upvoted 4 times

  **TariqKipkemei** 1 year ago

**Selected Answer: C**

"processing will require 1 GB of memory and will finish within 30 seconds", perfect for AWS Lambda.

upvoted 4 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: C**

I reckon C, but I would consider other well founded options.

upvoted 3 times

  **nosense** 1 year, 5 months ago

**Selected Answer: C**

c anyway the MOST cost-effectively

upvoted 3 times

A company hosts an online shopping application that stores all orders in an Amazon RDS for PostgreSQL Single-AZ DB instance. Management wants to eliminate single points of failure and has asked a solutions architect to recommend an approach to minimize database downtime without requiring any changes to the application code.

Which solution meets these requirements?

- A. Convert the existing database instance to a Multi-AZ deployment by modifying the database instance and specifying the Multi-AZ option.
- B. Create a new RDS Multi-AZ deployment. Take a snapshot of the current RDS instance and restore the new Multi-AZ deployment with the snapshot.
- C. Create a read-only replica of the PostgreSQL database in another Availability Zone. Use Amazon Route 53 weighted record sets to distribute requests across the databases.
- D. Place the RDS for PostgreSQL database in an Amazon EC2 Auto Scaling group with a minimum group size of two. Use Amazon Route 53 weighted record sets to distribute requests across instances.

**Correct Answer: A**

*Community vote distribution*

A (100%)

🗲️ 👤 **Guru4Cloud** Highly Voted 👍 1 year, 2 months ago

**Selected Answer: A**

A. Convert the existing database instance to a Multi-AZ deployment by modifying the database instance and specifying the Multi-AZ option  
upvoted 6 times

🗲️ 👤 **TariqKipkemei** Highly Voted 👍 1 year, 4 months ago

**Selected Answer: A**

Eliminate single points of failure = Multi-AZ deployment  
upvoted 5 times

🗲️ 👤 **awsgeek75** Most Recent 🕒 9 months, 2 weeks ago

**Selected Answer: A**

A is correct but reason need to be clarified:

<https://aws.amazon.com/blogs/database/best-practices-for-converting-a-single-az-amazon-rds-instance-to-a-multi-az-instance/>

The instance doesn't automatically convert to Multi-AZ immediately. By default it will convert at next maintenance window but you can convert it immediately. Compared to B this is much better. CD are too many changes overall so unsuitable.

upvoted 3 times

🗲️ 👤 **antropaws** 1 year, 5 months ago

**Selected Answer: A**

A) <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZSingleStandby.html#Concepts.MultiAZ.Migrating>

upvoted 2 times

🗲️ 👤 **Abrar2022** 1 year, 5 months ago

"minimize database downtime" so why create a new DB just modify the existing one so no time is wasted.

upvoted 5 times

🗲️ 👤 **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: A**

Compared to other solutions that involve creating new instances, restoring snapshots, or setting up replication manually, converting to a Multi-AZ deployment is a simpler and more streamlined approach with lower overhead.

Overall, option A offers a cost-effective and efficient way to minimize database downtime without requiring significant changes or additional complexities.

upvoted 3 times

🗲️ 👤 **Efren** 1 year, 5 months ago

A for HA, but also read replica can convert itself to master if the master is down... so not sure if C?

upvoted 1 times

🗲️ 👤 **Efren** 1 year, 5 months ago

Sorry, the Route 53 doesnt make sense to sent requests to RR , what if is a write?

upvoted 2 times



**nosense** 1 year, 5 months ago

**Selected Answer: A**

i guess aa

upvoted 3 times

A company is developing an application to support customer demands. The company wants to deploy the application on multiple Amazon EC2 Nitro-based instances within the same Availability Zone. The company also wants to give the application the ability to write to multiple block storage volumes in multiple EC2 Nitro-based instances simultaneously to achieve higher application availability.

Which solution will meet these requirements?

- A. Use General Purpose SSD (gp3) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach
- B. Use Throughput Optimized HDD (st1) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach
- C. Use Provisioned IOPS SSD (io2) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach
- D. Use General Purpose SSD (gp2) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach

**Correct Answer: C**

*Community vote distribution*

C (95%)2%

- potomac**

Highly Voted

12 months ago

Selected Answer: C

Multi-Attach is supported exclusively on Provisioned IOPS SSD (io1 and io2) volumes.  
upvoted 11 times
- awsgeek75**

Highly Voted

9 months, 2 weeks ago

Selected Answer: C

hdd<gp2<gp3<io2  
upvoted 7 times
- master9**

Most Recent

10 months, 1 week ago

Selected Answer: C

AWS IO2 does support Multi-Attach. Multi-Attach allows you to share access to an EBS data volume between up to 16 Nitro-based EC2 instances within the same Availability Zone. Each attached instance has full read and write permission to the shared volume. This feature is intended to make it easier to achieve higher application availability for customers that want to deploy applications that manage storage consistency from multiple writers in shared storage infrastructure. However, please note that Multi-Attach on io2 is available in certain regions only.  
upvoted 6 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: C

C. Use Provisioned IOPS SSD (io2) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach  
upvoted 5 times
- TariqKipkemei**

1 year, 4 months ago

Selected Answer: C

Multi-Attach is supported exclusively on Provisioned IOPS SSD (io1 and io2) volumes.  
  
[https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html#:~:text=Multi%2DAttach%20is%20supported%20exclusively%20on%20Provisioned%20IOPS%20SSD%20\(io1%20and%20io2\)%20volume](https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html#:~:text=Multi%2DAttach%20is%20supported%20exclusively%20on%20Provisioned%20IOPS%20SSD%20(io1%20and%20io2)%20volume)  
S.  
upvoted 3 times
- Axeashes**

1 year, 4 months ago

Multi-Attach is supported exclusively on Provisioned IOPS SSD (io1 and io2) volumes.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html>  
upvoted 3 times
- Uzi\_m**

1 year, 4 months ago

The correct answer is A.  
Currently, Multi Attach EBS feature is supported by gp3 volumes also.  
Multi-Attach is supported for certain EBS volume types, including io1, io2, gp3, st1, and sc1 volumes.  
upvoted 2 times
- Kp88**


1 year, 3 months ago

No , Read this --> <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html#considerations>  
upvoted 4 times
- AshishRocks**



1 year, 5 months ago



Answer should be D  
upvoted 1 times



  **Kp88** 1 year, 3 months ago  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html#considerations>  
upvoted 2 times



  **AshishRocks** 1 year, 5 months ago  
By ChatGPT - Create General Purpose SSD (gp2) volumes: Provision multiple gp2 volumes with the required capacity for your application.  
upvoted 1 times



  **Kp88** 1 year, 3 months ago  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html#considerations>  
upvoted 1 times



  **AshishRocks** 1 year, 5 months ago  
Multi-Attach does not support Provisioned IOPS SSD (io2) volumes. Multi-Attach is currently available only for General Purpose SSD (gp2), Throughput Optimized HDD (st1), and Cold HDD (sc1) EBS volumes.  
upvoted 1 times



  **Abrar2022** 1 year, 5 months ago  
Multi-Attach is supported exclusively on Provisioned IOPS SSD (io1 or io2) volumes.  
upvoted 2 times



  **elmogy** 1 year, 5 months ago  
**Selected Answer: C**  
only io1/io2 supports Multi-Attach  
upvoted 3 times



  **Uzi\_m** 1 year, 4 months ago  
Multi-Attach is supported for certain EBS volume types, including io1, io2, gp3, st1, and sc1 volumes.  
upvoted 1 times



  **Kp88** 1 year, 3 months ago  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html#considerations>  
upvoted 2 times

  **examtopictempacc** 1 year, 5 months ago  
**Selected Answer: C**  
only io1/io2 supports Multi-Attach  
upvoted 3 times

  **Vlad** 1 year, 5 months ago  
**Selected Answer: A**  
Option D suggests using General Purpose SSD (gp2) EBS volumes with Amazon EBS Multi-Attach. While gp2 volumes support multi-attach, gp3 volumes offer a more cost-effective solution with enhanced performance characteristics.  
upvoted 1 times



  **Vlad** 1 year, 5 months ago  
I'm sorry :  
  
Multi-Attach enabled volumes can be attached to up to 16 instances built on the Nitro System that are in the same Availability Zone. Multi-Attach is supported exclusively on Provisioned IOPS SSD (io1 or io2) volumes.  
upvoted 3 times

  **Vlad** 1 year, 5 months ago  
The answer is C:  
upvoted 2 times

  **EA100** 1 year, 5 months ago  
Answer - C  
C. Use Provisioned IOPS SSD (io2) EBS volumes with Amazon Elastic Block Store (Amazon EBS) Multi-Attach.

While both option C and option D can support Amazon EBS Multi-Attach, using Provisioned IOPS SSD (io2) EBS volumes provides higher performance and lower latency compared to General Purpose SSD (gp2) volumes. This makes io2 volumes better suited for demanding and mission-critical applications where performance is crucial.

If the goal is to achieve higher application availability and ensure optimal performance, using Provisioned IOPS SSD (io2) EBS volumes with Multi-Attach will provide the best results.  
upvoted 3 times

  **nosense** 1 year, 5 months ago  
**Selected Answer: C**  
c is right  
Amazon EBS Multi-Attach enables you to attach a single Provisioned IOPS SSD (io1 or io2) volume to multiple instances that are in the same

Availability Zone.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html>  
nothing about gp  
upvoted 3 times

  **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: D**

Given that the scenario does not mention any specific requirements for high-performance or specific IOPS needs, using General Purpose SSD (gp2) EBS volumes with Amazon EBS Multi-Attach (option D) is typically the more cost-effective and suitable choice. General Purpose SSD (gp2) volumes provide a good balance of performance and cost, making them well-suited for general-purpose workloads.  
upvoted 1 times



  **elmogy** 1 year, 5 months ago

the question has not mentioned anything about cost-effective solution.  
only io1/io2 supports Multi-Attach

plus fyi, gp3 is the one gives a good balance of performance and cost. so gp2 is wrong in every way  
upvoted 2 times

  **omoakin** 1 year, 5 months ago

I agree  
General Purpose SSD (gp2) volumes are the most common volume type. They were designed to be a cost-effective storage option for a wide variety of workloads. Gp2 volumes cover system volumes, dev and test environments, and various low-latency apps.  
upvoted 1 times

  **y0** 1 year, 5 months ago

gp2 - IOPS 16000  
Nitro - IOPS 64000 - supported by io2. C is correct  
upvoted 2 times

A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance. New company management wants to ensure the application is highly available.



What should a solutions architect do to meet this requirement?


- A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer
- B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region
- C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application
- D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer

**Correct Answer: A**

*Community vote distribution*

A (100%)



-   **nosense**


Highly Voted 

 11 months, 3 weeks ago

Selected Answer: A

it's A



upvoted 6 times
-   **cloudenthusiast**


Highly Voted 

 11 months, 3 weeks ago

Selected Answer: A

By combining Multi-AZ EC2 Auto Scaling and an Application Load Balancer, you achieve high availability for the EC2 instances hosting your stateless two-tier application.



upvoted 5 times
-   **Guru4Cloud**

Most Recent 

 8 months, 2 weeks ago



Selected Answer: A

A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer

upvoted 2 times
-   **TariqKipkemei** 10 months, 2 weeks ago

Selected Answer: A

Highly available = Multi-AZ EC2 Auto Scaling and Application Load Balancer.

upvoted 3 times
-   **antropaws** 11 months ago

Selected Answer: A

Most likely A.

upvoted 2 times

A company uses AWS Organizations. A member account has purchased a Compute Savings Plan. Because of changes in the workloads inside the member account, the account no longer receives the full benefit of the Compute Savings Plan commitment. The company uses less than 50% of its purchased compute power.

- A. Turn on discount sharing from the Billing Preferences section of the account console in the member account that purchased the Compute Savings Plan.
- B. Turn on discount sharing from the Billing Preferences section of the account console in the company's Organizations management account.
- C. Migrate additional compute workloads from another AWS account to the account that has the Compute Savings Plan.
- D. Sell the excess Savings Plan commitment in the Reserved Instance Marketplace.

Correct Answer: B

Community vote distribution

B (74%)

D (26%)

  **norris81** Highly Voted 1 year, 11 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html>

Sign in to the AWS Management Console and open the AWS Billing console at <https://console.aws.amazon.com/billing/>

Note

Ensure you're logged in to the management account of your AWS Organizations.

upvoted 11 times

  **baba365** Highly Voted 1 year, 7 months ago

So what exactly is the question?

upvoted 7 times

  **awsgeek75** 1 year, 3 months ago

It's an English test on complete the sentence...

upvoted 4 times

  **pentium75** 1 year, 4 months ago

What to do



upvoted 2 times

  **jingen11** Most Recent 6 months, 2 weeks ago

**Selected Answer: B**

I dont think there is a way for us to sell excess Savings plan. Only Selling Reserved instances is possible in marketplace

upvoted 2 times

  **jingen11** 6 months, 2 weeks ago

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html#ri-turn-on-process>

upvoted 2 times

  **Stranko** 1 year, 2 months ago

**Selected Answer: D**

I'd go with D, due to "The company uses less than 50% of its purchased compute power". Like, why are you sharing it between other accounts of the company, if the company itself doesn't need it? If you provisioned too much you can sell the overprovisioned capacity on the market. I'd understand B if it was about the account using about 50% of the plan and other accounts running similar workloads, but no such thing is stated.

upvoted 2 times

  **NayeraB** 1 year, 2 months ago

Option E, Take it out of the salary of the guy who made the decision to purchase an entire compute plan without studying the company's needs.

upvoted 6 times

  **mr123dd** 1 year, 4 months ago

**Selected Answer: D**

in the question, it does not clarify then number of accounts the company has, if they only has one account, I think it is D,

upvoted 1 times

🗨️ 👤 **Mujahid\_1** 1 year, 4 months ago

what are you guys doing  
this section is for discussion not for copy paste  
upvoted 4 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: B**

B, it's a generic Compute Savings Plan that can be used for compute workloads in the other accounts.

A doesn't work, discount sharing must be enabled for all accounts (at least for those that provide and share the discounts).

C is not possible, there's a reason why the workloads are in different accounts.

D would be a last resort if there wouldn't be any other workloads in the own organization, but here are.

upvoted 4 times

🗨️ 👤 **michalf84** 1 year, 7 months ago

**Selected Answer: D**

I saw similar question in older exam one can sell on the market unused capacity

upvoted 1 times

🗨️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

B. Turn on discount sharing from the Billing Preferences section of the account console in the company's Organizations management account

upvoted 3 times

🗨️ 👤 **Lx016** 1 year, 8 months ago

Bro, no need to copy paste the answer that is already written. Need an explanation, I see that you just copy pasting the potential answers without any explanation in each discussion.

upvoted 26 times

🗨️ 👤 **live\_reply\_developers** 1 year, 10 months ago

**Selected Answer: D**

"For example, you might want to sell Reserved Instances after moving instances to a new AWS Region, changing to a new instance type, ending projects before the term expiration, when your business needs change, or if you have unneeded capacity."

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ri-market-general.html>

upvoted 2 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

D would make sense if the company wouldn't have other accounts with workloads. Or if it would be EC2 Savings Plans that would not match the instance types in other accounts. But it's a generic Compute Savings Plan that surely can be used in another account. Thus B.

upvoted 2 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago

I am also confused between B and D as the last part of the question "The company uses less than 50% of its purchased compute power." could imply that the whole company (not just this member account) only uses 50% of the computer power. If they said the member account only uses 50% then it would be clear cut B.

upvoted 2 times

🗨️ 👤 **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: B**

answer is B.

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html#:~:text=choose%20Save.-,Turning%20on%20shared%20reserved%20instances%20and%20Savings%20Plans%20discounts,-You%20can%20use>

upvoted 2 times

🗨️ 👤 **Felix\_br** 1 year, 11 months ago

**Selected Answer: D**

The company uses less than 50% of its purchased compute power.

For this reason i believe D is the best solution : <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ri-market-general.html>

upvoted 3 times

🗨️ 👤 **Abrar2022** 1 year, 11 months ago

The company Organization's management account can turn on/off shared reserved instances.



upvoted 2 times

🗨️ 👤 **cloudenthusiast** 1 year, 11 months ago

**Selected Answer: B**

To summarize, option C (Migrate additional compute workloads from another AWS account to the account that has the Compute Savings Plan) is a valid solution to address the underutilization of the Compute Savings Plan. However, it involves workload migration and may require careful planning and coordination. Consider the feasibility and impact of migrating workloads before implementing this solution.



upvoted 3 times

  **EA100** 1 year, 11 months ago

Answer - C

If a member account within AWS Organizations has purchased a Compute Savings Plan

upvoted 1 times

  **EA100** 1 year, 11 months ago

Asnwer - C

upvoted 1 times

A company is developing a microservices application that will provide a search catalog for customers. The company must use REST APIs to present the frontend of the application to users. The REST APIs must access the backend services that the company hosts in containers in private VPC subnets.

Which solution will meet these requirements?

- A. Design a WebSocket API by using Amazon API Gateway. Host the application in Amazon Elastic Container Service (Amazon ECS) in a private subnet. Create a private VPC link for API Gateway to access Amazon ECS.
- B. Design a REST API by using Amazon API Gateway. Host the application in Amazon Elastic Container Service (Amazon ECS) in a private subnet. Create a private VPC link for API Gateway to access Amazon ECS.
- C. Design a WebSocket API by using Amazon API Gateway. Host the application in Amazon Elastic Container Service (Amazon ECS) in a private subnet. Create a security group for API Gateway to access Amazon ECS.
- D. Design a REST API by using Amazon API Gateway. Host the application in Amazon Elastic Container Service (Amazon ECS) in a private subnet. Create a security group for API Gateway to access Amazon ECS.

**Correct Answer: B**

*Community vote distribution*

B (100%)

- cloudenthusiast**

Highly Voted

1 year, 11 months ago

Selected Answer: B

REST API with Amazon API Gateway: REST APIs are the appropriate choice for providing the frontend of the microservices application. Amazon API Gateway allows you to design, deploy, and manage REST APIs at scale.

Amazon ECS in a Private Subnet: Hosting the application in Amazon ECS in a private subnet ensures that the containers are securely deployed within the VPC and not directly exposed to the public internet.

Private VPC Link: To enable the REST API in API Gateway to access the backend services hosted in Amazon ECS, you can create a private VPC link. This establishes a private network connection between the API Gateway and ECS containers, allowing secure communication without traversing the public internet.

upvoted 14 times
- MNotABot**

Highly Voted

1 year, 10 months ago

Question itself says: "The company must use REST APIs", hence WebSocket APIs are not applicable and such options are eliminated straight away.

upvoted 9 times
- MatAlves**

Most Recent

7 months, 3 weeks ago

Selected Answer: B

"VPC links enable you to create private integrations that connect your HTTP API routes to private resources in a VPC, such as Application Load Balancers or Amazon ECS container-based applications."

upvoted 3 times
- freedafeng**

9 months, 2 weeks ago

I think the connection should be from the application to the ECS in the private VPC, instead of from the API Gateway to the ECS in the private VPC. API Gateway only needs to connect to the application...

upvoted 1 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: B

AC are wrong as they are not REST API

D, you don't make SG for API Gateway to EC2, you have to make a VPC Link. More details at <https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-vpc-links.html>

upvoted 5 times
- Guru4Cloud**

1 year, 8 months ago

Selected Answer: B

To allow the REST APIs to securely access the backend, a private VPC link should be created from API Gateway to the ECS containers. A private VPC link provides private connectivity between API Gateway and the VPC without using public IP addresses or requiring an internet gateway/NAT

upvoted 5 times
- Axeashes**

1 year, 10 months ago

Selected Answer: B

<https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-private-integration.html>  
upvoted 2 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: B**

A VPC link is a resource in Amazon API Gateway that allows for connecting API routes to private resources inside a VPC.  
upvoted 3 times

  **samehpalass** 1 year, 10 months ago



B is the right choice  
upvoted 2 times

  **Yadav\_Sanjay** 1 year, 10 months ago

Why Not D  
upvoted 3 times

  **potomac** 1 year, 6 months ago

A security group acts as a firewall for associated EC2 instances, controlling both inbound and outbound traffic at the instance level.  
upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: B**

b is right, bcs vpc link provided security connection  
upvoted 4 times

A company stores raw collected data in an Amazon S3 bucket. The data is used for several types of analytics on behalf of the company's customers. The type of analytics requested determines the access pattern on the S3 objects.

The company cannot predict or control the access pattern. The company wants to reduce its S3 costs.



Which solution will meet these requirements?


- A. Use S3 replication to transition infrequently accessed objects to S3 Standard-Infrequent Access (S3 Standard-IA)
- B. Use S3 Lifecycle rules to transition objects from S3 Standard to Standard-Infrequent Access (S3 Standard-IA)
- C. Use S3 Lifecycle rules to transition objects from S3 Standard to S3 Intelligent-Tiering
- D. Use S3 Inventory to identify and transition objects that have not been accessed from S3 Standard to S3 Intelligent-Tiering

**Correct Answer: C**

*Community vote distribution*

C (100%)

 **nosense**

Highly Voted 

11 months, 3 weeks ago

Selected Answer: C

S3 Inventory can't to move files to another class  
upvoted 8 times



 **Murtadhaceit**


Highly Voted 

4 months, 3 weeks ago

Selected Answer: C

Unpredictable access pattern = Intelligent-Tiering.  
upvoted 5 times



 **Guru4Cloud**

Most Recent 

8 months, 2 weeks ago

Selected Answer: C



C. Use S3 Lifecycle rules to transition objects from S3 Standard to S3 Intelligent-Tiering  
upvoted 4 times

 **TariqKipkemei**

10 months, 2 weeks ago

Selected Answer: C

Cannot predict access pattern = S3 Intelligent-Tiering.  
upvoted 4 times

 **Efren**

11 months, 3 weeks ago

Selected Answer: C

Not known patterns, Intelligent Tier  
upvoted 4 times

A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet. However the company's security policy states that any external service cannot initiate a connection to the EC2 instances.

What should a solutions architect recommend to resolve this issue?

- A. Create a NAT gateway and make it the destination of the subnet's route table
- B. Create an internet gateway and make it the destination of the subnet's route table
- C. Create a virtual private gateway and make it the destination of the subnet's route table
- D. Create an egress-only internet gateway and make it the destination of the subnet's route table

**Correct Answer: D**

Community vote distribution

D (100%)

  **wRhIH** Highly Voted 1 year, 10 months ago

For exam,  
egress-only internet gateway: IPv6  
NAT gateway: IPv4  
upvoted 54 times

  **MatAlves** 7 months, 3 weeks ago

Good stuff.

"An egress-only internet gateway is for use with IPv6 traffic only. To enable outbound-only internet communication over IPv4, use a NAT gateway instead."

<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>  
upvoted 2 times

  **b82faaf** 1 year, 4 months ago

This is very helpful, thanks.  
upvoted 4 times

  **RDM10** 1 year, 7 months ago

thanks a lot  
upvoted 4 times

  **cloudenthusiast** Highly Voted 1 year, 11 months ago

**Selected Answer: D**

An egress-only internet gateway (EIGW) is specifically designed for IPv6-only VPCs and provides outbound IPv6 internet access while blocking inbound IPv6 traffic. It satisfies the requirement of preventing external services from initiating connections to the EC2 instances while allowing the instances to initiate outbound communications.  
upvoted 9 times

  **cloudenthusiast** 1 year, 11 months ago

Since the company's security policy explicitly states that external services cannot initiate connections to the EC2 instances, using a NAT gateway (option A) would not be suitable. A NAT gateway allows outbound connections from private subnets to the internet, but it does not restrict inbound connections from external sources.  
upvoted 5 times

  **pentium75** 1 year, 4 months ago

"A NAT gateway ... does not restrict inbound connections from external sources." Actually it does, but only for IPv4.  
upvoted 2 times

  **[Removed]** 1 year, 11 months ago

Enable outbound IPv6 traffic using an egress-only internet gateway  
<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>  
upvoted 3 times


  **MatAlves** Most Recent 7 months, 3 weeks ago

**Selected Answer: D**

"An egress-only internet gateway is for use with IPv6 traffic only. To enable outbound-only internet communication over IPv4, use a NAT gateway instead."

<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

D. Create an egress-only internet gateway and make it the destination of the subnet's route table



upvoted 2 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: D**

Outbound traffic only = Create an egress-only internet gateway and make it the destination of the subnet's route table

upvoted 2 times

  **radev** 1 year, 11 months ago

**Selected Answer: D**

Egress-Only internet Gateway

upvoted 4 times

A company is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket. During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.

Which solution will meet these requirements?

- A. Enable S3 Intelligent-Tiering for the S3 bucket
- B. Enable S3 Transfer Acceleration for the S3 bucket
- C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC
- D. Create an interface endpoint for Amazon S3 in the VPC. Associate this endpoint with all route tables in the VPC

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **litos168** Highly Voted 1 year, 3 months ago

Amazon S3 supports both gateway endpoints and interface endpoints. With a gateway endpoint, you can access Amazon S3 from your VPC, without requiring an internet gateway or NAT device for your VPC, and with no additional cost. However, gateway endpoints do not allow access from on-premises networks, from peered VPCs in other AWS Regions, or through a transit gateway. For those scenarios, you must use an interface endpoint, which is available for an additional cost.

upvoted 14 times

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: C**

**Gateway VPC Endpoint:** A gateway VPC endpoint enables private connectivity between a VPC and Amazon S3. It allows direct access to Amazon S3 without the need for internet gateways, NAT devices, VPN connections, or AWS Direct Connect.

**Minimize Internet Traffic:** By creating a gateway VPC endpoint for Amazon S3 and associating it with all route tables in the VPC, the traffic between the VPC and Amazon S3 will be kept within the AWS network. This helps in minimizing data transfer costs and prevents the need for traffic to traverse the internet.

**Cost-Effective:** With a gateway VPC endpoint, the data transfer between the application running in the VPC and the S3 bucket stays within the AWS network, reducing the need for data transfer across the internet. This can result in cost savings, especially when dealing with large amounts of data.

upvoted 8 times

  **cloudenthusiast** 1 year, 5 months ago

Option B (Enable S3 Transfer Acceleration for the S3 bucket) is a feature that uses the CloudFront global network to accelerate data transfers to and from Amazon S3. While it can improve data transfer speed, it still involves traffic traversing the internet and doesn't directly address the goal of minimizing costs and preventing internet traffic whenever possible.

upvoted 3 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

**Selected Answer: C**

<https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/>

A: Storage cost is not described as an issue here

B: Tx Accelerator is for external (global user) traffic acceleration

D: Interface endpoint is on-prem to S3

C: gateway VPC is specifically for S3 to AWS resources

upvoted 4 times

  **dkw2342** 7 months, 4 weeks ago

Interface endpoints are not exclusively for on-prem to S3.

The only reason why option D is wrong is because "Associate this endpoint with all route tables in the VPC" makes no sense.

upvoted 1 times

  **bsbs1234** 1 year, 1 month ago

I think both C&D will work.

But D will have extra cost. So C is correct.

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC

upvoted 2 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: C**



Prevent traffic from traversing the internet = Gateway VPC endpoint for S3.

upvoted 3 times

  **Anmol\_1010** 1 year, 5 months ago

Key word transversing to internet

upvoted 2 times

  **Efren** 1 year, 5 months ago

**Selected Answer: C**

Gateway endpoint for S3

upvoted 3 times

  **nosense** 1 year, 5 months ago

**Selected Answer: C**

vpc endpoint for s3

upvoted 5 times

A company has a mobile chat application with a data store based in Amazon DynamoDB. Users would like new messages to be read with as little latency as possible. A solutions architect needs to design an optimal solution that requires minimal application changes.

Which method should the solutions architect select?

- A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages table. Update the code to use the DAX endpoint.
- B. Add DynamoDB read replicas to handle the increased read load. Update the application to point to the read endpoint for the read replicas.
- C. Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint.
- D. Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB.

**Correct Answer: A**

*Community vote distribution*

A (95%)

5%

- pentium75** 10 months, 1 week ago

Selected Answer: A

B and C do not reduce latency. D would reduce latency but require significant application changes.

upvoted 2 times
- wizcloudifa** 6 months, 2 weeks ago

D would not reduce latency, all the messages are new, they wont be stored in cache and they are unique messages(dynamic content) so Elasticache would be pointless here, its more preferable for static content/frequently accessed content.... A makes perfect sense

upvoted 3 times
- Cyberkayu** 10 months, 2 weeks ago

Selected Answer: C

0 code change @C

ABD. In memory cache, read replica, elasticache. Chat application and content is dynamic, cache will still pull data from prod database

upvoted 1 times
- pentium75** 10 months, 1 week ago

C has 0 codes changes but doesn't address the issue.

upvoted 6 times
- danielmakita** 1 year ago

Would go for A.

Minimal application changes != No application changes

upvoted 2 times
- thanhnv142** 1 year ago

"requires minimal application changes" - Do not choose A because it requires updates of codes.

upvoted 1 times
- thanhnv142** 1 year ago

C is correct

A, B and D all require code changes to the app.

upvoted 1 times
- Guru4Cloud** 1 year, 2 months ago

Selected Answer: A

A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages table. Update the code to use the DAX endpoint.

upvoted 2 times
- haoAWS** 1 year, 4 months ago

Selected Answer: A

Read replica does improve the read speed, but it cannot improve the latency because there is always latency between replicas. So A works and B not work.

upvoted 2 times
- mattcl** 1 year, 4 months ago

C , "requires minimal application changes"

upvoted 1 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: A**

little latency = Amazon DynamoDB Accelerator (DAX) .

upvoted 4 times

  **DrWatson** 1 year, 5 months ago

**Selected Answer: A**

I go with A <https://aws.amazon.com/blogs/mobile/building-a-full-stack-chat-application-with-aws-and-nextjs/> but I have some doubts about this <https://aws.amazon.com/blogs/database/how-to-build-a-chat-application-with-amazon-elasticache-for-redis/>

upvoted 3 times

  **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: A**

Amazon DynamoDB Accelerator (DAX): DAX is an in-memory cache for DynamoDB that provides low-latency access to frequently accessed data. By configuring DAX for the new messages table, read requests for the table will be served from the DAX cache, significantly reducing the latency.

Minimal Application Changes: With DAX, the application code can be updated to use the DAX endpoint instead of the standard DynamoDB endpoint. This change is relatively minimal and does not require extensive modifications to the application's data access logic.



Low Latency: DAX caches frequently accessed data in memory, allowing subsequent read requests for the same data to be served with minimal latency. This ensures that new messages can be read by users with minimal delay.

upvoted 4 times

  **cloudenthusiast** 1 year, 5 months ago


Option B (Add DynamoDB read replicas) involves creating read replicas to handle the increased read load, but it may not directly address the requirement of minimizing latency for new message reads.

upvoted 1 times

  **Efren** 1 year, 5 months ago

Tricky one, in doubt also with B, read replicas.

upvoted 1 times

  **awsgeek75** 10 months ago

Yes it's tricky but least code changes is the tie breaker. DAX has zero code changes.

upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: A**

a is valid

upvoted 3 times

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing, and the company is concerned about a potential increase in cost.

- A. Create an Amazon CloudFront distribution to cache static files at edge locations
- B. Create an Amazon ElastiCache cluster. Connect the ALB to the ElastiCache cluster to serve cached files
- C. Create an AWS WAF web ACL and associate it with the ALB. Add a rule to the web ACL to cache static files
- D. Create a second ALB in an alternative AWS Region. Route user traffic to the closest Region to minimize data transfer costs

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

**Selected Answer: A**

Amazon CloudFront: CloudFront is a content delivery network (CDN) service that caches content at edge locations worldwide. By creating a CloudFront distribution, static content from the website can be cached at edge locations, reducing the load on the EC2 instances and improving the overall performance.

Caching Static Files: Since the website serves static content, caching these files at CloudFront edge locations can significantly reduce the number of requests forwarded to the EC2 instances. This helps to lower the overall cost by offloading traffic from the instances and reducing the data transfer costs.

upvoted 6 times

  **awsgeek75** Highly Voted  10 months ago

**Selected Answer: A**

The problem with this question is that no sane AWS architect will chose any of these options and go for S3 caching. But given the choices, A is the only one which will solve the problem within reasonable cost.

upvoted 5 times

  **Guru4Cloud** Most Recent  1 year, 2 months ago

**Selected Answer: A**

A. Create an Amazon CloudFront distribution to cache static files at edge locations

upvoted 3 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: A**

Serves static content = Amazon CloudFront distribution.

upvoted 3 times

  **nosense** 1 year, 5 months ago

**Selected Answer: A**

a for me

upvoted 3 times

A company has multiple VPCs across AWS Regions to support and run workloads that are isolated from workloads in other Regions. Because of a recent application launch requirement, the company's VPCs must communicate with all other VPCs across all Regions.

Which solution will meet these requirements with the LEAST amount of administrative effort?

- A. Use VPC peering to manage VPC communication in a single Region. Use VPC peering across Regions to manage VPC communications.
- B. Use AWS Direct Connect gateways across all Regions to connect VPCs across regions and manage VPC communications.
- C. Use AWS Transit Gateway to manage VPC communication in a single Region and Transit Gateway peering across Regions to manage VPC communications.
- D. Use AWS PrivateLink across all Regions to connect VPCs across Regions and manage VPC communications

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Felix\_br** Highly Voted 1 year, 5 months ago

The correct answer is: C. Use AWS Transit Gateway to manage VPC communication in a single Region and Transit Gateway peering across Regions to manage VPC communications.

AWS Transit Gateway is a network hub that you can use to connect your VPCs and on-premises networks. It provides a single point of control for managing your network traffic, and it can help you to reduce the number of connections that you need to manage.

Transit Gateway peering allows you to connect two Transit Gateways in different Regions. This can help you to create a global network that spans multiple Regions.

To use Transit Gateway to manage VPC communication in a single Region, you would create a Transit Gateway in each Region. You would then attach your VPCs to the Transit Gateway.

To use Transit Gateway peering to manage VPC communication across Regions, you would create a Transit Gateway peering connection between the Transit Gateways in each Region.

upvoted 25 times

  **TariqKipkemei** 1 year, 4 months ago

thank you for this comprehensive explanation

upvoted 3 times

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: C**

AWS Transit Gateway: Transit Gateway is a highly scalable service that simplifies network connectivity between VPCs and on-premises networks. By using a Transit Gateway in a single Region, you can centralize VPC communication management and reduce administrative effort.

Transit Gateway Peering: Transit Gateway supports peering connections across AWS Regions, allowing you to establish connectivity between VPCs in different Regions without the need for complex VPC peering configurations. This simplifies the management of VPC communications across Regions.

upvoted 8 times

  **awsgeek75** Most Recent 10 months ago

**Selected Answer: C**

C is like a managed solution for A. A can work but with a lot of overhead (CIDR blocks uniqueness requirement). B and D are not the right products

upvoted 2 times

  **potomac** 12 months ago

**Selected Answer: C**

multiple regions + multiple VPCs --> Transit Gateway

upvoted 3 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: C**

Definitely C.

Very well explained by @Felix\_br

upvoted 2 times

  **omoakin** 1 year, 5 months ago

Ccccccccccccccccccc

if you have services in multiple Regions, a Transit Gateway will allow you to access those services with a simpler network configuration.

upvoted 3 times

A company is designing a containerized application that will use Amazon Elastic Container Service (Amazon ECS). The application needs to access a shared file system that is highly durable and can recover data to another AWS Region with a recovery point objective (RPO) of 8 hours. The file system needs to provide a mount target in each Availability Zone within a Region.

A solutions architect wants to use AWS Backup to manage the replication to another Region.

Which solution will meet these requirements?

- A. Amazon FSx for Windows File Server with a Multi-AZ deployment
- B. Amazon FSx for NetApp ONTAP with a Multi-AZ deployment
- C. Amazon Elastic File System (Amazon EFS) with the Standard storage class
- D. Amazon FSx for OpenZFS

Correct Answer: C

Community vote distribution

C (89%)

11%

  **elmogy** Highly Voted 1 year, 5 months ago

Selected Answer: C

<https://aws.amazon.com/efs/faq/>

Q: What is Amazon EFS Replication?

EFS Replication can replicate your file system data to another Region or within the same Region without requiring additional infrastructure or a custom process. Amazon EFS Replication automatically and transparently replicates your data to a second file system in a Region or AZ of your choice. You can use the Amazon EFS console, AWS CLI, and APIs to activate replication on an existing file system. EFS Replication is continual and provides a recovery point objective (RPO) and a recovery time objective (RTO) of minutes, helping you meet your compliance and business continuity goals.



upvoted 13 times

  **wizcloudifa** Most Recent 6 months, 2 weeks ago

Selected Answer: C

The key thing to notice here in the question is "with a recovery point objective (RPO) of 8 hours", as it is 8 hours of time it recovery can be easily managed by EFS, no need to go for costlier and not built for this use-case(share file system) options like NetApp ONTAP(proprietary data cluster), OpenZFS(not a built in filesystem in AWS) or FSx for windows(file system for windows compatible workloads)

upvoted 3 times

  **awsgeek75** 10 months ago

Selected Answer: C

A: ECS is not Windows File Server so won't work

B: ONTAP is proprietary data cluster completely unrelated to this question

D: OpenZFS needs a Linux kind of host for access. Not a built-in filesystem in AWS by default

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Selected Answer: C

"The file system<https://www.examttopics.com/exams/amazon/aws-certified-solutions-architect-associate-saa-c03/view/48/#> needs to provide a mount target in each (!) Availability Zone within a Region", most regions have three AZs, but FSx Multi-AZ provides only nodes "spread across two AZs". While "or Amazon EFS file systems that use Regional storage classes [such as Standard], you can create a mount target in each Availability Zone in an AWS Region."

upvoted 3 times

  **pentium75** 10 months, 1 week ago

Huh, comment has been scrambled a bit. Anyway

FSx Multi-AZ: Mount targets in two AZs

EFS Standard: Can create mount target in each AZ

upvoted 3 times

  **Goutham4981** 11 months, 3 weeks ago

Selected Answer: C

In the absence of this information, we can only make an assumption based on the provided requirements. The requirement for a shared file system that can recover data to another AWS Region with a recovery point objective (RPO) of 8 hours, and the need for a mount target in each Availability Zone within a Region, are all natively supported by Amazon EFS with the Standard storage class.

While Amazon FSx for NetApp ONTAP does provide shared file systems and supports both Windows and Linux, it does not natively support replication to another region through AWS Backup.



upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C. Amazon Elastic File System (Amazon EFS) with the Standard storage class

upvoted 2 times

  **cd93** 1 year, 2 months ago

**Selected Answer: B**

B or C, but since question didn't mention operating system type, I guess we should go with B because it is more versatile (EFS supports Linux only), although ECS containers do support windows instances...

upvoted 2 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: C**

Both option B and C will support this requirement.

<https://aws.amazon.com/efs/faq/#:~:text=What%20is%20Amazon%20EFS%20Replication%3F>

<https://aws.amazon.com/fsx/netapp-ontap/faqs/#:~:text=How%20do%20I%20configure%20cross%2Dregion%20replication%20for%20the%20data%20in%20my%20file%20system%3F>

upvoted 2 times

  **omoakin** 1 year, 5 months ago

BBBBBBBBBBBBBBBB

upvoted 1 times

  **[Removed]** 1 year, 5 months ago

Both B and C are feasible.

Amazon FSx for NetApp ONTAP is just way overpriced for a backup storage solution. The keyword to look out for is sub milli seconds latency  
In real life env, Amazon Elastic File System (Amazon EFS) with the Standard storage class is good enough.


upvoted 4 times

  **Anmol\_1010** 1 year, 5 months ago

Efs, can be mounted only in 1 region

So the answer is B



upvoted 3 times

  **Rob1L** 1 year, 5 months ago

**Selected Answer: C**

C: EFS

upvoted 3 times

  **y0** 1 year, 5 months ago

Selected Answer: C

AWS Backup can manage replication of EFS to another region as mentioned below

<https://docs.aws.amazon.com/efs/latest/ug/awsbackup.html>

upvoted 2 times

  **norris81** 1 year, 5 months ago

<https://aws.amazon.com/efs/faq/>



During a disaster or fault within an AZ affecting all copies of your data, you might experience loss of data that has not been replicated using Amazon EFS Replication. EFS Replication is designed to meet a recovery point objective (RPO) and recovery time objective (RTO) of minutes. You can use AWS Backup to store additional copies of your file system data and restore them to a new file system in an AZ or Region of your choice. Amazon EFS file system backup data created and managed by AWS Backup is replicated to three AZs and is designed for 99.999999999% (11 nines) durability.

upvoted 2 times

  **nosense** 1 year, 5 months ago



Amazon EFS is a scalable and durable elastic file system that can be used with Amazon ECS. However, it does not support replication to another AWS Region.

upvoted 1 times

  **fakrap** 1 year, 5 months ago

To use EFS replication in a Region that is disabled by default, you must first opt in to the Region, so it does support.

upvoted 2 times

  **elmogy** 1 year, 5 months ago

it does support replication to another AWS Region

check the same link you are replying to :/

<https://aws.amazon.com/efs/faq/>

Q: What is Amazon EFS Replication?

EFS Replication can replicate your file system data to another Region or within the same Region without requiring additional infrastructure or a custom process. Amazon EFS Replication automatically and transparently replicates your data to a second file system in a Region or AZ of your choice. You can use the Amazon EFS console, AWS CLI, and APIs to activate replication on an existing file system. EFS Replication is continual and provides a recovery point objective (RPO) and a recovery time objective (RTO) of minutes, helping you meet your compliance and business continuity goals.



upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: B**

shared file system that is highly durable and can recover data

upvoted 2 times

  **Efren** 1 year, 5 months ago

Why not EFS?

upvoted 2 times

A company is expecting rapid growth in the near future. A solutions architect needs to configure existing users and grant permissions to new users on AWS. The solutions architect has decided to create IAM groups. The solutions architect will add the new users to IAM groups based on department.

Which additional action is the MOST secure way to grant permissions to the new users?

- A. Apply service control policies (SCPs) to manage access permissions
- B. Create IAM roles that have least privilege permission. Attach the roles to the IAM groups
- C. Create an IAM policy that grants least privilege permission. Attach the policy to the IAM groups
- D. Create IAM roles. Associate the roles with a permissions boundary that defines the maximum permissions

**Correct Answer: C**

Community vote distribution

C (95%)



5%

  **Rob1L** Highly Voted 1 year, 11 months ago

**Selected Answer: C**

Option B is incorrect because IAM roles are not directly attached to IAM groups.

upvoted 9 times

  **RoroJ** 1 year, 11 months ago

IAM Roles can be attached to IAM Groups:

[https://docs.aws.amazon.com/directoryservice/latest/admin-guide/assign\\_role.html](https://docs.aws.amazon.com/directoryservice/latest/admin-guide/assign_role.html)

upvoted 4 times

  **antropaws** 1 year, 11 months ago

Read your own link: You can assign an existing IAM role to an AWS Directory Service user or group. Not to IAM groups.

upvoted 11 times

  **Efren** Highly Voted 1 year, 11 months ago

**Selected Answer: C**

Agreed with C

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_groups\\_manage\\_attach-policy.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups_manage_attach-policy.html)

Attaching a policy to an IAM user group

upvoted 7 times

  **MatAlves** Most Recent 7 months, 3 weeks ago


**Selected Answer: C**

"Manage access in AWS by creating policies and attaching them to IAM identities (users, groups of users, or roles) or AWS resources."

"An IAM role is an identity within your AWS account that has specific permissions. It's similar to an IAM user, but isn't associated with a specific person."

"IAM roles do not have any permanent credentials associated with them and are instead assumed by IAM users, AWS services, or applications that need temporary security credentials to access AWS resources"

upvoted 2 times

  **MatAlves** 7 months, 3 weeks ago

[https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html)

<https://docs.aws.amazon.com/IAM/latest/UserGuide/id.html>

<https://blog.awsfundamentals.com/aws-iam-roles-terms-concepts-and-examples>

upvoted 2 times

  **zinabu** 1 year ago

create role=for resource like EC2 and lambda ....

create a Policy =for groups or user access policy for the resources like S3 bucket

upvoted 4 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: C**

Not A or D because this is not about restricting maximum permissions, it is about securely granting permissions

Not B because IAM roles are not attached to IAM groups.

C because IAM policies are attached to IAM groups.

upvoted 5 times

  **potomac** 1 year, 6 months ago

**Selected Answer: C**

A is wrong

SCPs are mainly used along with AWS Organizations organizational units (OUs). SCPs do not replace IAM Policies such that they do not provide actual permissions. To perform an action, you would still need to grant appropriate IAM Policy permissions.

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

Create an IAM policy that grants least privilege permission. Attach the policy to the IAM groups

upvoted 2 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: C**

An IAM policy is an object in AWS that, when associated with an identity or resource, defines their permissions. Permissions in the policies determine whether a request is allowed or denied. You manage access in AWS by creating policies and attaching them to IAM identities (users, groups of users, or roles) or AWS resources.

So, option B will also work.

But Since I can only choose one, C would be it.


upvoted 3 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: C**

You can attach up to 10 IAM policy for a 'user group'.

upvoted 2 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: C**

C is the correct one.

upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: B**

should be b

upvoted 2 times

  **imazsyed** 1 year, 11 months ago

it should be C

upvoted 4 times

  **nosense** 1 year, 11 months ago

Option C is not as secure as option B because IAM policies are attached to individual users and cannot be used to manage permissions for groups of users.

upvoted 2 times

  **omoakin** 1 year, 11 months ago

IAM Roles manage who has access to your AWS resources, whereas IAM policies control their permissions. A Role with no Policy attached to it won't have to access any AWS resources. A Policy that is not attached to an IAM role is effectively unused.

upvoted 5 times

  **Clouddon** 1 year, 8 months ago

[https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html)

upvoted 2 times

  **pentium75** 1 year, 4 months ago

IAM roles are not attached to IAM groups.

IAM policies are attached to IAM roles, IAM groups or IAM users. IAM roles are used by services.

upvoted 2 times

A group requires permissions to list an Amazon S3 bucket and delete objects from that bucket. An administrator has created the following IAM policy to provide access to the bucket and applied that policy to the group. The group is not able to delete objects in the bucket. The company follows least-privilege access rules.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:DeleteObject"
      ],
      "Resource": [
        "arn:aws:s3:::bucket-name"
      ],
      "Effect": "Allow"
    }
  ]
}
```

Which statement should a solutions architect add to the policy to correct bucket access?

- "Action": [
 "s3:\*Object"
 ],
 "Resource": [
 "arn:aws:s3:::bucket-name/\*"
 ],
 "Effect": "Allow"
- A.
- "Action": [
 "s3:\*"
 ],
 "Resource": [
 "arn:aws:s3:::bucket-name/\*"
 ],
 "Effect": "Allow"
- B.
- "Action": [
 "s3:DeleteObject"
 ],
 "Resource": [
 "arn:aws:s3:::bucket-name\*"
 ],
 "Effect": "Allow"
- C.
- "Action": [
 "s3:DeleteObject"
 ],
 "Resource": [
 "arn:aws:s3:::bucket-name/\*"
 ],
 "Effect": "Allow"
- D.

**Correct Answer:** D

Community vote distribution

D (100%)

 **Guru4Cloud** Highly Voted 1 year, 2 months ago



**Selected Answer: D**

option B action is S3:\*. this means all actions. The company follows least-privilege access rules. Hence option D  
upvoted 6 times

  **nosense** Highly Voted 1 year, 5 months ago

**Selected Answer: D**

d work  
upvoted 5 times

  **Efren** 1 year, 5 months ago

Agreed  
upvoted 2 times

  **TariqKipkemei** Most Recent 1 year, 4 months ago

**Selected Answer: D**

D is the answer  
upvoted 4 times

  **AncaZalog** 1 year, 4 months ago


what's the difference between B and D? on B the statements are just placed in another order  
upvoted 1 times

  **TariqKipkemei** 1 year, 4 months ago

option B action is S3:\*. this means all actions. The company follows least-privilege access rules. Hence option D  
upvoted 2 times

  **serepetru** 1 year, 5 months ago



What is the difference between C and D?  
upvoted 2 times

  **Ta\_Les** 1 year, 4 months ago

the "/" at the end of the last line on D  
upvoted 6 times

  **sheilawu** 5 months, 1 week ago

so annoying when you oversee this"/"  
upvoted 1 times

  **Rob1L** 1 year, 5 months ago

**Selected Answer: D**

D for sure  
upvoted 2 times

A law firm needs to share information with the public. The information includes hundreds of files that must be publicly readable. Modifications or deletions of the files by anyone before a designated future date are prohibited.

Which solution will meet these requirements in the MOST secure way?

- A. Upload all files to an Amazon S3 bucket that is configured for static website hosting. Grant read-only IAM permissions to any AWS principals that access the S3 bucket until the designated date.
- B. Create a new Amazon S3 bucket with S3 Versioning enabled. Use S3 Object Lock with a retention period in accordance with the designated date. Configure the S3 bucket for static website hosting. Set an S3 bucket policy to allow read-only access to the objects.
- C. Create a new Amazon S3 bucket with S3 Versioning enabled. Configure an event trigger to run an AWS Lambda function in case of object modification or deletion. Configure the Lambda function to replace the objects with the original versions from a private S3 bucket.
- D. Upload all files to an Amazon S3 bucket that is configured for static website hosting. Select the folder that contains the files. Use S3 Object Lock with a retention period in accordance with the designated date. Grant read-only IAM permissions to any AWS principals that access the S3 bucket.

**Correct Answer: B**

*Community vote distribution*

B (100%)

  **Lin878**  10 months, 3 weeks ago

**Selected Answer: B**

Object Lock works only in buckets that have S3 Versioning enabled.  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html>  
upvoted 6 times

  **nosense**  1 year, 11 months ago

**Selected Answer: B**

Option A allows the files to be modified or deleted by anyone with read-only IAM permissions. Option C allows the files to be modified or deleted by anyone who can trigger the AWS Lambda function.  
Option D allows the files to be modified or deleted by anyone with read-only IAM permissions to the S3 bucket  
upvoted 5 times

  **wizcloudifa** 1 year ago

no it doesnt, did you not notice this part: "S3 Object Lock with a retention period in accordance with the designated date", this part avoids deletion/modification of files  
upvoted 2 times

  **MatAlves**  7 months, 3 weeks ago

**Selected Answer: B**

Versioning Enabled + Object Lock = B  
upvoted 2 times

  **potomac** 1 year, 6 months ago

**Selected Answer: B**

S3 bucket policy  
upvoted 4 times

  **thanhnv142** 1 year, 6 months ago



B is correct.  
A doesnot have S3 object lock, but deletion is prohibited, which implies object lock  
C does not have S3 as static web, but have to share the s3 with the public  
D mentions files - but S3 manages objects, not file  
upvoted 2 times

  **hydro143** 1 year, 6 months ago

D?  
Its like B, but also with read-only access limitations for anyone with IAM permissions. Also versioning in B doesn't help with anything.  
upvoted 2 times

  **ManikRoy** 11 months, 4 weeks ago

Enabling versioning is a pre-requisite for object lock.  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Create a new Amazon S3 bucket with S3 Versioning enabled. Use S3 Object Lock with a retention period in accordance with the designated date. Configure the S3 bucket for static website hosting. Set an S3 bucket policy to allow read-only access to the objects.

upvoted 3 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: B**

Create a new Amazon S3 bucket with S3 Versioning enabled. Use S3 Object Lock with a retention period in accordance with the designated date. Configure the S3 bucket for static website hosting. Set an S3 bucket policy to allow read-only access to the objects.



upvoted 4 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: B**

Clearly B.

upvoted 3 times

  **dydzah** 1 year, 11 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html>

upvoted 5 times

A company is making a prototype of the infrastructure for its new website by manually provisioning the necessary infrastructure. This infrastructure includes an Auto Scaling group, an Application Load Balancer and an Amazon RDS database. After the configuration has been thoroughly validated, the company wants the capability to immediately deploy the infrastructure for development and production use in two Availability Zones in an automated fashion.

What should a solutions architect recommend to meet these requirements?

- A. Use AWS Systems Manager to replicate and provision the prototype infrastructure in two Availability Zones
- B. Define the infrastructure as a template by using the prototype infrastructure as a guide. Deploy the infrastructure with AWS CloudFormation.
- C. Use AWS Config to record the inventory of resources that are used in the prototype infrastructure. Use AWS Config to deploy the prototype infrastructure into two Availability Zones.
- D. Use AWS Elastic Beanstalk and configure it to use an automated reference to the prototype infrastructure to automatically deploy new environments in two Availability Zones.

**Correct Answer: B**

*Community vote distribution*

B (84%)

A (16%)

🗳️ 👤 **Guru4Cloud** Highly Voted 👍 1 year, 8 months ago

Just Think Infrastructure as Code=== Cloud Formation  
upvoted 9 times

🗳️ 👤 **nosense** Highly Voted 👍 1 year, 11 months ago

**Selected Answer: B**

b obvious  
upvoted 5 times

🗳️ 👤 **MatAlves** Most Recent ⌚ 7 months, 3 weeks ago

**Selected Answer: A**

The difference between CloudFormation and Beanstalk might be trick, but just for the exam think:

Cloudformation -> Infra as Code  
Beanstalk -> deploy and manage applications  
upvoted 5 times

🗳️ 👤 **LeonSauveterre** 5 months ago

That's right. And I think you've voted for A by mistake. I think you wanna vote for B ;)  
upvoted 1 times

🗳️ 👤 **awsgeek75** 1 year, 3 months ago

**Selected Answer: B**

A: Wrong product  
C: Wrong product  
D: EBS can only handle EC2 so RDS won't be replicated automatically  
B: CloudFormation = IaaS  
upvoted 4 times

🗳️ 👤 **capino** 1 year, 8 months ago

**Selected Answer: B**

Just Think Infrastructure as Code=== Cloud Formation  
upvoted 5 times

🗳️ 👤 **haoAWS** 1 year, 10 months ago

Why D is not correct?  
upvoted 2 times

🗳️ 👤 **Kiki\_Pass** 1 year, 9 months ago

I guess it's because Beanstalk is PaaS (platform as a service) while CloudFormation is IaC (infrastructure as code). The question emphasis more on infrastructure  
upvoted 3 times

🗳️ 👤 **wRhIH** 1 year, 10 months ago

I guess "TEMPLATE" leads to CloudFormation  
upvoted 3 times

  **TariqKipkemei** 1 year, 10 months ago


**Selected Answer: B**

Infrastructure as code = AWS CloudFormation  
upvoted 5 times

  **antropaws** 1 year, 11 months ago

**Selected Answer: B**

Clearly B.  
upvoted 3 times

  **Felix\_br** 1 year, 11 months ago

**Selected Answer: B**

AWS CloudFormation is a service that allows you to define and provision infrastructure as code. This means that you can create a template that describes the resources you want to create, and then use CloudFormation to deploy those resources in an automated fashion.

In this case, the solutions architect should define the infrastructure as a template by using the prototype infrastructure as a guide. The template should include resources for an Auto Scaling group, an Application Load Balancer, and an Amazon RDS database. Once the template is created, the solutions architect can use CloudFormation to deploy the infrastructure in two Availability Zones.

upvoted 4 times

  **omoakin** 1 year, 11 months ago

B

Define the infrastructure as a template by using the prototype infrastructure as a guide. Deploy the infrastructure with AWS CloudFormation  
upvoted 2 times

A business application is hosted on Amazon EC2 and uses Amazon S3 for encrypted object storage. The chief information security officer has directed that no application traffic between the two services should traverse the public internet.

Which capability should the solutions architect use to meet the compliance requirements?

- A. AWS Key Management Service (AWS KMS)
- B. VPC endpoint
- C. Private subnet
- D. Virtual private gateway

Correct Answer: B

Community vote distribution

B (100%)

- cloudenthusiast

Highly Voted

11 months, 3 weeks ago

Selected Answer: B

A VPC endpoint enables you to privately access AWS services without requiring internet gateways, NAT gateways, VPN connections, or AWS Direct Connect connections. It allows you to connect your VPC directly to supported AWS services, such as Amazon S3, over a private connection within the AWS network.  
  
By creating a VPC endpoint for Amazon S3, the traffic between your EC2 instances and S3 will stay within the AWS network and won't traverse the public internet. This provides a more secure and compliant solution, as the data transfer remains within the private network boundaries.  
upvoted 10 times
- Tendu

Most Recent

2 months, 1 week ago

Selected Answer: B

Answer B -VPC Endpoint  
upvoted 1 times
- TariqKipkemei

10 months, 2 weeks ago

Selected Answer: B

Prevent traffic from traversing the internet = VPC endpoint for S3.  
upvoted 4 times
- antropaws

11 months ago

Selected Answer: B

B until proven contrary.  
upvoted 3 times
- handsonlabsaws

11 months ago

Selected Answer: B

B for sure  
upvoted 3 times
- Blingy

11 months, 1 week ago

BBBBBBBBB

upvoted 2 times

A company hosts a three-tier web application in the AWS Cloud. A Multi-AZ Amazon RDS for MySQL server forms the database layer. Amazon ElastiCache forms the cache layer. The company wants a caching strategy that adds or updates data in the cache when a customer adds an item to the database. The data in the cache must always match the data in the database.

Which solution will meet these requirements?

- A. Implement the lazy loading caching strategy
- B. Implement the write-through caching strategy
- C. Implement the adding TTL caching strategy
- D. Implement the AWS AppConfig caching strategy

**Correct Answer: B**

*Community vote distribution*

B (100%)

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

**Selected Answer: B**

In the write-through caching strategy, when a customer adds or updates an item in the database, the application first writes the data to the database and then updates the cache with the same data. This ensures that the cache is always synchronized with the database, as every write operation triggers an update to the cache.

upvoted 29 times



  **cloudenthusiast** 1 year, 5 months ago

Lazy loading caching strategy (option A) typically involves populating the cache only when data is requested, and it does not guarantee that the data in the cache always matches the data in the database.

Adding TTL (Time-to-Live) caching strategy (option C) involves setting an expiration time for cached data. It is useful for scenarios where the data can be considered valid for a specific period, but it does not guarantee that the data in the cache is always in sync with the database.

AWS AppConfig caching strategy (option D) is a service that helps you deploy and manage application configurations. It is not specifically designed for caching data synchronization between a database and cache layer.

upvoted 36 times

  **Kp88** 1 year, 3 months ago

Great explanation , thanks

upvoted 3 times

  **zinabu** Most Recent  6 months, 4 weeks ago

write-through caching strategy

upvoted 2 times

  **dikshya1233** 9 months, 2 weeks ago

In exam

upvoted 4 times

  **awsgeek75** 9 months, 2 weeks ago

**Selected Answer: B**

More helpful reading for why B is the answer:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html#Strategies.WriteThrough>

upvoted 4 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

B. Implement the write-through caching strategy

upvoted 4 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: B**

The answer is definitely B.

I couldn't provide any more details than what has been shared by @cloudenthusiast.

upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: B**

write-through caching strategy updates the cache at the same time as the database  
upvoted 3 times

A company wants to migrate 100 GB of historical data from an on-premises location to an Amazon S3 bucket. The company has a 100 megabits per second (Mbps) internet connection on premises. The company needs to encrypt the data in transit to the S3 bucket. The company will store new data directly in Amazon S3.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the s3 sync command in the AWS CLI to move the data directly to an S3 bucket
- B. Use AWS DataSync to migrate the data from the on-premises location to an S3 bucket
- C. Use AWS Snowball to move the data to an S3 bucket
- D. Set up an IPsec VPN from the on-premises location to AWS. Use the s3 cp command in the AWS CLI to move the data directly to an S3 bucket

**Correct Answer:** B

*Community vote distribution*

B (61%)

A (39%)

🗲️ 👤 **cloudenthusiast** Highly Voted 👍 1 year, 11 months ago

**Selected Answer: B**

AWS DataSync is a fully managed data transfer service that simplifies and automates the process of moving data between on-premises storage and Amazon S3. It provides secure and efficient data transfer with built-in encryption, ensuring that the data is encrypted in transit.

By using AWS DataSync, the company can easily migrate the 100 GB of historical data from their on-premises location to an S3 bucket. DataSync will handle the encryption of data in transit and ensure secure transfer.

upvoted 13 times

🗲️ 👤 **luiscc** Highly Voted 👍 1 year, 11 months ago

**Selected Answer: B**

Using DataSync, the company can easily migrate the 100 GB of historical data to an S3 bucket. DataSync will handle the encryption of data in transit, so the company does not need to set up a VPN or worry about managing encryption keys.

Option A, using the s3 sync command in the AWS CLI to move the data directly to an S3 bucket, would require more operational overhead as the company would need to manage the encryption of data in transit themselves. Option D, setting up an IPsec VPN from the on-premises location to AWS, would also require more operational overhead and would be overkill for this scenario. Option C, using AWS Snowball, could work but would require more time and resources to order and set up the physical device.

upvoted 7 times

🗲️ 👤 **bora4motion** Most Recent ⌚ 3 days, 21 hours ago

**Selected Answer: B**

I will not use cli (A) to send 100GB to s3 from web. This means I will have to stay by computer? or use a different ec2 to login into it and initiate cli from there? hell no!!

B Datasync is the elegant way.

upvoted 1 times

🗲️ 👤 **dollykidurian** 1 month, 3 weeks ago

My counter to all those, picking choice A, you will have to manually re-run the command if on-premise internet goes down however with datasync the service auto resume. (less operational overhead and less monitoring)

upvoted 1 times

🗲️ 👤 **zdi561** 2 months, 4 weeks ago

**Selected Answer: B**

You can also copy or synchronize data from on-premises to AWS manually with the AWS CLI, but doing so requires a fair bit of system tuning see <https://aws.amazon.com/blogs/storage/migrating-and-managing-large-datasets-on-amazon-s3/>

upvoted 1 times

🗲️ 👤 **FlyingHawk** 4 months, 3 weeks ago

**Selected Answer: A**

For A, it is more straightforward approach, with 100Mbps = 12.5MBs, 100 GB = 100 \* 1024 MB = 102,400 MB, 102400 / 12.5 = 8192s = 2.3 hours, so if the network is stable, it should be less operational overhead than B which requires the installation of agent and network configuration, so compared to S3 sync, DataSync requires more upfront configuration but provides more robust, automated transfer capabilities. In an AWS certification exam, they typically expect you to choose the solution that provides the most straightforward, efficient approach with the least operational overhead while meeting all specified requirements.

For this scenario, they would likely expect you to choose A.

upvoted 1 times

🗨️ 👤 **FlyingHawk** 4 months, 3 weeks ago

In the context of AWS exam, AWS may expect we choose B as AWS DataSync is a managed service designed specifically for data migration tasks, A is more ad-hoc approach suit for small data-set. 100 GB will take 2.22 hours, although it is acceptable.

upvoted 2 times

🗨️ 👤 **LeonSauveterre** 5 months ago

**Selected Answer: A**

Option A meets all the requirements with the LEAST operational overhead (while option B will introduce unnecessary complexity for a one-time transfer). The data will be encrypted in transit (via HTTPS), and a 100 Mbps connection is sufficient to transfer data of 100 GB.

If the task involves periodic transfers or error recovery, DataSync (option B) could be considered as an alternative.

upvoted 2 times

🗨️ 👤 **1e22522** 9 months ago

**Selected Answer: B**

Why would u u se the CLI

upvoted 2 times

🗨️ 👤 **bujuman** 1 year, 1 month ago

**Selected Answer: A**

Assertions:

- needs to encrypt the data in transit to the S3 bucket.
- The company will store new data directly in Amazon S3.

Requirements:

- with the LEAST operational overhead

Even Though options A and B could do the job, option A requires VM maintenance because it is not a once-off migration (The company will store new data directly in Amazon S3)

NB: According to me, we must stuck to the question and avoid to interpret

upvoted 3 times

🗨️ 👤 **bujuman** 1 year, 1 month ago

Erratum:

Assertions:

- needs to encrypt the data in transit to the S3 bucket.
- The company will store new data directly in Amazon S3.

Requirements:

- with the LEAST operational overhead

Even Though options A and B could do the job, option B requires VM maintenance because it is not a once-off migration (The company will store new data directly in Amazon S3)

NB: According to me, we must stuck to the question and avoid to interpret

upvoted 1 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: A**

A - one single command, uses encryption automatically

B - Must install, configure and eventually decommission DataSync

C - Overkill

D - No need for VPN

upvoted 5 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago

I agree, A is a million times simpler than B in terms of operational setup. AWS CLI is just one install on a server on client side and one command (literally) to sync the data.

upvoted 3 times

🗨️ 👤 **1rob** 1 year, 4 months ago

**Selected Answer: A**

By default, all data transmitted from the client computer running the AWS CLI and AWS service endpoints is encrypted by sending everything through a HTTPS/TLS connection. You don't need to do anything to enable the use of HTTPS/TLS. It is always enabled unless you explicitly disable it for an individual command by using the --no-verify-ssl command line option.

This is simpler compared to datasync, which will cost operational overhead to configure.

upvoted 2 times

🗨️ 👤 **potomac** 1 year, 6 months ago

**Selected Answer: B**

storage data (including metadata) is encrypted in transit, but how it's encrypted throughout the transfer depends on your source and destination locations.

upvoted 2 times

🗨️ 👤 **thanhnv142** 1 year, 6 months ago

B is correct to migrate

A is incorrect because is it only used to upload minor files (about a few GB) to AWS. 100 GB is not appropriate.

upvoted 2 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago


There is no limitation on AWS CLI s3 sync command transfer size. Not that I can find in the docs.  
<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/s3/sync.html>

Happy to be corrected!  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Use AWS DataSync to migrate the data from the on-premises location to an S3 bucket  
upvoted 4 times

  **HectorLeon2099** 1 year, 9 months ago

**Selected Answer: A**

B is a good option but as the volume is not large and the speed is not bad, A requires less operational overhead  
upvoted 5 times

  **VellaDevil** 1 year, 10 months ago

**Selected Answer: B**

Answer A and B both are correct and with least operational overhead. But since the question says from an "On-premise Location" hence I would go with DataSync.  
upvoted 1 times

  **TariqKipkemei** 1 year, 10 months ago

**Selected Answer: B**

AWS DataSync is a secure, online service that automates and accelerates moving data between on premises and AWS Storage services.  
upvoted 2 times

A company containerized a Windows job that runs on .NET 6 Framework under a Windows container. The company wants to run this job in the AWS Cloud. The job runs every 10 minutes. The job’s runtime varies between 1 minute and 3 minutes.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Lambda function based on the container image of the job. Configure Amazon EventBridge to invoke the function every 10 minutes.
- B. Use AWS Batch to create a job that uses AWS Fargate resources. Configure the job scheduling to run every 10 minutes.
- C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job. Create a scheduled task based on the container image of the job to run every 10 minutes.
- D. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job. Create a standalone task based on the container image of the job. Use Windows task scheduler to run the job every 10 minutes.

**Correct Answer: C**

*Community vote distribution*

C (51%)

B (29%)

A (20%)

- baba365**

Highly Voted

1 year, 7 months ago

Lambda supports only Linux-based container images.

<https://docs.aws.amazon.com/lambda/latest/dg/images-create.html>

upvoted 16 times
- awsgeek75**

1 year, 3 months ago

Not really. Lambda supports .Net 6 directly: <https://aws.amazon.com/blogs/compute/introducing-the-net-6-runtime-for-aws-lambda/>

upvoted 7 times
- AmrFawzy93**

Highly Voted

1 year, 11 months ago

Selected Answer: C

By using Amazon ECS on AWS Fargate, you can run the job in a containerized environment while benefiting from the serverless nature of Fargate, where you only pay for the resources used during the job's execution. Creating a scheduled task based on the container image of the job ensures that it runs every 10 minutes, meeting the required schedule. This solution provides flexibility, scalability, and cost-effectiveness.

upvoted 8 times

**bora4motion**

3 days, 21 hours ago

why not lambda?

upvoted 1 times

**zdi561**

Most Recent

2 months, 4 weeks ago

Selected Answer: A

Lambda is the best for cost , the job is kicked off every 10min and it only takes several min.Set 10min timeout , there is no cost for idling . Batch and ECS cost more for sure.

upvoted 1 times

**zdi561**

3 months, 2 weeks ago

Selected Answer: A

Lambda can be in .net container, and it is the cheaper than any other options.

upvoted 1 times

**LeonSauveterre**

5 months ago

Selected Answer: C

For B: Even though AWS Batch is free, I think it's more suitable for large-scale or relatively complex workloads, not lightweight periodic jobs. So B is not cost-effective (The cost depends on the compute resources you provision to execute the jobs, and maybe "cost-effective" somewhat involves less operational overhead because you have spent your money)

If you require queueing for task prioritization or processing a backlog of tasks, or work on data processing over datasets where Batch excels, AWS Batch would be better!

upvoted 1 times

**muhammadahmer36**

9 months ago

A A A A A A

upvoted 2 times

🗳️ 👤 **zinabu** 1 year ago

selected answer : A

AWS Lambda now supports .NET 6 as both a managed runtime and a container base image

upvoted 5 times

🗳️ 👤 **xBUGx** 1 year, 1 month ago

**Selected Answer: A**

<https://aws.amazon.com/about-aws/whats-new/2022/02/aws-lambda-adds-support-net6/>

upvoted 3 times

🗳️ 👤 **awsgeek75** 1 year, 3 months ago

The question is weirdly phrased for .Net based containers. "A company containerized a Windows job that runs on .NET 6 Framework under a Windows container." This could mean that the job requires .Net 6 Framework OR it could mean the job requires Windows and .Net Framework 6. If the job is just based on .Net 6 then Lambda can run it. I am just a bit cautious about language because other parameters fall under Lambda.

Question may have been wrongly quoted here.

upvoted 4 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: C**

I guess this is an old question from before August 2023, when AWS Batch did not support Windows containers, while ECS already did since September 2021. Thus it would be C, though now B does also work. Since both Batch and ECS are free, we'd pay only for the Fargate resources (which are identical in both cases), now B and C would be correct.

A doesn't work because Lambda still does not support Windows containeres.

D doesn't make sense because the container would have to run 24/7

upvoted 7 times

🗳️ 👤 **ftaws** 1 year, 4 months ago

I think that Batch with Fargate is more cheaper than ECS.

upvoted 1 times

🗳️ 👤 **pentium75** 1 year, 4 months ago

Both Batch and ECS are free.

<https://aws.amazon.com/de/ecs/pricing/>

<https://aws.amazon.com/de/batch/pricing/>

upvoted 2 times

🗳️ 👤 **kt7** 1 year, 5 months ago

**Selected Answer: B**

Batch supports fargate now

upvoted 6 times

🗳️ 👤 **ccmc** 1 year, 6 months ago

**Selected Answer: B**

aws batch supports fargate

upvoted 3 times

🗳️ 👤 **deechean** 1 year, 8 months ago

**Selected Answer: C**

C works. For A, the lambda support container image, but the container image much implement the Lambda Runtime API.

upvoted 1 times

🗳️ 👤 **markoniz** 1 year, 7 months ago

Absolutely agree with this one ... Lambda do not support Windows container, on the other hand ECS is adequate solution

upvoted 3 times

🗳️ 👤 **Hades2231** 1 year, 8 months ago

**Selected Answer: B**

As they support Batch on Fargate now (Aug 2023), the correct answer should be B?

upvoted 3 times

🗳️ 👤 **RDM10** 1 year, 7 months ago

that's exactly my question too.

In one of the discussions, they same lambda is for jobs for 15 min. But for other question, they same batch is the best. I do not understand why we cant use batch?

upvoted 1 times

🗳️ 👤 **Smart** 1 year, 8 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/lambda/latest/dg/csharp-image.html#csharp-image-clients>

upvoted 1 times

  **pentium75** 1 year, 4 months ago

But it's clearly "a Windows job". Lambda does not support Windows containers. (.NET 6 could also run under Linux, but they'd need to modify the container in any case.)

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

C is the most cost-effective solution for running a short-lived Windows container job on a schedule.

Using Amazon ECS scheduled tasks on Fargate eliminates the need to provision EC2 resources. You pay only for the duration the task runs.

Scheduled tasks handle scheduling the jobs and scaling resources automatically. This is lower cost than managing your own scaling via Lambda or Batch.

ECS also supports Windows containers natively unlike Lambda (option A).

Option D still requires provisioning and paying for full time EC2 resources to run a task scheduler even when tasks are not running.

upvoted 2 times

A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture. The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service.



Which combination of actions should a solutions architect recommend to meet these requirements? (Choose two.)


- A. Create a new organization in AWS Organizations with all features turned on. Create the new AWS accounts in the organization.
- B. Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication.
- C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service.
- D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly.
- E. Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

**Correct Answer: AE**

*Community vote distribution*

AE (100%)

 **cloudenthusiast**

Highly Voted 

 11 months, 3 weeks ago



Selected Answer: AE


A. By creating a new organization in AWS Organizations, you can establish a consolidated multi-account architecture. This allows you to create and manage multiple AWS accounts for different business units under a single organization.

E. Setting up AWS IAM Identity Center (AWS Single Sign-On) within the organization enables you to integrate it with the company's corporate directory service. This integration allows for centralized authentication, where users can sign in using their corporate credentials and access the AWS accounts within the organization.

Together, these actions create a centralized, multi-account architecture that leverages AWS Organizations for account management and AWS IAM Identity Center (AWS Single Sign-On) for authentication and access control.

upvoted 11 times

 **Guru4Cloud**

Most Recent 

 8 months, 2 weeks ago



Selected Answer: AE

A) Using AWS Organizations allows centralized management of multiple AWS accounts in a single organization. New accounts can easily be created within the organization.

E) Integrating AWS IAM Identity Center (AWS SSO) with the company's corporate directory enables federated single sign-on. Users can log in once to access accounts and resources across AWS.

Together, Organizations and IAM Identity Center provide consolidated management and authentication for multiple accounts using existing corporate credentials.

upvoted 3 times



 **samehpalass** 10 months, 2 weeks ago

Selected Answer: AE

A:AWS Organization

E:Authentication because option C (SCP) for Authorization



upvoted 4 times

 **baba365** 9 months, 4 weeks ago

Ans: CD

‘centralized corporate directory service’ with new accounts in AWS Organizations

upvoted 2 times

 **TariqKipkemei** 10 months, 2 weeks ago

Selected Answer: AE



Create a new organization in AWS Organizations with all features turned on. Create the new AWS accounts in the organization.

Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

AWS IAM Identity Center (successor to AWS Single Sign-On) helps you securely create or connect your workforce identities and manage their access centrally across AWS accounts and applications.

[https://aws.amazon.com/iam/identity-center/#:~:text=AWS%20IAM%20Identity%20Center%20\(successor%20to%20AWS%20Single%20Sign%2DOn\)%20helps%20you%20securely%20create%20or%20connect%20your%20workforce%20identities%20and%20manage%20their%20access%20centrally%20across%20AWS%20accounts%20and%20applications.](https://aws.amazon.com/iam/identity-center/#:~:text=AWS%20IAM%20Identity%20Center%20(successor%20to%20AWS%20Single%20Sign%2DOn)%20helps%20you%20securely%20create%20or%20connect%20your%20workforce%20identities%20and%20manage%20their%20access%20centrally%20across%20AWS%20accounts%20and%20applications.)

upvoted 2 times

  **nosense** 11 months, 3 weeks ago

ae is right

upvoted 2 times

A company is looking for a solution that can store video archives in AWS from old news footage. The company needs to minimize costs and will rarely need to restore these files. When the files are needed, they must be available in a maximum of five minutes.

What is the MOST cost-effective solution?

- A. Store the video archives in Amazon S3 Glacier and use Expedited retrievals.
- B. Store the video archives in Amazon S3 Glacier and use Standard retrievals.
- C. Store the video archives in Amazon S3 Standard-Infrequent Access (S3 Standard-IA).
- D. Store the video archives in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA).

Correct Answer: A

Community vote distribution

A (97%)

- cloudenthusiast

Highly Voted

1 year, 5 months ago

Selected Answer: A

By choosing Expedited retrievals in Amazon S3 Glacier, you can reduce the retrieval time to minutes, making it suitable for scenarios where quick access is required. Expedited retrievals come with a higher cost per retrieval compared to standard retrievals but provide faster access to your archived data.  
upvoted 12 times
- nosense

Highly Voted

1 year, 5 months ago

glacier expedited retrieval times of typically 1-5 minutes.  
upvoted 5 times
- wsdasdasdqwdaw

1 year ago

Fully agree. Check here for evidences: <https://aws.amazon.com/s3/storage-classes/glacier/#:~:text=S3%20Glacier%20Flexible%20Retrieval%20provides,amounts%20of%20data%20typically%20in>  
upvoted 2 times
- Ravan

Most Recent

8 months, 1 week ago

Selected Answer: A

The most cost-effective solution that also meets the requirement of having the files available within a maximum of five minutes when needed is:  
  
A. Store the video archives in Amazon S3 Glacier and use Expedited retrievals.  
  
Amazon S3 Glacier is designed for long-term storage of data archives, providing a highly durable and secure solution at a low cost. With Expedited retrievals, data can be retrieved within a few minutes, which meets the requirement of having the files available within five minutes when needed. This option provides the balance between cost-effectiveness and retrieval speed, making it the best choice for the company's needs.  
upvoted 2 times

pentium75

10 months, 1 week ago

Selected Answer: A

Occasional cost for retrieval from Glacier is nothing compared to the huge storage cost savings compared to C. Still meets the five minute requirement.  
upvoted 2 times

master9

10 months, 1 week ago

Selected Answer: C

The retrieval price will play an important role here. I selected the "C" option because in "Glacier and use Expedited retrievals" its around \$0.004 per GB/month and for STD-IA \$0.0125 per GB/month  
<https://www.cloudforecast.io/blog/aws-s3-pricing-and-optimization-guide/>  
upvoted 1 times

pentium75

10 months, 1 week ago

But they "will rarely need to restore their files", thus the low cost for occasional expedited retrievals will be nothing compared to the huge storage cost savings.  
upvoted 2 times

ngo01214



1 year ago

s3 expedited can only be applied on glacier flexible retrieval storage class and s3 intelligent tiering archive access tier. so the answer should be C  
upvoted 3 times

pentium75

10 months, 1 week ago

A mentions "G3 Glacier" which has been renamed to "S3 Glacier Flexible Retrieval" and meets the requirements.  
upvoted 2 times

  **Smart** 1 year, 2 months ago

**Selected Answer: A**

I am going with option A, but it is a poorly written question. "For all but the largest archives (more than 250 MB), data accessed by using Expedited retrievals is typically made available within 1–5 minutes. "  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

Answer - A  
Fast availability: Although retrieval times for objects stored in Amazon S3 Glacier typically range from minutes to hours, you can use the Expedited retrievals option to expedite access to your archives. By using Expedited retrievals, the files can be made available in a maximum of five minutes when needed. However, Expedited retrievals do incur higher costs compared to standard retrievals.  
upvoted 2 times

  **hsinchang** 1 year, 3 months ago

**Selected Answer: A**

Expedited retrievals are designed for urgent requests and can provide access to data in as little as 1-5 minutes for most archive objects. Standard retrievals typically finish within 3-5 hours for objects stored in the S3 Glacier Flexible Retrieval storage class or S3 Intelligent-Tiering Archive Access tier. These retrievals typically finish within 12 hours for objects stored in the S3 Glacier Deep Archive storage class or S3 Intelligent-Tiering Deep Archive Access tier. So A.  
upvoted 3 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: A**

Expedited retrievals allow you to quickly access your data that's stored in the S3 Glacier Flexible Retrieval storage class or the S3 Intelligent-Tiering Archive Access tier when occasional urgent requests for restoring archives are required. Data accessed by using Expedited retrievals is typically made available within 1–5 minutes.  
upvoted 2 times



  **MrAWSAssociate** 1 year, 4 months ago

**Selected Answer: A**

A for sure!  
upvoted 2 times



  **Doyin8807** 1 year, 5 months ago

C because A is not the most cost effective  
upvoted 1 times

  **luiscc** 1 year, 5 months ago


**Selected Answer: A**

Expedited retrieval typically takes 1-5 minutes to retrieve data, making it suitable for the company's requirement of having the files available in a maximum of five minutes.  
upvoted 5 times

  **Efren** 1 year, 5 months ago

**Selected Answer: A**

Glacier expedite  
upvoted 3 times

  **EA100** 1 year, 5 months ago

Answer - A  
Fast availability: Although retrieval times for objects stored in Amazon S3 Glacier typically range from minutes to hours, you can use the Expedited retrievals option to expedite access to your archives. By using Expedited retrievals, the files can be made available in a maximum of five minutes when needed. However, Expedited retrievals do incur higher costs compared to standard retrievals.  
upvoted 2 times

A company is building a three-tier application on AWS. The presentation tier will serve a static website The logic tier is a containerized application. This application will store data in a relational database. The company wants to simplify deployment and to reduce operational costs.

Which solution will meet these requirements?

- A. Use Amazon S3 to host static content. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute power. Use a managed Amazon RDS cluster for the database.
- B. Use Amazon CloudFront to host static content. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 for compute power. Use a managed Amazon RDS cluster for the database.
- C. Use Amazon S3 to host static content. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute power. Use a managed Amazon RDS cluster for the database.
- D. Use Amazon EC2 Reserved Instances to host static content. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 for compute power. Use a managed Amazon RDS cluster for the database.

Correct Answer: A

Community vote distribution

A (100%)

  **Yadav\_Sanjay** Highly Voted 1 year, 5 months ago

**Selected Answer: A**

ECS is slightly cheaper than EKS  
upvoted 12 times

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: A**

Amazon S3 is a highly scalable and cost-effective storage service that can be used to host static website content. It provides durability, high availability, and low latency access to the static files.

Amazon ECS with AWS Fargate eliminates the need to manage the underlying infrastructure. It allows you to run containerized applications without provisioning or managing EC2 instances. This reduces operational overhead and provides scalability.

By using a managed Amazon RDS cluster for the database, you can offload the management tasks such as backups, patching, and monitoring to AWS. This reduces the operational burden and ensures high availability and durability of the database.  
upvoted 5 times

  **awsgeek75** Most Recent 10 months ago

**Selected Answer: A**

B: CloudFront = Extra cost for something they don't want (CDN)  
C: Kubernetes is more operationally complex than ECS containers on Fargate.  
D: EC2 expensive  
A: S3 is cheap for static content. ECS with Fargate is easiest implantation. Managed RDS is very low op overhead  
upvoted 5 times

  **wsdasdasdqwdaw** 1 year ago

Why not B ?  
upvoted 1 times

  **wsdasdasdqwdaw** 1 year ago

Aaa I got it. With CF we are adding additional cost => A.  
upvoted 2 times

  **cyber\_bedouin** 11 months ago

A is better because ECS Fargate = "containerized application"  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

Use Amazon S3 to host static content. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute power. Use a managed Amazon RDS cluster for the database.  
upvoted 3 times

  **jaydesai8** 1 year, 3 months ago

**Selected Answer: A**

S3= hosting static contents  
Ecs = Little cheaper than EKS  
RDS = Database  
upvoted 4 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: A**

Use Amazon S3 to host static content. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute power. Use a managed Amazon RDS cluster for the database  
upvoted 3 times

A company seeks a storage solution for its application. The solution must be highly available and scalable. The solution also must function as a file system be mountable by multiple Linux instances in AWS and on premises through native protocols, and have no minimum size requirements. The company has set up a Site-to-Site VPN for access from its on-premises network to its VPC.

Which storage solution meets these requirements?

- A. Amazon FSx Multi-AZ deployments
- B. Amazon Elastic Block Store (Amazon EBS) Multi-Attach volumes
- C. Amazon Elastic File System (Amazon EFS) with multiple mount targets
- D. Amazon Elastic File System (Amazon EFS) with a single mount target and multiple access points

Correct Answer: C

Community vote distribution

C (100%)

  **Felix\_br** Highly Voted 1 year, 5 months ago

**Selected Answer: C**

The other options are incorrect for the following reasons:

- A. Amazon FSx Multi-AZ deployments Amazon FSx is a managed file system service that provides access to file systems that are hosted on Amazon EC2 instances. Amazon FSx does not support native protocols, such as NFS.
  - B. Amazon Elastic Block Store (Amazon EBS) Multi-Attach volumes Amazon EBS is a block storage service that provides durable, block-level storage volumes for use with Amazon EC2 instances. Amazon EBS Multi-Attach volumes can be attached to multiple EC2 instances at the same time, but they cannot be mounted by multiple Linux instances through native protocols, such as NFS.
  - D. Amazon Elastic File System (Amazon EFS) with a single mount target and multiple access points A single mount target can only be used to mount the file system on a single EC2 instance. Multiple access points are used to provide access to the file system from different VPCs.
- upvoted 12 times

  **dkw2342** 7 months, 3 weeks ago

"A single mount target can only be used to mount the file system on a single EC2 instance. Multiple access points are used to provide access to the file system from different VPCs."

This is clearly wrong. You can have exactly one EFS mount target per subnet (AZ), and of course this mount target can be used by many clients (EC2 instances, containers etc.) - see diagram here for example: <https://docs.aws.amazon.com/efs/latest/ug/accessing-fs.html>

In my opinion, C and D are equally valid answers.

upvoted 2 times

  **Salilgen** 4 months, 1 week ago

D is not valid answer because "the solution must be highly available" then it must be multi-AZ: every AZ requires a mount target

upvoted 1 times

  **unbendable** 1 year ago

Amazon FSx ONTAP supports clients mounting it with NFS. <https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/attach-linux-client.html>. Though A is not clear about which FSx product is used

upvoted 2 times

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: C**

Amazon EFS is a fully managed file system service that provides scalable, shared storage for Amazon EC2 instances. It supports the Network File System version 4 (NFSv4) protocol, which is a native protocol for Linux-based systems. EFS is designed to be highly available, durable, and scalable.

upvoted 9 times

  **awsgeek75** Most Recent 10 months ago



**Selected Answer: C**

- A: FSx is a File Server, not a mountable file system
  - B: EBS can't be mounted on on-prem devices
  - D: Access points are not same as mount points
  - C: EFS support multi mount targets and on-prem devices: <https://docs.aws.amazon.com/efs/latest/ug/mounting-fs-mount-helper-direct.html>
- upvoted 3 times

  **iwannabeawsgod** 1 year ago


EFS POSIX LINUX

upvoted 3 times



  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C. Amazon Elastic File System (Amazon EFS) with multiple mount targets  
upvoted 3 times

  **boubie44** 1 year, 5 months ago

i don't understand why not D?  
upvoted 1 times

  **lucdt4** 1 year, 5 months ago

the requirement is mountable by multiple Linux  
-> C (multiple mount targets)  
upvoted 3 times

A 4-year-old media company is using the AWS Organizations all features feature set to organize its AWS accounts. According to the company's finance team, the billing information on the member accounts must not be accessible to anyone, including the root user of the member accounts.

Which solution will meet these requirements?

- A. Add all finance team users to an IAM group. Attach an AWS managed policy named Billing to the group.
- B. Attach an identity-based policy to deny access to the billing information to all users, including the root user.
- C. Create a service control policy (SCP) to deny access to the billing information. Attach the SCP to the root organizational unit (OU).
- D. Convert from the Organizations all features feature set to the Organizations consolidated billing feature set.

**Correct Answer: C**

Community vote distribution

C (100%)

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

**Selected Answer: C**

Service Control Policies (SCP): SCPs are an integral part of AWS Organizations and allow you to set fine-grained permissions on the organizational units (OUs) within your AWS Organization. SCPs provide central control over the maximum permissions that can be granted to member accounts, including the root user.

Denying Access to Billing Information: By creating an SCP and attaching it to the root OU, you can explicitly deny access to billing information for all accounts within the organization. SCPs can be used to restrict access to various AWS services and actions, including billing-related services.


Granular Control: SCPs enable you to define specific permissions and restrictions at the organizational unit level. By denying access to billing information at the root OU, you can ensure that no member accounts, including root users, have access to the billing information.

upvoted 8 times

  **Kiki\_Pass** Highly Voted  1 year, 3 months ago

but SCP do not apply to the management account (full admin power)?

upvoted 5 times

  **TwinSpark** 5 months, 3 weeks ago

i can understand this information coming from the famous course in udemy. I tought same, but after some research i now think it is a wrong information.

"SCPs affect all users and roles in attached accounts, including the root user. The only exceptions are those described in Tasks and entities not restricted by SCPs."

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps.html#:~:text=SCPs%20affect%20all%20users%20and,affect%20any%20service%2Dlinked%20role.](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html#:~:text=SCPs%20affect%20all%20users%20and,affect%20any%20service%2Dlinked%20role.)

upvoted 3 times

  **potomac** Most Recent  12 months ago

**Selected Answer: C**

SCP is for authorization

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C. Create a service control policy (SCP) to deny access to the billing information. Attach the SCP to the root organizational unit (OU)

upvoted 3 times

  **PRASAD180** 1 year, 4 months ago

C Crt 100%

upvoted 2 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: C**

Service control policy are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization. SCPs help you to ensure your accounts stay within your organization's access control guidelines. SCPs are available only in an organization that has all features enabled.

upvoted 4 times

  **Abrar2022** 1 year, 5 months ago

By denying access to billing information at the root OU, you can ensure that no member accounts, including root users, have access to the billing information.

upvoted 2 times



**nosense** 1 year, 5 months ago

**Selected Answer: C**

c for me

upvoted 2 times

An ecommerce company runs an application in the AWS Cloud that is integrated with an on-premises warehouse solution. The company uses Amazon Simple Notification Service (Amazon SNS) to send order messages to an on-premises HTTPS endpoint so the warehouse application can process the orders. The local data center team has detected that some of the order messages were not received.

A solutions architect needs to retain messages that are not delivered and analyze the messages for up to 14 days.

Which solution will meet these requirements with the LEAST development effort?

- A. Configure an Amazon SNS dead letter queue that has an Amazon Kinesis Data Stream target with a retention period of 14 days.
- B. Add an Amazon Simple Queue Service (Amazon SQS) queue with a retention period of 14 days between the application and Amazon SNS.
- C. Configure an Amazon SNS dead letter queue that has an Amazon Simple Queue Service (Amazon SQS) target with a retention period of 14 days.
- D. Configure an Amazon SNS dead letter queue that has an Amazon DynamoDB target with a TTL attribute set for a retention period of 14 days.

**Correct Answer: C**

Community vote distribution

C (80%)

B (20%)

  **pentium75** Highly Voted 10 months, 1 week ago

**Selected Answer: C**

"Configuring an Amazon SNS dead-letter queue for a subscription ...

A dead-letter queue is an Amazon SQS queue that an Amazon SNS subscription can target for messages that can't be delivered to subscribers successfully", this is exactly what C says. <https://docs.aws.amazon.com/sns/latest/dg/sns-configure-dead-letter-queue.html>

B, an SQS queue "between the application and Amazon SNS" would change the application logic. SQS cannot push messages to the "on-premises https endpoint", rather the destination would have to retrieve messages from the queue. Besides, option B would eventually deliver the messages that failed on the first attempt, which is NOT what is asked for. The goal is to retain undeliverable messages for analysis (NOT to deliver them), and this is typically achieved with a dead letter queue.

upvoted 11 times

  **awsgeek75** Highly Voted 10 months ago

**Selected Answer: C**

LEAST development effort!

A: Custom dead letter queue using Kinesis Data Stream (laughable solution!) so lots of coding

B: Change app logic to put SQS between SNS and the app. Also too much coding

D: Same as A, too much code change

C: SNS dead letter queue is by default a SQS que so no coding required

upvoted 5 times

  **osmk** Most Recent 9 months, 3 weeks ago

A dead-letter queue is an Amazon SQS queue that an Amazon SNS subscription can target for messages that can't be delivered to subscribers successfully.<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>

upvoted 3 times

  **Mikado211** 10 months, 3 weeks ago

**Selected Answer: C**

Problem here SNS dead letter queue is a SQS queue, so technically speaking both B and C are right. But I suppose that they want us to speak about SNS dead letter queue, that nobody do... meh, frustrating.

upvoted 4 times


  **Mikado211** 10 months, 3 weeks ago

Aaaah ok.

So with B == you place the SQS queue between the application and the SNS topic  
with C == you place the SQS queue as a DLQ for the SNS topic

Of course it's C !

upvoted 6 times

  **aws94** 10 months, 3 weeks ago

**Selected Answer: C**

<https://docs.aws.amazon.com/sns/latest/dg/sns-configure-dead-letter-queue.html>

upvoted 3 times

  **daniel1** 1 year ago

**Selected Answer: C**

GPT4 to the rescue:


The most appropriate solution would be to configure an Amazon SNS dead letter queue with an Amazon Simple Queue Service (Amazon SQS) target with a retention period of 14 days (Option C). This setup would ensure that any undelivered messages are retained in the SQS queue for up to 14 days for analysis, with minimal development effort required.

upvoted 2 times

  **ealpuche** 11 months, 1 week ago

ChatGPT is not a reliable source.

upvoted 9 times

  **Wayne23Fang** 1 year ago

**Selected Answer: B**

I like (B) since it is put SQS before SNS so we could prepare for retention. (C) dead letter Queue is kind of "rescue" effort. Also (C) should mention reprocessing dead letter.

upvoted 1 times

  **pentium75** 10 months, 1 week ago

"Reprocessing dead letters" is not desired here. They want to "retain messages that are not delivered and analyze the messages for up to 14 days", which is what C does.

upvoted 1 times

  **thanhnv142** 1 year ago

C is correct. It used a combination of SNS and SQS so it better than B.

upvoted 2 times

  **iwannabeawsgod** 1 year ago

**Selected Answer: C**

C is the answer

upvoted 2 times

  **Devsin2000** 1 year, 1 month ago

B is correct Answer. SQS Retain messages in queues for up to 14 days

C is incorrect because there is nothing called Amazon SNS dead letter queue

upvoted 2 times

  **pentium75** 10 months, 1 week ago

C "Configure an Amazon SNS dead letter queue"

AWS "Configuring an Amazon SNS dead-letter queue"

<https://docs.aws.amazon.com/sns/latest/dg/sns-configure-dead-letter-queue.html>

upvoted 2 times

  **RDM10** 1 year, 1 month ago

<https://docs.aws.amazon.com/sns/latest/dg/sns-configure-dead-letter-queue.html>

upvoted 6 times

  **lemur88** 1 year, 2 months ago

**Selected Answer: C**

<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C. Configure an Amazon SNS dead letter queue that has an Amazon Simple Queue Service (Amazon SQS) target with a retention period of 14 days. By using an Amazon SQS queue as the target for the dead letter queue, you ensure that the undelivered messages are reliably stored in a queue for up to 14 days. Amazon SQS allows you to specify a retention period for messages, which meets the retention requirement without additional development effort.

upvoted 2 times

  **mtmayer** 1 year, 2 months ago

**Selected Answer: B**

Dead Letter is a SQS feature not SNS.

A dead-letter queue is an Amazon SQS queue that an Amazon SNS subscription can target for messages that can't be delivered to subscribers successfully. Messages that can't be delivered due to client errors or server errors are held in the dead-letter queue for further analysis or reprocessing. For more information, see [Configuring an Amazon SNS dead-letter queue for a subscription](#) and [Amazon SNS message delivery retries](#).

<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>

upvoted 3 times

  **pentium75** 10 months, 1 week ago

"See Configuring an Amazon SNS (!) dead-letter queue", exactly, thus C.

upvoted 2 times

  **xyb** 1 year, 2 months ago

**Selected Answer: B**

In SNS, DLQs store the messages that failed to be delivered to subscribed endpoints. For more information, see Amazon SNS Dead-Letter Queues.

In SQS, DLQs store the messages that failed to be processed by your consumer application. This failure mode can happen when producers and consumers fail to interpret aspects of the protocol that they use to communicate. In that case, the consumer receives the message from the queue, but fails to process it, as the message doesn't have the structure or content that the consumer expects. The consumer can't delete the message from the queue either. After exhausting the receive count in the redrive policy, SQS can sideline the message to the DLQ. For more information, see Amazon SQS Dead-Letter Queues.

<https://aws.amazon.com/blogs/compute/designing-durable-serverless-apps-with-dlqs-for-amazon-sns-amazon-sqs-aws-lambda/>  
upvoted 2 times

  **pentium75** 10 months, 1 week ago

"Configuring an Amazon SNS dead-letter queue for a subscription

A dead-letter queue is an Amazon SQS queue that an Amazon SNS subscription can target for messages that can't be delivered to subscribers successfully. "


upvoted 2 times

  **TariqKipkemei** 1 year, 4 months ago

C is best to handle this requirement. Although good to note that dead-letter queue is an SQS queue.

"A dead-letter queue is an Amazon SQS queue that an Amazon SNS subscription can target for messages that can't be delivered to subscribers successfully. Messages that can't be delivered due to client errors or server errors are held in the dead-letter queue for further analysis or reprocessing."

<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html#:~:text=A%20dead%2Dletter%20queue%20is%20an%20Amazon%20SQS%20queue>  
upvoted 2 times

  **Felix\_br** 1 year, 5 months ago

C - Amazon SNS dead letter queues are used to handle messages that are not delivered to their intended recipients. When a message is sent to an Amazon SNS topic, it is first delivered to the topic's subscribers. If a message is not delivered to any of the subscribers, it is sent to the topic's dead letter queue.

Amazon SQS is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. Amazon SQS queues can be configured to have a retention period, which is the amount of time that messages will be kept in the queue before they are deleted.

To meet the requirements of the company, you can configure an Amazon SNS dead letter queue that has an Amazon SQS target with a retention period of 14 days. This will ensure that any messages that are not delivered to the on-premises warehouse application will be stored in the Amazon SQS queue for up to 14 days. The company can then analyze the messages in the Amazon SQS queue to determine why they were not delivered.

upvoted 3 times

  **Yadav\_Sanjay** 1 year, 5 months ago

**Selected Answer: C**

<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>  
upvoted 3 times

A gaming company uses Amazon DynamoDB to store user information such as geographic location, player data, and leaderboards. The company needs to configure continuous backups to an Amazon S3 bucket with a minimal amount of coding. The backups must not affect availability of the application and must not affect the read capacity units (RCUs) that are defined for the table.

Which solution meets these requirements?



- A. Use an Amazon EMR cluster. Create an Apache Hive job to back up the data to Amazon S3.
- B. Export the data directly from DynamoDB to Amazon S3 with continuous backups. Turn on point-in-time recovery for the table.
- C. Configure Amazon DynamoDB Streams. Create an AWS Lambda function to consume the stream and export the data to an Amazon S3 bucket.
- D. Create an AWS Lambda function to export the data from the database tables to Amazon S3 on a regular basis. Turn on point-in-time recovery for the table.


**Correct Answer: B**

*Community vote distribution*

B (89%)

11%

-   **elmogy**



Highly Voted 


 1 year, 5 months ago

**Selected Answer: B**

Continuous backups is a native feature of DynamoDB, it works at any scale without having to manage servers or clusters and allows you to export data across AWS Regions and accounts to any point-in-time in the last 35 days at a per-second granularity. Plus, it doesn't affect the read capacity or the availability of your production tables.

<https://aws.amazon.com/blogs/aws/new-export-amazon-dynamodb-table-data-to-data-lake-amazon-s3/>



upvoted 12 times
-   **awsgeek75**


Highly Voted 

 10 months ago

**Selected Answer: B**

A: Impacts RCU  
C: Requires coding of Lambda to read from stream to S3  
D: More coding in Lambda  
B: AWS Managed solution with no coding



upvoted 5 times
-   **potomac**

Most Recent 



 12 months ago

**Selected Answer: B**

DynamoDB export to S3 is a fully managed solution for exporting DynamoDB data to an Amazon S3 bucket at scale.



upvoted 4 times
-   **baba365** 1 year, 1 month ago

A DynamoDB stream is an ordered flow of information about changes to items in a DynamoDB table... for C.U.D events ( Create, Update, Delete) and its logs are retained for only 24hrs .

upvoted 4 times
-   **Guru4Cloud** 1 year, 2 months ago



**Selected Answer: B**

Export the data directly from DynamoDB to Amazon S3 with continuous backups. Turn on point-in-time recovery for the table.



upvoted 3 times
-   **ukivanlamlpi** 1 year, 2 months ago

**Selected Answer: C**

continous backup, no impact to availability ==> DynamoDB stream  
B. export is one off, noy continuous and demand on read capacity

upvoted 4 times
-   **hsinchang** 1 year, 3 months ago

minimal amount of coding rules out Lambda



upvoted 4 times
-   **Chris22usa** 1 year, 4 months ago

ChatGpt answer is C and it indicates continuous backup process uses DynamoDB stream actually

upvoted 2 times

  **pentium75** 10 months, 1 week ago

ChatGPT is usually wrong on these topics.  
upvoted 2 times

  **Gajendr** 10 months, 2 weeks ago

Wrong.  
"DynamoDB full exports are charged based on the size of the DynamoDB table (table data and local secondary indexes) at the point in time for which the export is done. DynamoDB incremental exports are charged based on the size of data processed from your continuous backups for the time period being exported."  
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/S3DataExport.HowItWorks.html>  
upvoted 1 times

  **TariqKipkemei** 1 year, 4 months ago

**Selected Answer: B**

Using DynamoDB table export, you can export data from an Amazon DynamoDB table from any time within your point-in-time recovery window to an Amazon S3 bucket. Exporting a table does not consume read capacity on the table, and has no impact on table performance and availability.  
upvoted 1 times

  **norris81** 1 year, 5 months ago

**Selected Answer: B**

<https://repost.aws/knowledge-center/back-up-dynamodb-s3>  
<https://aws.amazon.com/blogs/aws/new-amazon-dynamodb-continuous-backups-and-point-in-time-recovery-pitr/>  
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.Lambda.html>

There is no edit  
upvoted 2 times

  **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: B**

Continuous Backups: DynamoDB provides a feature called continuous backups, which automatically backs up your table data. Enabling continuous backups ensures that your table data is continuously backed up without the need for additional coding or manual interventions.

Export to Amazon S3: With continuous backups enabled, DynamoDB can directly export the backups to an Amazon S3 bucket. This eliminates the need for custom coding to export the data.



Minimal Coding: Option B requires the least amount of coding effort as continuous backups and the export to Amazon S3 functionality are built-in features of DynamoDB.

No Impact on Availability and RCUs: Enabling continuous backups and exporting data to Amazon S3 does not affect the availability of your application or the read capacity units (RCUs) defined for the table. These operations happen in the background and do not impact the table's performance or consume additional RCUs.  
upvoted 4 times

  **Efren** 1 year, 5 months ago

**Selected Answer: B**



DynamoDB Export to S3 feature  
Using this feature, you can export data from an Amazon DynamoDB table anytime within your point-in-time recovery window to an Amazon S3 bucket.  
upvoted 3 times

  **Efren** 1 year, 5 months ago

B also for me  
upvoted 3 times

  **norris81** 1 year, 5 months ago

<https://repost.aws/knowledge-center/back-up-dynamodb-s3>  
<https://aws.amazon.com/blogs/aws/new-amazon-dynamodb-continuous-backups-and-point-in-time-recovery-pitr/>  
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.Lambda.html>  
upvoted 1 times

  **Efren** 1 year, 5 months ago

you could mention what is the best answer from you :)  
upvoted 2 times

A solutions architect is designing an asynchronous application to process credit card data validation requests for a bank. The application must be secure and be able to process each request at least once.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Lambda event source mapping. Set Amazon Simple Queue Service (Amazon SQS) standard queues as the event source. Use AWS Key Management Service (SSE-KMS) for encryption. Add the kms:Decrypt permission for the Lambda execution role.
- B. Use AWS Lambda event source mapping. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues as the event source. Use SQS managed encryption keys (SSE-SQS) for encryption. Add the encryption key invocation permission for the Lambda function.
- C. Use the AWS Lambda event source mapping. Set Amazon Simple Queue Service (Amazon SQS) FIFO queues as the event source. Use AWS KMS keys (SSE-KMS). Add the kms:Decrypt permission for the Lambda execution role.
- D. Use the AWS Lambda event source mapping. Set Amazon Simple Queue Service (Amazon SQS) standard queues as the event source. Use AWS KMS keys (SSE-KMS) for encryption. Add the encryption key invocation permission for the Lambda function.

Correct Answer: A

Community vote distribution

A (77%)

B (21%)

  **elmogy** Highly Voted 1 year, 11 months ago

Selected Answer: A

SQS FIFO is slightly more expensive than standard queue  
<https://calculator.aws/#/addService/SQS>

I would still go with the standard because of the keyword "at least once" because FIFO process "exactly once". That leaves us with A and D, I believe that lambda function only needs to decrypt so I would choose A  
upvoted 12 times

  **pentium75** Highly Voted 1 year, 4 months ago

Selected Answer: A

"Process each request at least once" = Standard queue, rules out B and C which use more expensive FIFO queue  
  
Permissions are added to Lambda execution roles, not Lambda functions, thus D is out.  
upvoted 11 times

  **emakid** Most Recent 10 months, 1 week ago

Selected Answer: A

Use AWS Lambda event source mapping. Set Amazon Simple Queue Service (Amazon SQS) standard queues as the event source. Use AWS Key Management Service (SSE-KMS) for encryption. Add the kms permission for the Lambda execution role.  
upvoted 2 times



  **JackyCCK** 1 year, 2 months ago

D is not FIFO either  
upvoted 1 times

  **EdenWang** 1 year, 5 months ago

Selected Answer: B

With the SSE-SQS encryption type, you do not need to create, manage, or pay for SQS-managed encryption keys.  
upvoted 1 times

  **pentium75** 1 year, 4 months ago  
And what the hell is "encryption key invocation permission for the Lambda function"?  
upvoted 5 times

  **wsdasdasdwdaw** 1 year, 6 months ago

Initially though it is B, but it is said that the messages should be processed at lest once, not the same order, and Standard SQS is "almost" FIFO, which changed my opinion and I would go with A as correct.  
upvoted 4 times

  **BrijMohan08** 1 year, 8 months ago

Selected Answer: A

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/standard-queues.html>

upvoted 5 times

 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Using SQS FIFO queues ensures each message is processed at least once in order. SSE-SQS provides encryption that is handled entirely by SQS without needing decrypt permissions.

Standard SQS queues (Options A and D) do not guarantee order.

Using KMS keys (Options C and D) requires providing the Lambda role with decrypt permissions, adding complexity.

SQS FIFO queues with SSE-SQS encryption provide orderly, secure, server-side message processing that Lambda can consume without needing to manage decryption. This is the most efficient and cost-effective approach.

upvoted 8 times

 **Clouddon** 1 year, 7 months ago

Amazon SQS offers standard as the default queue type. Standard queues support a nearly unlimited number of API calls per second, per API action (SendMessage, ReceiveMessage, or DeleteMessage). Standard queues support at-least-once message delivery. However, occasionally (because of the highly distributed architecture that allows nearly unlimited throughput), more than one copy of a message might be delivered out of order. Standard queues provide best-effort ordering which ensures that messages are generally delivered in the same order as they're sent. Whereas, FIFO (First-In-First-Out) queues have all the capabilities of the standard queues, but are designed to enhance messaging between applications when the order of operations and events is critical, or where duplicates can't be tolerated. ( is correct)

upvoted 5 times

 **pentium75** 1 year, 4 months ago

But permissions are added to Lambda execution roles, not functions

upvoted 5 times

 **hsinchang** 1 year, 9 months ago

Least Privilege Policy leads to A over D.

upvoted 2 times

 **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: B**

Considering this is credit card validation process, there needs to be a strict 'process exactly once' policy offered by the SQS FIFO, and also SQS already supports server-side encryption with customer-provided encryption keys using the AWS Key Management Service (SSE-KMS) or using SQS-owned encryption keys (SSE-SQS). Both encryption options greatly reduce the operational burden and complexity involved in protecting data. Additionally, with the SSE-SQS encryption type, you do not need to create, manage, or pay for SQS-managed encryption keys. Therefore option B stands out for me.

upvoted 1 times

 **TariqKipkemei** 1 year, 5 months ago

I retract my answer and change it to A, there is a requirement to process each request 'at least once'. Only standard queues can deliver messages at least once.

There is also a requirement for the most 'cost-effective' option. Standard queues are the cheaper option.

<https://aws.amazon.com/sqs/pricing/#:~:text=SQS%20requests%20priced%3F>

upvoted 3 times

 **darren\_song** 1 year, 10 months ago

**Selected Answer: A**

[https://docs.aws.amazon.com/zh\\_tw/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-least-privilege-policy.html](https://docs.aws.amazon.com/zh_tw/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-least-privilege-policy.html)


upvoted 2 times

 **Abrar2022** 1 year, 11 months ago

**Selected Answer: A**

at least once and cost effective suggests SQS standard

upvoted 2 times

 **Felix\_br** 1 year, 11 months ago

**Selected Answer: B**

Solution B is the most cost-effective solution to meet the requirements of the application.

Amazon Simple Queue Service (SQS) FIFO queues are a good choice for this application because they guarantee that messages are processed in the order in which they are received. This is important for credit card data validation because it ensures that fraudulent transactions are not processed before legitimate transactions.


SQS managed encryption keys (SSE-SQS) are a good choice for encrypting the messages in the SQS queue because they are free to use. AWS Key Management Service (KMS) keys (SSE-KMS) are also a good choice for encrypting the messages, but they do incur a cost.

upvoted 2 times

 **pentium75** 1 year, 4 months ago

"They guarantee that messages are processed in the order in which they are received. This is important" but not asked for!

upvoted 3 times



 **omoakin** 1 year, 11 months ago

AAAAA  
upvoted 1 times



  **Yadav\_Sanjay** 1 year, 11 months ago

**Selected Answer: A**

should be A. Key word - at least once and cost effective suggests SQS standard  
upvoted 3 times

  **Efren** 1 year, 11 months ago

It has to be default, no FIFO. It doesnt say just one, it says at least once, so that is default queue that is cheaper than FIFO. Between the default options, nto sure to be honest  
upvoted 4 times



  **jayce5** 1 year, 11 months ago

No, when it comes to "credit card data validation," it should be FIFO. If you use the standard approach, there is a chance that people who come after will get processed before those who come first.  
upvoted 1 times

  **pentium75** 1 year, 4 months ago



Question clearly says "process each request at least once" which is the description of a standard queue. Your opinion how these transactions should be processed does not matter if it contradicts the requirements given.

Besides, it is about "credit card data validation", NOT payments. Nothing happens if they check twice is your credit card is valid.  
upvoted 2 times

  **awwass** 1 year, 11 months ago

**Selected Answer: A**

I guess A  
upvoted 2 times

  **awwass** 1 year, 11 months ago

This solution uses standard queues in Amazon SQS, which are less expensive than FIFO queues. It also uses AWS Key Management Service (SSE-KMS) for encryption, which is a cost-effective way to encrypt data at rest and in transit. The kms:Decrypt permission is added to the Lambda execution role to allow it to decrypt messages from the queue  
upvoted 2 times

A company has multiple AWS accounts for development work. Some staff consistently use oversized Amazon EC2 instances, which causes the company to exceed the yearly budget for the development accounts. The company wants to centrally restrict the creation of AWS resources in these accounts.

Which solution will meet these requirements with the LEAST development effort?

- A. Develop AWS Systems Manager templates that use an approved EC2 creation process. Use the approved Systems Manager templates to provision EC2 instances.
- B. Use AWS Organizations to organize the accounts into organizational units (OUs). Define and attach a service control policy (SCP) to control the usage of EC2 instance types.
- C. Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is created. Stop disallowed EC2 instance types.
- D. Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types. Ensure that staff can deploy EC2 instances only by using the Service Catalog products.

**Correct Answer: B**

*Community vote distribution*

B (93%)

7%

- alexandercamachop

Highly Voted

1 year, 5 months ago

Selected Answer: B

Anytime you see Multiple AWS Accounts, and needs to consolidate is AWS Organization. Also anytime we need to restrict anything in an organization, it is SCP policies.

upvoted 7 times
- omarshaban

Highly Voted

9 months, 3 weeks ago

IN MY EXAM

upvoted 5 times
- Cyberkayu

Most Recent

10 months, 2 weeks ago

Selected Answer: B

B. Multiple AWS account, consolidate under one AWS Organization, top down policy (SCP) to all member account to restrict EC2 Type.

upvoted 3 times
- Guru4Cloud

1 year, 2 months ago

Selected Answer: B

Use AWS Organizations to organize the accounts into organizational units (OUs). Define and attach a service control policy (SCP) to control the usage of EC2 instance types.

upvoted 3 times
- Ale1973

1 year, 2 months ago

Selected Answer: D

I have a question regarding this answer, what do they mean by "development effort"?:  
If they mean the work it takes to implement the solution (using develop as implement), option B achieves the constraint with little administrative overhead (there is less to do to configure this option).  
If by "development effort", they mean less effort for the development team, when development team try to deploy instances and gets errors because they are not allowed, this generates overhead. In this case the best option is D.  
What did you think?

upvoted 2 times
- pentium75

10 months, 1 week ago

"Development effort" = Develop the solution that the question asks for. We don't care about the developers whose permissions we want to restrict.

upvoted 4 times
- TariqKipkemei

1 year, 3 months ago



Selected Answer: B

Use AWS Organizations to organize the accounts into organizational units (OUs). Define and attach a service control policy (SCP) to control the usage of EC2 instance types

upvoted 3 times
- Blingy

1 year, 5 months ago

BBBBBBBBBB  
upvoted 2 times

  **elmogy** 1 year, 5 months ago

**Selected Answer: B**

I would choose B  
The other options would require some level of programming or custom resource creation:  
A. Developing Systems Manager templates requires development effort  
C. Configuring EventBridge rules and Lambda functions requires development effort  
D. Creating Service Catalog products requires development effort to define the allowed EC2 configurations.

Option B - Using Organizations service control policies - requires no custom development. It involves:  
Organizing accounts into OUs  
Creating an SCP that defines allowed/disallowed EC2 instance types  
Attaching the SCP to the appropriate OUs  
This is a native AWS service with a simple UI for defining and managing policies. No coding or resource creation is needed.  
So option B, using Organizations service control policies, will meet the requirements with the least development effort.  
upvoted 4 times



  **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: B**

AWS Organizations: AWS Organizations is a service that helps you centrally manage multiple AWS accounts. It enables you to group accounts into organizational units (OUs) and apply policies across those accounts.

Service Control Policies (SCPs): SCPs in AWS Organizations allow you to define fine-grained permissions and restrictions at the account or OU level. By attaching an SCP to the development accounts, you can control the creation and usage of EC2 instance types.

Least Development Effort: Option B requires minimal development effort as it leverages the built-in features of AWS Organizations and SCPs. You can define the SCP to restrict the use of oversized EC2 instance types and apply it to the appropriate OUs or accounts.  
upvoted 5 times

  **Efren** 1 year, 5 months ago

B for me as well  
upvoted 2 times

A company wants to use artificial intelligence (AI) to determine the quality of its customer service calls. The company currently manages calls in four different languages, including English. The company will offer new languages in the future. The company does not have the resources to regularly maintain machine learning (ML) models.

The company needs to create written sentiment analysis reports from the customer service call recordings. The customer service call recording text must be translated into English.

Which combination of steps will meet these requirements? (Choose three.)

- A. Use Amazon Comprehend to translate the audio recordings into English.
- B. Use Amazon Lex to create the written sentiment analysis reports.
- C. Use Amazon Polly to convert the audio recordings into text.
- D. Use Amazon Transcribe to convert the audio recordings in any language into text.
- E. Use Amazon Translate to translate text in any language to English.
- F. Use Amazon Comprehend to create the sentiment analysis reports.

**Correct Answer:** DEF

*Community vote distribution*

DEF (100%)

- cloudenthusiast**

Highly Voted

1 year, 5 months ago

Selected Answer: DEF

Amazon Transcribe will convert the audio recordings into text, Amazon Translate will translate the text into English, and Amazon Comprehend will perform sentiment analysis on the translated text to generate sentiment analysis reports.

upvoted 7 times
- awsgeek75**

Highly Voted

10 months ago

Selected Answer: DEF

A: Comprehend cannot translate  
B: Lex is like a chatbot so not useful  
C: Polly converts text to audio (polly the parrot!) so this is wrong  
D: Can convert audio to text  
E: Can translate  
F: Can do sentiment analysis reports

upvoted 6 times
- wsdasdasdqwda**

Most Recent

1 year ago

It is: DEF

upvoted 2 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: DEF

D. Use Amazon Transcribe to convert the audio recordings in any language into text.  
E. Use Amazon Translate to translate text in any language to English.  
F. Use Amazon Comprehend to create the sentiment analysis reports.

upvoted 3 times
- TariqKipkemei**

1 year, 3 months ago

Selected Answer: DEF

Amazon Transcribe to convert speech to text. Amazon Translate to translate text to english. Amazon Comprehend to perform sentiment analysis on translated text.

upvoted 2 times
- HareshPrajapati**

1 year, 5 months ago



afree with DEF

upvoted 2 times
- Blingy**

1 year, 5 months ago



I'd go with DEF too

upvoted 2 times

  **elmogy** 1 year, 5 months ago

**Selected Answer: DEF**

agree with DEF  
upvoted 3 times

  **Efren** 1 year, 5 months ago

agreed as well, weird  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

@efren - It is not weird - This need to know the services for it  
upvoted 3 times

A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message.

The administrator is using an IAM role that has the following IAM policy attached:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Condition": {
        "NotIpAddress": {
          "aws:SourceIp": [
            "192.0.2.0/24",
            "203.0.113.0/24"
          ]
        }
      },
      "Resource": ["*"]
    }
  ]
}
```

What is the cause of the unsuccessful request?

- A. The EC2 instance has a resource-based policy with a Deny statement.
- B. The principal has not been specified in the policy statement.
- C. The "Action" field does not grant the actions that are required to terminate the EC2 instance.
- D. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0.113.0/24.

**Correct Answer:** D

*Community vote distribution*

D (100%)

 **TariqKipkemei** Highly Voted 1 year, 9 months ago

**Selected Answer: D**

the command is coming from a source IP which is not in the allowed range.

upvoted 5 times

 **chasingsummer** Most Recent 1 year, 3 months ago

**Selected Answer: D**

I ran a Policy Simulator and indeed, D is right answer.

Here is the JSON policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:TerminateInstances",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
```

```
"Action": "ec2:TerminateInstances",
"Condition": {
  "NotIpAddress": {
    "aws:SourceIp" : [
      "192.0.2.0/24",
      "203.0.113.0/24"
    ]
  }
},
"Resource": "*"
}
```

upvoted 2 times

  **chasingsummer** 1 year, 3 months ago

The condition operator is "NotIpAddress" so I am not sure about D as right answer.

upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

Deny when IP address is not in (NotIPAddress). AWS has a weird way of stating Deny and it almost sound like double negative meaning positive. But read this doc for more clarity:  
[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_examples\\_aws\\_deny-ip.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-ip.html)

It has the exact same example! Good luck!

upvoted 2 times

  **LeonSauveterre** 5 months ago

This is exactly what I need, thanks!



upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: D**

If you want to read more about this, see how it works: [https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_examples\\_aws\\_deny-ip.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-ip.html)  
Same policy as in this question with almost same use case.  
D is correct answer.



upvoted 4 times

  **elmogy** 1 year, 11 months ago

**Selected Answer: D**

" aws:SourceIP " indicates the IP address that is trying to perform the action.

upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: D**

d for sure

upvoted 3 times

A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company’s AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (PII) or financial information, including passport numbers and credit card numbers.



Which solution will meet these requirements?


- A. Configure AWS Audit Manager on the account. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
- B. Configure Amazon S3 Inventory on the S3 bucket Configure Amazon Athena to query the inventory.
- C. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
- D. Use Amazon S3 Select to run a report across the S3 bucket.

**Correct Answer: C**

*Community vote distribution*



C (100%)


-   **Blingy**

Highly Voted 

 1 year, 5 months ago

Macie = Sensitive PII



upvoted 5 times
-   **elmogy**


Highly Voted 

 1 year, 5 months ago

**Selected Answer: C**

agree with C



upvoted 5 times
-   **awsgeek75**

Most Recent 

 10 months ago



**Selected Answer: C**

PII or sensitive data = Macie

upvoted 4 times
-   **Guru4Cloud** 1 year, 2 months ago



**Selected Answer: C**

Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.

upvoted 3 times
-   **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: C**

Amazon Macie is a data security service that uses machine learning (ML) and pattern matching to discover and help protect your sensitive data.

upvoted 3 times
-   **cloudenthusiast** 1 year, 5 months ago

**Selected Answer: C**

Amazon Macie is a service that helps discover, classify, and protect sensitive data stored in AWS. It uses machine learning algorithms and managed identifiers to detect various types of sensitive information, including personally identifiable information (PII) and financial information. By configuring Amazon Macie to run a data discovery job with the appropriate managed identifiers for the required data types (such as passport numbers and credit card numbers), the company can identify and classify any sensitive data present in the S3 bucket.

upvoted 5 times

A company uses on-premises servers to host its applications. The company is running out of storage capacity. The applications use both block storage and NFS storage. The company needs a high-performing solution that supports local caching without re-architecting its existing applications.

Which combination of actions should a solutions architect take to meet these requirements? (Choose two.)

- A. Mount Amazon S3 as a file system to the on-premises servers.
- B. Deploy an AWS Storage Gateway file gateway to replace NFS storage.
- C. Deploy AWS Snowball Edge to provision NFS mounts to on-premises servers.
- D. Deploy an AWS Storage Gateway volume gateway to replace the block storage.
- E. Deploy Amazon Elastic File System (Amazon EFS) volumes and mount them to on-premises servers.

**Correct Answer:** BD

Community vote distribution

BD (100%)

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

**Selected Answer:** BD

By combining the deployment of an AWS Storage Gateway file gateway and an AWS Storage Gateway volume gateway, the company can address both its block storage and NFS storage needs, while leveraging local caching capabilities for improved performance.

upvoted 7 times

  **awsgeek75** Most Recent  10 months ago

**Selected Answer:** BD

A: Not possible

C: Snowball edge is snowball with computing. It's not a NAS!

E: Technically yes but requires VPN or Direct Connect so re-architecture

B & D both use Storage Gateway which can be used as NFS and Block storage

<https://aws.amazon.com/storagegateway/>

upvoted 4 times

  **ftaws** 10 months, 2 weeks ago

Use the Storage Gateway -> It means that use S3 for storage ?

upvoted 3 times

  **thanhnv142** 1 year ago

DE

B is not correct be cause NFS is a file system while storage gw is a storage. To replace a file system, need another file system which is EFS.

upvoted 3 times

  **Tekk97** 11 months, 2 weeks ago

That's what i thought. but I think B is work too.

upvoted 2 times

  **wizcloudifa** 6 months, 1 week ago

if you focus on the wording of option B its "Storage gateway File gateway" not volume gateway hence it is perfect replacement for NFS files.

upvoted 1 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer:** BD

Deploy an AWS Storage Gateway file gateway to replace NFS storage

Deploy an AWS Storage Gateway volume gateway to replace the block storage

upvoted 3 times

  **elmogy** 1 year, 5 months ago

**Selected Answer:** BD

local caching is a key feature of AWS Storage Gateway solution

<https://aws.amazon.com/storagegateway/features/>

<https://aws.amazon.com/blogs/storage/aws-storage-gateway-increases-cache-4x-and-enhances-bandwidth-throttling/#:~:text=AWS%20Storage%20Gateway%20increases%20cache%204x%20and%20enhances,for%20Volume%20Gateway%20customers%200...%205%20Conclusion%20>

0...%205%20Conclusion%20

upvoted 4 times

  **Piccalo** 1 year, 5 months ago

**Selected Answer: BD**

B and D is the correct answer  
upvoted 2 times

Question #497

Topic 1

A company has a service that reads and writes large amounts of data from an Amazon S3 bucket in the same AWS Region. The service is deployed on Amazon EC2 instances within the private subnet of a VPC. The service communicates with Amazon S3 over a NAT gateway in the public subnet. However, the company wants a solution that will reduce the data output costs.

Which solution will meet these requirements MOST cost-effectively?

- A. Provision a dedicated EC2 NAT instance in the public subnet. Configure the route table for the private subnet to use the elastic network interface of this instance as the destination for all S3 traffic.
- B. Provision a dedicated EC2 NAT instance in the private subnet. Configure the route table for the public subnet to use the elastic network interface of this instance as the destination for all S3 traffic.
- C. Provision a VPC gateway endpoint. Configure the route table for the private subnet to use the gateway endpoint as the route for all S3 traffic.
- D. Provision a second NAT gateway. Configure the route table for the private subnet to use this NAT gateway as the destination for all S3 traffic.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **cloudenthusiast** **Highly Voted**  1 year, 5 months ago

**Selected Answer: C**

A VPC gateway endpoint allows you to privately access Amazon S3 from within your VPC without using a NAT gateway or NAT instance. By provisioning a VPC gateway endpoint for S3, the service in the private subnet can directly communicate with S3 without incurring data transfer costs for traffic going through a NAT gateway.  
upvoted 10 times

  **awsgeek75** **Most Recent**  10 months ago

**Selected Answer: C**

As a rule of thumb, EC2<->S3 in your workload should always try to use a VPC gateway unless there is an explicit restriction (account etc.) which disallows it.  
upvoted 4 times

  **Guru4Cloud** 1 year, 2 months ago



**Selected Answer: C**

Using a VPC endpoint for S3 allows the EC2 instances to access S3 directly over the Amazon network without traversing the internet. This significantly reduces data output charges.  
upvoted 3 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: C**

use VPC gateway endpoint to route traffic internally and save on costs.  
upvoted 2 times

  **elmogy** 1 year, 5 months ago

**Selected Answer: C**

private subnet needs to communicate with S3 --> VPC endpoint right away  
upvoted 3 times

A company uses Amazon S3 to store high-resolution pictures in an S3 bucket. To minimize application changes, the company stores the pictures as the latest version of an S3 object. The company needs to retain only the two most recent versions of the pictures.

The company wants to reduce costs. The company has identified the S3 bucket as a large expense.

Which solution will reduce the S3 costs with the LEAST operational overhead?

- A. Use S3 Lifecycle to delete expired object versions and retain the two most recent versions.
- B. Use an AWS Lambda function to check for older versions and delete all but the two most recent versions.
- C. Use S3 Batch Operations to delete noncurrent object versions and retain only the two most recent versions.
- D. Deactivate versioning on the S3 bucket and retain the two most recent versions.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: A**

S3 Lifecycle policies allow you to define rules that automatically transition or expire objects based on their age or other criteria. By configuring an S3 Lifecycle policy to delete expired object versions and retain only the two most recent versions, you can effectively manage the storage costs while maintaining the desired retention policy. This solution is highly automated and requires minimal operational overhead as the lifecycle management is handled by S3 itself.

upvoted 6 times

  **VellaDevil** Highly Voted 1 year, 3 months ago

**Selected Answer: A**

A --> "you can also provide a maximum number of noncurrent versions to retain."  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/intro-lifecycle-rules.html>

upvoted 5 times

  **awsgeek75** Most Recent 10 months ago

**Selected Answer: A**

B: Too much work with Lambda  
C: Possible but requires lot of work  
D: Oxymoron statement... i.e. how do you remove version and retain version at same time without additional overhead? Custom solution may be more work.

A: S3 Lifecycle is designed to retain object and version with set criteria

upvoted 4 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

Use S3 Lifecycle to delete expired object versions and retain the two most recent versions.

upvoted 3 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: A**

S3 Lifecycle to the rescue...whooooosh


upvoted 3 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: A**

A is correct.

upvoted 3 times

  **Konb** 1 year, 5 months ago

**Selected Answer: A**

Agree with LONGMEN

upvoted 4 times

A company needs to minimize the cost of its 1 Gbps AWS Direct Connect connection. The company's average connection utilization is less than 10%. A solutions architect must recommend a solution that will reduce the cost without compromising security.

Which solution will meet these requirements?

- A. Set up a new 1 Gbps Direct Connect connection. Share the connection with another AWS account.
- B. Set up a new 200 Mbps Direct Connect connection in the AWS Management Console.
- C. Contact an AWS Direct Connect Partner to order a 1 Gbps connection. Share the connection with another AWS account.
- D. Contact an AWS Direct Connect Partner to order a 200 Mbps hosted connection for an existing AWS account.

**Correct Answer: D**

*Community vote distribution*

D (82%)

B (18%)

- Abrar2022**

Highly Voted

1 year, 5 months ago

Selected Answer: D

Hosted Connection 50 Mbps, 100 Mbps, 200 Mbps,  
Dedicated Connection 1 Gbps, 10 Gbps, and 100 Gbps  
upvoted 11 times
- norris81**

Highly Voted

1 year, 5 months ago

Selected Answer: D

D  
  
For Dedicated Connections, 1 Gbps, 10 Gbps, and 100 Gbps ports are available. For Hosted Connections, connection speeds of 50 Mbps, 100 Mbps, 200 Mbps, 300 Mbps, 400 Mbps, 500 Mbps, 1 Gbps, 2 Gbps, 5 Gbps and 10 Gbps may be ordered from approved AWS Direct Connect Partners. See AWS Direct Connect Partners for more information.  
upvoted 5 times
- Ravan**

Most Recent

8 months, 1 week ago

Selected Answer: D

No, you cannot directly adjust the speed of an existing Direct Connect connection through the AWS Management Console.  
  
To adjust the speed of an existing Direct Connect connection, you typically need to contact your Direct Connect service provider. They can assist you in modifying the speed of your connection based on your requirements. Depending on the provider, this process may involve submitting a request or contacting their support team to initiate the necessary changes. Keep in mind that adjusting the speed of your Direct Connect connection may also involve contractual and billing considerations.  
upvoted 4 times
- awsgeek75**

10 months ago

Selected Answer: D

A: Not secure as sharing with another account  
B: I don't think this possible as you need ISP to setup Direct Connect  
C: Less secure due to sharing  
D: Direct connect partners can provide hosted solutions for existing accounts so correct answer  
upvoted 4 times
- awsgeek75**

10 months ago

For B I'm wrong above, it's because you cannot order 200MB connection through management console.  
upvoted 2 times
- pentium75**

10 months, 1 week ago

Selected Answer: D

< 1 Gbps = Hosted (through partner)  
upvoted 3 times
- Guru4Cloud**



1 year, 2 months ago



Selected Answer: B



If you already have an existing AWS Direct Connect connection configured at 1 Gbps, and you wish to reduce the connection bandwidth to 200 Mbps to minimize costs, you should indeed contact your AWS Direct Connect Partner and request to lower the connection speed to 200 Mbps.  
upvoted 4 times
- Guru4Cloud**



1 year, 2 months ago



I meant D.. DDDDDDDDDDD  
upvoted 6 times



  **omoakin** 1 year, 5 months ago  
BBBBBBBBBBBBBBBB  
upvoted 2 times



  **elmogy** 1 year, 5 months ago  
**Selected Answer: D**  
company need to setup a cheaper connection (200 M) but B is incorrect because you can only order port speeds of 1, 10, or 100 Gbps for more flexibility you can go with hosted connection, You can order port speeds between 50 Mbps and 10 Gbps.  
  
<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect.html>  
upvoted 5 times

  **cloudenthusiast** 1 year, 5 months ago  
**Selected Answer: B**  
By opting for a lower capacity 200 Mbps connection instead of the 1 Gbps connection, the company can significantly reduce costs. This solution ensures a dedicated and secure connection while aligning with the company's low utilization, resulting in cost savings.  
upvoted 4 times

  **pentium75** 10 months, 1 week ago  
But 200M cannot be ordered through Management Console, only partners.  
upvoted 3 times

  **nosense** 1 year, 5 months ago  
**Selected Answer: D**  
A hosted connection is a lower-cost option that is offered by AWS Direct Connect Partners  
upvoted 5 times

  **Efren** 1 year, 5 months ago  
Also, there are not 200 MBps direct connection speed.  
upvoted 1 times

  **nosense** 1 year, 5 months ago  
Hosted Connection 50 Mbps, 100 Mbps, 200 Mbps,  
Dedicated Connection 1 Gbps, 10 Gbps, and 100 Gbps  
B would require the company to purchase additional hardware or software  
upvoted 3 times

A company has multiple Windows file servers on premises. The company wants to migrate and consolidate its files into an Amazon FSx for Windows File Server file system. File permissions must be preserved to ensure that access rights do not change.

Which solutions will meet these requirements? (Choose two.)

- A. Deploy AWS DataSync agents on premises. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- B. Copy the shares on each file server into Amazon S3 buckets by using the AWS CLI. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- C. Remove the drives from each file server. Ship the drives to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- D. Order an AWS Snowcone device. Connect the device to the on-premises network. Launch AWS DataSync agents on the device. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- E. Order an AWS Snowball Edge Storage Optimized device. Connect the device to the on-premises network. Copy data to the device by using the AWS CLI. Ship the device back to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.

**Correct Answer:** AD

Community vote distribution

AD (96%)

4%

- pentium75**

Highly Voted

10 months, 1 week ago

Selected Answer: AD

B, C and E would copy the files to S3 first where permissions would be lost

upvoted 10 times
- cloudenthusiast**

Highly Voted

1 year, 5 months ago

Selected Answer: AD

A This option involves deploying DataSync agents on your on-premises file servers and using DataSync to transfer the data directly to the FSx for Windows File Server. DataSync ensures that file permissions are preserved during the migration process.

D This option involves using an AWS Snowcone device, a portable data transfer device. You would connect the Snowcone device to your on-premises network, launch DataSync agents on the device, and schedule DataSync tasks to transfer the data to FSx for Windows File Server. DataSync handles the migration process while preserving file permissions.

upvoted 8 times
- Guru4Cloud**

Most Recent

1 year, 2 months ago

Selected Answer: BD

Why not - BD?

upvoted 1 times
- pentium75**

10 months, 1 week ago

B copies the data to S3 first where file permissions would be lost.

upvoted 3 times
- Guru4Cloud**

1 year, 2 months ago

° This option uses S3 as an intermediary, ensuring that file permissions are preserved during the initial data copy. DataSync can then transfer the data from S3 to FSx while maintaining the permissions.

° This option uses a Snowcone device with DataSync agents to replicate the on-premises permission structure directly to FSx. This approach is suitable for maintaining file permissions during migration.

upvoted 2 times
- elmogy**

1 year, 5 months ago

Selected Answer: AD

the key is file permissions are preserved during the migration process. only datasync supports that

upvoted 5 times
- coolkidsclubvip**

1 year, 2 months ago

Bro,all 5 answers mentioned Datasync.....

upvoted 5 times
- Devsin2000**

1 year, 1 month ago

Yes but AD have only DataSync, whereas others have others have AWS CLI used.  
upvoted 3 times

  **nosense** 1 year, 5 months ago

**Selected Answer: AD**

Option B would require copy the data to Amazon S3 before transferring it to Amazon FSx for Windows File Server  
Option C would require the company to remove the drives from each file server and ship them to AWS  
upvoted 2 times

  **barracouto** 1 year, 2 months ago

Also, S3 doesn't retain permissions because it isn't a file system.  
upvoted 4 times

A company wants to ingest customer payment data into the company's data lake in Amazon S3. The company receives payment data every minute on average. The company wants to analyze the payment data in real time. Then the company wants to ingest the data into the data lake.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Use Amazon Kinesis Data Streams to ingest data. Use AWS Lambda to analyze the data in real time.
- B. Use AWS Glue to ingest data. Use Amazon Kinesis Data Analytics to analyze the data in real time.
- C. Use Amazon Kinesis Data Firehose to ingest data. Use Amazon Kinesis Data Analytics to analyze the data in real time.
- D. Use Amazon API Gateway to ingest data. Use AWS Lambda to analyze the data in real time.

**Correct Answer: C**

Community vote distribution

C (95%)

5%

- Axeashes

Highly Voted

1 year, 4 months ago

Kinesis Data Firehose is near real time (min. 60 sec). - The question is focusing on real time processing/analysis + efficiency -> Kinesis Data Stream is real time ingestion.

<https://www.amazonaws.cn/en/kinesis/data-firehose/#:~:text=Near%20real%2Dtime,is%20sent%20to%20the%20service.>

upvoted 12 times
- Axeashes

1 year, 4 months ago

Unless the intention is real time analytics not real time ingestion !

upvoted 5 times
- cloudenthusiast

Highly Voted

1 year, 5 months ago

Selected Answer: C

By leveraging the combination of Amazon Kinesis Data Firehose and Amazon Kinesis Data Analytics, you can efficiently ingest and analyze the payment data in real time without the need for manual processing or additional infrastructure management. This solution provides a streamlined and scalable approach to handle continuous data ingestion and analysis requirements.

upvoted 11 times

zdi561

Most Recent

3 months, 2 weeks ago

Selected Answer: C

The question says that the data is analyzed first then ingest into s3 for data lake

upvoted 1 times

wizcloudifa

6 months, 1 week ago

Selected Answer: C

Kinesis Firehouse = ingesting  
Kinesis Datastreams = storing  
Kinesis analytics = doing analysis

upvoted 6 times

awsgeek75

10 months ago

Selected Answer: C

Data is stored on S3 so real-time data analytics can be done with Kinesis Data Analytics which rules out Lambda solutions (A and D) as they are more operationally complex.

B is not useful it is more of ETL.

Firehose is actually to distribute data but given that company is already receiving data somehow so Firehose can basically distribute it to S3 with minimum latency. I have to admit this was confusing. I would have used Kinesis Streams to store on S3 and Data Analytics but combination is confusing!

upvoted 3 times

Mr\_Marcus

5 months ago

"Data is stored on S3..."

Nope. Re-read the first sentence. S3 is the destination, not the source.

The task is to ingest, analyze in real time, and store in S3.

upvoted 2 times

1rob

11 months, 1 week ago

Selected Answer: C

"payment data every minute on average" is a good-to-go- for firehose.

Also firehose is more operational efficient compared to Data Streams.

upvoted 3 times



**lucasbg** 11 months, 1 week ago

**Selected Answer: A**

I think this is A. The purpose of Firehose is to ingest and deliver to a data store, no to an analytics service. And in fact you can use lambda for real time analysis, such I find A more aligned.

upvoted 2 times



**pentium75** 10 months, 1 week ago

But developing and maintaining a custom Lambda function "to analyze the data in real time" is surely not as 'operationally efficient' as using an existing service such as Kinesis Data Analytics.

upvoted 3 times



**DDongi** 1 year ago

Firehose has a 60 sec delay so real time analytics should be without real time data isn't that problematic? Why would you have then real time analytics then in the first place?

upvoted 2 times



**Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

Kinesis Data Streams focuses on ingesting and storing data streams while Kinesis Data Firehose focuses on delivering data streams to select destinations, as the motive of the question is to do analytics, the answer should be C.

upvoted 4 times



**hsinchang** 1 year, 3 months ago

**Selected Answer: C**

Kinesis Data Streams focuses on ingesting and storing data streams while Kinesis Data Firehose focuses on delivering data streams to select destinations, as the motive of the question is to do analytics, the answer should be C.

upvoted 2 times



**james2033** 1 year, 3 months ago

**Selected Answer: C**

Quote "Connect with 30+ fully integrated AWS services and streaming destinations such as Amazon Simple Storage Service (S3)" at <https://aws.amazon.com/kinesis/data-firehose/> . Amazon Kinesis Data Analytics <https://aws.amazon.com/kinesis/data-analytics/>

upvoted 2 times



**TariqKipkemei** 1 year, 3 months ago

**Selected Answer: C**

Use Kinesis Firehose to capture and deliver the data to Kinesis Analytics to perform analytics.

upvoted 2 times



**Anmol\_1010** 1 year, 5 months ago

Did anyone took tge exam recently,  
How many questiona were there

upvoted 2 times



**omoakin** 1 year, 5 months ago

Can we understand why admin's answers are mostly wrong? Or is this done on purpose?

upvoted 3 times



**nosense** 1 year, 5 months ago

**Selected Answer: C**

Amazon Kinesis Data Firehose the most optimal variant

upvoted 4 times



**kailu** 1 year, 5 months ago

Shouldn't C be more appropriate?

upvoted 5 times



**MostofMichelle** 1 year, 5 months ago

You're right. I believe the answers are wrong on purpose, so good thing votes can be made on answers and discussions are allowed.

upvoted 2 times

A company runs a website that uses a content management system (CMS) on Amazon EC2. The CMS runs on a single EC2 instance and uses an Amazon Aurora MySQL Multi-AZ DB instance for the data tier. Website images are stored on an Amazon Elastic Block Store (Amazon EBS) volume that is mounted inside the EC2 instance.

Which combination of actions should a solutions architect take to improve the performance and resilience of the website? (Choose two.)

- A. Move the website images into an Amazon S3 bucket that is mounted on every EC2 instance
- B. Share the website images by using an NFS share from the primary EC2 instance. Mount this share on the other EC2 instances.
- C. Move the website images onto an Amazon Elastic File System (Amazon EFS) file system that is mounted on every EC2 instance.
- D. Create an Amazon Machine Image (AMI) from the existing EC2 instance. Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling group. Configure the Auto Scaling group to maintain a minimum of two instances. Configure an accelerator in AWS Global Accelerator for the website
- E. Create an Amazon Machine Image (AMI) from the existing EC2 instance. Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling group. Configure the Auto Scaling group to maintain a minimum of two instances. Configure an Amazon CloudFront distribution for the website.

**Correct Answer:** CE

*Community vote distribution*

CE (75%)

AE (25%)

- cloudenthusiast**

Highly Voted

1 year, 11 months ago

Selected Answer: CE

By combining the use of Amazon EFS for shared file storage and Amazon CloudFront for content delivery, you can achieve improved performance and resilience for the website.

upvoted 14 times
- wizcloudifa**

Highly Voted

1 year ago

Selected Answer: CE

First of all you should understand, a website using CMS is a dynamic one not static, so A is out, B is more complicated than C, so C, and between global accelerator and cloudfront, Cloudfront suits better as there is no legacy protocols data(UDP, etc) that needs to be accessed, hence E

upvoted 5 times
- foha2012**

Most Recent

1 year, 4 months ago

I choose AE. Although I don't know if s3 can be mounted on ec2 ?? Maybe wrong wording. Efs is a better choice but its not a natural selection for strong images.

upvoted 2 times
- awsgeek75**

1 year, 3 months ago

I made the same mistake but mounting S3 on EC2 is a painful operation so EFS makes more sense (C). Option E takes care of caching static images on CDN so that problem is solved along with resilience etc.

upvoted 4 times
- pentium75**

1 year, 4 months ago

Selected Answer: CE

Not A because you can't mount an S3 bucket on an EC2 instance. You could use a file gateway and share an S3 bucket via NFS and mount that on EC2, but that is not mentioned here and would also not make sense.

upvoted 4 times
- seldiora**

7 months ago

it is possible to mount s3 to EC2, just quite difficult: <https://aws.amazon.com/blogs/storage/mounting-amazon-s3-to-an-amazon-ec2-instance-using-a-private-connection-to-s3-file-gateway/>

upvoted 2 times
- potomac**

1 year, 6 months ago

Selected Answer: CE

You can mount EFS file systems to multiple Amazon EC2 instances remotely and securely without having to log in to the instances by using the AWS Systems Manager Run Command.

upvoted 3 times
- wsdasdasdqwda**

1 year, 6 months ago

A is out of the game for sure. Mount S3 to EC2 ... madness. The question is CE or DE, but it is CE because of AWS Global Accelerator is match with NLB, not ALB as it is staded in option D, thus CE as many of all here noted.

upvoted 5 times

  **thanhnv142** 1 year, 6 months ago

A and E is correct. We have a cloud fromt + S3 combo

upvoted 1 times

  **wsdasdasdqwdaw** 1 year, 6 months ago

S3 can't be mounted on EC2 it is not A for sure!

upvoted 4 times

  **NickGordon** 1 year, 5 months ago



<https://aws.amazon.com/blogs/storage/mounting-amazon-s3-to-an-amazon-ec2-instance-using-a-private-connection-to-s3-file-gateway/>

upvoted 3 times

  **pentium75** 1 year, 4 months ago

"Using a ... S3 file gateway"

upvoted 2 times

  **thanhnv142** 1 year, 6 months ago

A and E.

C is not correct because You dont mount a new EFS onto existing EC2. If you do that, you have to migrate all exising data in EBS into EFS. Then remove all the EBS. Should never do this.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

I can't follow. EFS provides NFS mount points, how can you not mount those onto existing EC2?



upvoted 2 times

  **franbarberan** 1 year, 7 months ago

**Selected Answer: CE**

<https://bluexp.netapp.com/blog/ebs-efs-amazons3-best-cloud-storage-system>

upvoted 2 times

  **Smart** 1 year, 8 months ago

**Selected Answer: CE**

Not A - S3 cannot be mounted (up until few months ago). Exam does not test for the updates in last 6 months.

upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: AE**

You have summarized the reasons why options A and E are the best choices very well.

Migrating static website assets like images to Amazon S3 enables high scalability, durability and shared access across instances. This improves performance.

Using Auto Scaling with load balancing provides elasticity and resilience. Adding a CloudFront distribution further boosts performance through caching and content delivery.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

You can't directly mount an S3 bucket on EC2.

upvoted 2 times

  **Ale1973** 1 year, 8 months ago

**Selected Answer: AE**

Both options AE and CE would work, but I choose AE, because, on my opinion, S3 is best suited for performance and resilience.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

You can't directly mount an S3 bucket on EC2

upvoted 2 times

  **MicketyMouse** 1 year, 9 months ago

**Selected Answer: CE**

EFS, unlike EBS, can be mounted across multiple EC2 instances and hence C over A.

upvoted 2 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: AE**



Technically both options AE and CE would work. But S3 is best suited for unstructured data, and the key benefit of mounting S3 on EC2 is that it provides a cost-effective alternative of using object storage for applications dealing with large files, as compared to expensive file or block storage. At the same time it provides more performant, scalable and highly available storage for these applications.

Even though there is no mention of 'cost efficient' in this question, in the real world cost is the no.1 factor.  
In the exam I believe both options would be a pass.

<https://aws.amazon.com/blogs/storage/mounting-amazon-s3-to-an-amazon-ec2-instance-using-a-private-connection-to-s3-file-gateway/>  
upvoted 4 times



  **pentium75** 1 year, 4 months ago

You can't directly mount an S3 bucket on EC2, only through file gateway  
upvoted 1 times

  **AshutoshSingh1923** 1 year, 10 months ago

**Selected Answer: CE**



Option C provides moving the website images onto an Amazon EFS file system that is mounted on every EC2 instance. Amazon EFS provides a scalable and fully managed file storage solution that can be accessed concurrently from multiple EC2 instances. This ensures that the website images can be accessed efficiently and consistently by all instances, improving performance  
In Option E The Auto Scaling group maintains a minimum of two instances, ensuring resilience by automatically replacing any unhealthy instances. Additionally, configuring an Amazon CloudFront distribution for the website further improves performance by caching content at edge locations closer to the end-users, reducing latency and improving content delivery.  
Hence combining these actions, the website's performance is improved through efficient image storage and content delivery  
upvoted 3 times

  **Vadbro7** 1 year, 10 months ago

Which answe is correct?the most voted ones or the Suggested answers?  
upvoted 1 times

  **9be0170** 10 months ago

the most voted always  
upvoted 2 times

  **mattcl** 1 year, 10 months ago

A and E: S3 is perfect for images. Besides is the perfect partner of cloudfront  
upvoted 2 times

A company runs an infrastructure monitoring service. The company is building a new feature that will enable the service to monitor data in customer AWS accounts. The new feature will call AWS APIs in customer accounts to describe Amazon EC2 instances and read Amazon CloudWatch metrics.

What should the company do to obtain access to customer accounts in the MOST secure way?

- A. Ensure that the customers create an IAM role in their account with read-only EC2 and CloudWatch permissions and a trust policy to the company's account.
- B. Create a serverless API that implements a token vending machine to provide temporary AWS credentials for a role with read-only EC2 and CloudWatch permissions.
- C. Ensure that the customers create an IAM user in their account with read-only EC2 and CloudWatch permissions. Encrypt and store customer access and secret keys in a secrets management system.
- D. Ensure that the customers create an Amazon Cognito user in their account to use an IAM role with read-only EC2 and CloudWatch permissions. Encrypt and store the Amazon Cognito user and password in a secrets management system.

Correct Answer: A

Community vote distribution

A (94%)

6%

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

Selected Answer: A

By having customers create an IAM role with the necessary permissions in their own accounts, the company can use AWS Identity and Access Management (IAM) to establish cross-account access. The trust policy allows the company's AWS account to assume the customer's IAM role temporarily, granting access to the specified resources (EC2 instances and CloudWatch metrics) within the customer's account. This approach follows the principle of least privilege, as the company only requests the necessary permissions and does not require long-term access keys or user credentials from the customers.

upvoted 15 times

  **Piccalo** Highly Voted 1 year, 5 months ago

Selected Answer: A

A. Roles give temporary credentials

upvoted 9 times

  **Efren** 1 year, 5 months ago

Agreed . Role is the keyword

upvoted 3 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

Selected Answer: A

B: Sharing credentials, even temporary, is insecure

C: Access and secret keys. That won't work and sharing secrets outside of account is not secure for this use case

A: Keyword "trust policy"

D: Again, sharing username and pwd and sharing in any way is not secure

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Selected Answer: A

Not B (would be about access to the company's account, not the customers' accounts)

Not C (storing credentials in a custom system is a big nono)

Not D (Cognito has nothing to do here and "user and password" is terrible)

upvoted 3 times

  **1rob** 11 months, 1 week ago

Selected Answer: D

The company's infrastructure monitoring service needs to call aws API's in the MOST secure way. So you have to focus on restricting access to the APIs and there is where cognito comes in to play.

upvoted 2 times

  **awsgeek75** 10 months ago

Are you suggesting to restrict CloudWatch API with Cognito roles?

upvoted 2 times

  **pentium75** 10 months, 1 week ago

What is unsecure with A?

upvoted 2 times

  **1rob** 9 months, 4 weeks ago

The company runs an infrastructure monitoring service. Nowhere is stated that this service lives in an aws account. So A and C I wouldn't choose.

B is a bit too vague. So I end up with D.

upvoted 1 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

A is the most secure approach for accessing customer accounts.

Having customers create a cross-account IAM role with the appropriate permissions, and configuring the trust policy to allow the monitoring service principal account access, implements secure delegation and least privilege access.

upvoted 2 times

A company needs to connect several VPCs in the us-east-1 Region that span hundreds of AWS accounts. The company's networking team has its own AWS account to manage the cloud network.

What is the MOST operationally efficient solution to connect the VPCs?

- A. Set up VPC peering connections between each VPC. Update each associated subnet's route table
- B. Configure a NAT gateway and an internet gateway in each VPC to connect each VPC through the internet
- C. Create an AWS Transit Gateway in the networking team's AWS account. Configure static routes from each VPC.
- D. Deploy VPN gateways in each VPC. Create a transit VPC in the networking team's AWS account to connect to each VPC.

**Correct Answer: C**

Community vote distribution

C (100%)

  **hsinchang** Highly Voted  1 year, 3 months ago

**Selected Answer: C**

The main difference between AWS Transit Gateway and VPC peering is that AWS Transit Gateway is designed to connect multiple VPCs together in a hub-and-spoke model, while VPC peering is designed to connect two VPCs together in a peer-to-peer model.

As we have several VPCs here, the answer should be C.

upvoted 19 times

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

**Selected Answer: C**

AWS Transit Gateway is a highly scalable and centralized hub for connecting multiple VPCs, on-premises networks, and remote networks. It simplifies network connectivity by providing a single entry point and reducing the number of connections required. In this scenario, deploying an AWS Transit Gateway in the networking team's AWS account allows for efficient management and control over the network connectivity across multiple VPCs.

upvoted 7 times

  **awsgeek75** Most Recent  9 months, 2 weeks ago

**Selected Answer: C**

A: This option is suggesting hundreds of peering connection for EACH VPC. Nope!

B: NAT gateway is for network translation not VPC interconnectivity so this is wrong

C: Transit GW + static routes will connect all VPCs <https://aws.amazon.com/transit-gateway/>

D: VPN gateway is for on-prem to VPN for a VPC. There is no on-prem here so this is wrong

upvoted 3 times

  **TariqKipkemei** 12 months ago

**Selected Answer: C**

Connect, Monitor and Manage Multiple VPCs in one place = AWS Transit Gateway

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: C**

C is the most operationally efficient solution for connecting a large number of VPCs across accounts.

Using AWS Transit Gateway allows all the VPCs to connect to a central hub without needing to create a mesh of VPC peering connections between each VPC pair.

This significantly reduces the operational overhead of managing the network topology as new VPCs are added or changed.

The networking team can centrally manage the Transit Gateway routing and share it across accounts using Resource Access Manager.

upvoted 3 times

  **MirKhobaeb** 1 year, 5 months ago

Answer is C

upvoted 2 times

  **MirKhobaeb** 1 year, 5 months ago

A transit gateway is a network transit hub that you can use to interconnect your virtual private clouds (VPCs) and on-premises networks. As your cloud infrastructure expands globally, inter-Region peering connects transit gateways together using the AWS Global Infrastructure. Your data is automatically encrypted and never travels over the public internet.

upvoted 3 times

  **nosense** 1 year, 5 months ago

**Selected Answer: C**

I voted for c  
upvoted 3 times

  **nosense** 1 year, 5 months ago

An AWS Transit Gateway is a highly scalable and secure way to connect VPCs in multiple AWS accounts. It is a central hub that routes traffic between VPCs, on-premises networks, and AWS services.  
upvoted 4 times

A company has Amazon EC2 instances that run nightly batch jobs to process data. The EC2 instances run in an Auto Scaling group that uses On-Demand billing. If a job fails on one instance, another instance will reprocess the job. The batch jobs run between 12:00 AM and 06:00 AM local time every day.

Which solution will provide EC2 instances to meet these requirements MOST cost-effectively?

- A. Purchase a 1-year Savings Plan for Amazon EC2 that covers the instance family of the Auto Scaling group that the batch job uses.
- B. Purchase a 1-year Reserved Instance for the specific instance type and operating system of the instances in the Auto Scaling group that the batch job uses.
- C. Create a new launch template for the Auto Scaling group. Set the instances to Spot Instances. Set a policy to scale out based on CPU usage.
- D. Create a new launch template for the Auto Scaling group. Increase the instance size. Set a policy to scale out based on CPU usage.

**Correct Answer:** C

*Community vote distribution*

C (100%)

  **cloudenthusiast** Highly Voted 1 year, 5 months ago

**Selected Answer: C**

Purchasing a 1-year Savings Plan (option A) or a 1-year Reserved Instance (option B) may provide cost savings, but they are more suitable for long-running, steady-state workloads. Since your batch jobs run for a specific period each day, using Spot Instances with the ability to scale out based on CPU usage is a more cost-effective choice.

upvoted 14 times

  **Guru4Cloud** Highly Voted 1 year, 2 months ago

**Selected Answer: C**

C is the most cost-effective solution in this scenario.

Using Spot Instances allows EC2 capacity to be purchased at significant discounts compared to On-Demand prices. The auto scaling group can scale out to add Spot Instances when needed for the batch jobs.

If Spot Instances become unavailable, regular On-Demand Instances will be launched instead to maintain capacity. The potential for interruptions is acceptable since failed jobs can be re-run.

upvoted 7 times

  **sandordini** Most Recent 6 months, 2 weeks ago

**Selected Answer: C**

Stateless, most cost-effective >> Spot

upvoted 3 times

  **awsgeek75** 10 months ago

**Selected Answer: C**

You don't need any scaling really as the job runs on another EC2 instance if it fails on first one. A. B. D are all more expensive than C due to spot instance being cheaper than reserved instances.



upvoted 3 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: C**

Spot Instances to the rescue....whooosh

upvoted 2 times

  **wRhIH** 1 year, 4 months ago

" If a job fails on one instance, another instance will reprocess the job". This ensures Spot Instances are enough for this case



upvoted 4 times

  **Abrar2022** 1 year, 5 months ago



**Selected Answer: C**



Since your batch jobs run for a specific period each day, using Spot Instances with the ability to scale out based on CPU usage is a more cost-effective choice.

upvoted 2 times

  **Blingy** 1 year, 5 months ago

C FOR ME COS OF SPOT INSTACES  
upvoted 3 times

  **udo2020** 1 year, 5 months ago  
First I think it is B but because of cost saving I think it should be C spot instances.  
upvoted 2 times

  **nosense** 1 year, 5 months ago  

Selected Answer: C

  
c for me  
upvoted 2 times

A social media company is building a feature for its website. The feature will give users the ability to upload photos. The company expects significant increases in demand during large events and must ensure that the website can handle the upload traffic from users.

Which solution meets these requirements with the MOST scalability?

- A. Upload files from the user's browser to the application servers. Transfer the files to an Amazon S3 bucket.
- B. Provision an AWS Storage Gateway file gateway. Upload files directly from the user's browser to the file gateway.
- C. Generate Amazon S3 presigned URLs in the application. Upload files directly from the user's browser into an S3 bucket.
- D. Provision an Amazon Elastic File System (Amazon EFS) file system. Upload files directly from the user's browser to the file system.

Correct Answer: C

Community vote distribution

C (91%)

9%

  **cloudenthusiast** Highly Voted  1 year, 5 months ago

**Selected Answer: C**

This approach allows users to upload files directly to S3 without passing through the application servers, reducing the load on the application and improving scalability. It leverages the client-side capabilities to handle the file uploads and offloads the processing to S3.

upvoted 16 times

  **Guru4Cloud** Highly Voted  1 year, 2 months ago

**Selected Answer: C**

C is the best solution to meet the scalability requirements.

Generating S3 presigned URLs allows users to upload directly to S3 instead of application servers. This removes the application servers as a bottleneck for upload traffic.


S3 can scale to handle very high volumes of uploads with no limits on storage or throughput. Using presigned URLs leverages this scalability.

upvoted 5 times

  **hro** Most Recent  7 months, 1 week ago

C - You may use presigned URLs to allow someone to upload an object to your Amazon S3 bucket. Using a presigned URL will allow an upload without requiring another party to have AWS security credentials or permissions.

upvoted 3 times

  **awsgeek75** 9 months, 4 weeks ago

**Selected Answer: C**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-presigned-url.html>

"You can also use presigned URLs to allow someone to upload a specific object to your Amazon S3 bucket. This allows an upload without requiring another party to have AWS security credentials or permissions. "

upvoted 2 times

  **Goutham4981** 11 months, 2 weeks ago

**Selected Answer: A**

S3 presigned url is used for sharing objects from an s3 bucket and not for uploading to an s3 bucket

upvoted 3 times

  **Murtadhaceit** 10 months, 3 weeks ago

No. It allows to download and upload.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-presigned-url.html>

upvoted 4 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: C**

You may use presigned URLs to allow someone to upload an object to your Amazon S3 bucket. Using a presigned URL will allow an upload without requiring another party to have AWS security credentials or permissions.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/PresignedUrlUploadObject.html>

upvoted 2 times

  **baba365** 1 year, 3 months ago

Hello Moderator. This question and answer should be rephrased because:

1. S3 pre-signed URLs are used to share objects FROM S3 buckets

2. How scalable are pre-signed URLs when they are time constrained?

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.html>

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Both is wrong

Presigned URLs can be used for upload

The solution is scalable because you can issue thousands of pre-signed URLs, and thousands of users can upload images at the same time.

User wants to upload picture -> server generates presigned URL and sends it to the app -> app uploads file

upvoted 3 times

  **nosense** 1 year, 5 months ago

**Selected Answer: C**

the most scalable because it allows users to upload files directly to Amazon S3,

upvoted 4 times

A company has a web application for travel ticketing. The application is based on a database that runs in a single data center in North America. The company wants to expand the application to serve a global user base. The company needs to deploy the application to multiple AWS Regions. Average latency must be less than 1 second on updates to the reservation database.

The company wants to have separate deployments of its web platform across multiple Regions. However, the company must maintain a single primary reservation database that is globally consistent.

Which solution should a solutions architect recommend to meet these requirements?

- A. Convert the application to use Amazon DynamoDB. Use a global table for the center reservation table. Use the correct Regional endpoint in each Regional deployment.
- B. Migrate the database to an Amazon Aurora MySQL database. Deploy Aurora Read Replicas in each Region. Use the correct Regional endpoint in each Regional deployment for access to the database.
- C. Migrate the database to an Amazon RDS for MySQL database. Deploy MySQL read replicas in each Region. Use the correct Regional endpoint in each Regional deployment for access to the database.
- D. Migrate the application to an Amazon Aurora Serverless database. Deploy instances of the database to each Region. Use the correct Regional endpoint in each Regional deployment to access the database. Use AWS Lambda functions to process event streams in each Region to synchronize the databases.

Correct Answer: A

Community vote distribution

A (51%)

B (49%)

  **cloudenthusiast** Highly Voted  1 year, 11 months ago

**Selected Answer: A**

Using DynamoDB's global tables feature, you can achieve a globally consistent reservation database with low latency on updates, making it suitable for serving a global user base. The automatic replication provided by DynamoDB eliminates the need for manual synchronization between Regions.  
upvoted 21 times

  **FlyingHawk** 4 months, 3 weeks ago

DynamoDB Global Tables rely on eventual consistency for replication between Regions. This means that updates in one Region may take milliseconds to seconds to propagate to other Regions.  
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadConsistency.html>

The requirement specifies global consistency, which DynamoDB cannot guarantee.  
upvoted 1 times

  **upliftinghut** Highly Voted  1 year, 3 months ago

**Selected Answer: B**

Aurora: less than 1 second: <https://aws.amazon.com/rds/aurora/global-database/>  
DynamoDB: from 0.5 to 2.5 second: [https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables\\_HowItWorks.html](https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_HowItWorks.html)  
upvoted 6 times

  **TheLaPlanta** 1 year, 1 month ago

B doesn't say Aurora Global  
upvoted 11 times

  **MatAlves** 7 months, 3 weeks ago

DynamoDB doesn't meet the <1s req though.  
upvoted 2 times

  **tch** Most Recent  1 month, 3 weeks ago

**Selected Answer: A**



answer should be A.  
we can consider if B change to "Amazon Aurora Global Database" ...  
upvoted 1 times

  **tch** 1 month, 3 weeks ago

Global tables - multi-Region replication for DynamoDB  
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.htm>Amazon DynamoDB global tables are a fully

managed, multi-Region, and multi-active database option that delivers fast and localized read and write performance for massively scaled global applications.

upvoted 1 times

  **Dantecito** 2 months, 3 weeks ago

**Selected Answer: A**

A is the only option that allows you to write from any region. I understand the < 1 second latency, but DynamoDB global tables use a last-writer-wins reconciliation between concurrent updates.

B and C only supports read scaling, not low-latency global writes. It will be much much much much than 1 second to write on the database.

D I stopped reading at lambda so NO!.

upvoted 1 times

  **zdi561** 2 months, 4 weeks ago

**Selected Answer: A**

Dynamodb is consistent and read/write in all node, Aurora db , only the master is writable. For reservation , writing is critical.

upvoted 1 times

  **khatingarun** 3 months, 1 week ago

**Selected Answer: A**

A is the correct answer

upvoted 1 times



  **zdi561** 3 months, 1 week ago

**Selected Answer: B**

dynamodb is not globally consistent. DynamoDB doesn't support strongly consistent reads across Regions.

[https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables\\_HowItWorks.html](https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_HowItWorks.html)

upvoted 2 times

  **hilker1983** 4 months ago

**Selected Answer: B**

correct answer bbbb

upvoted 1 times



  **EllenLiu** 4 months, 2 weeks ago

**Selected Answer: A**

A: dynamoDB global table preview strong consistent

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/multi-region-strong-consistency-gt.html>

upvoted 1 times


  **Penjerla** 4 months, 3 weeks ago

**Selected Answer: A**

Aurora replicas have to be in the same region as primary

<https://docs.aws.amazon.com/prescriptive-guidance/latest/aurora-replication-options/compare-solutions.html>

upvoted 2 times

  **FlyingHawk** 4 months, 3 weeks ago

**Selected Answer: B**

Since DynamoDB Global Tables allow writes in all Regions, conflicts can occur if the same item is updated simultaneously in multiple Regions.



DynamoDB uses a "last writer wins" conflict resolution strategy, which may lead to data integrity issues:

[https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables\\_HowItWorks.html#V2globaltables\\_HowItWorks.conflict-resolution](https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_HowItWorks.html#V2globaltables_HowItWorks.conflict-resolution)

If you require less than 1-second latency for write updates across Regions, DynamoDB may not consistently meet this requirement due to the asynchronous nature of its replication.

For this scenario, DynamoDB Global Tables do not fully meet the requirements due to their reliance on eventual consistency. If strong consistency is a must, go with Amazon Aurora Global Database.

upvoted 2 times

  **MatAlves** 7 months, 2 weeks ago

**Selected Answer: B**

The question asks "Average latency must be less than 1 second on updates to the reservation database."


A is incorrect:

" Changes to a DynamoDB global tables are replicated asynchronously, with typical latency of between 0.5 - 2.5 seconds between AWS Regions in the same geographic area."

B is the answer:

"All Aurora Replicas return the same data for query results with minimal replica lag. This lag is usually much less than 100 milliseconds after the primary instance has written an update."



upvoted 2 times

  **MatAlves** 7 months, 3 weeks ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>

<https://community.aws/content/2drxEND7MtOb2bWs2J0NICGewP/ddb-globaltables-lag?lang=en>

upvoted 3 times

  **SVDK** 1 year, 2 months ago

**Selected Answer: A**

How can you update your database in the different regions with read replicas? You need to be able to read and write to the database from the different regions.

upvoted 3 times

  **Milivoje** 1 year, 3 months ago

**Selected Answer: A**

In my Opinion it is A. The reason is that Aurora Read Replicas support up to 5 Read replicas in different regions . We don't have that limitation with Dynamo DB Global tables, hence I vote for A.

upvoted 1 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: B**



Purely from the wording, seems B.

DynamoDB "usually within one second"

Aurora "usually less than one second"

Question asks for "less than one second" thus Aurora

upvoted 3 times

  **Salilgen** 4 months, 1 week ago

Answer is A.

Option B is about cross-region Aurora replicas not about Aurora Global Database.

"Cross-region Aurora replicas have a lag that depends on transaction volume. Typically, a few seconds for most systems."

<https://docs.aws.amazon.com/prescriptive-guidance/latest/aurora-replication-options/compare-solutions.html>

" In a global table, a newly written item is typically propagated to all replica tables within a second."

<https://aws.amazon.com/dynamodb/global-tables/>

upvoted 1 times

  **pentium75** 1 year, 3 months ago

We need "a single primary reservation database that is globally consistent" -> A is out (DynamoDB is eventually consistent with "last writer wins" and "usually" updates "within [not: less than] one second"). D is out because it mentions multiple databases (and RDS Event Streams to not guarantee the order of events).

C is out because RDS has higher replication delay, only Aurora can guarantee "less than one second". So we'd have "a single primary reservation database that is globally consistent" in one region, and we'd have read replicas with "less than 1 second on updates" latency in other regions.

upvoted 5 times

  **numark** 1 year, 5 months ago

"a web application for travel ticketing". This would be a transaction, so DynamoDB is not the answer.

upvoted 1 times

  **pentium75** 1 year, 4 months ago



So you can't write to DynamoDB tables at all because tables writes are transactions?

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

There are no assumptions about the application here. The choices are related to the database that has one primary source of truth but multi-region presence. No requirement for transaction is given or implied.

upvoted 2 times

  **Goutham4981** 1 year, 5 months ago

**Selected Answer: A**

Dynamo DB global table acts as a single table. It does not consist of primary and standby databases. It is one single global table which is synchronously updated. Users can write to any of the regional endpoints and the write will be automatically updated across regions. To have a single primary database that is consistent does not align with dynamo db global tables.

Option B is even more dumb compared to A since read replicas does not provide failover capability or fast updates from the primary database.

The answer almost close to the requirement is Option A even though it is a misfit

upvoted 1 times

A company has migrated multiple Microsoft Windows Server workloads to Amazon EC2 instances that run in the us-west-1 Region. The company manually backs up the workloads to create an image as needed.

In the event of a natural disaster in the us-west-1 Region, the company wants to recover workloads quickly in the us-west-2 Region. The company wants no more than 24 hours of data loss on the EC2 instances. The company also wants to automate any backups of the EC2 instances.

Which solutions will meet these requirements with the LEAST administrative effort? (Choose two.)

- A. Create an Amazon EC2-backed Amazon Machine Image (AMI) lifecycle policy to create a backup based on tags. Schedule the backup to run twice daily. Copy the image on demand.
- B. Create an Amazon EC2-backed Amazon Machine Image (AMI) lifecycle policy to create a backup based on tags. Schedule the backup to run twice daily. Configure the copy to the us-west-2 Region.
- C. Create backup vaults in us-west-1 and in us-west-2 by using AWS Backup. Create a backup plan for the EC2 instances based on tag values. Create an AWS Lambda function to run as a scheduled job to copy the backup data to us-west-2.
- D. Create a backup vault by using AWS Backup. Use AWS Backup to create a backup plan for the EC2 instances based on tag values. Define the destination for the copy as us-west-2. Specify the backup schedule to run twice daily.
- E. Create a backup vault by using AWS Backup. Use AWS Backup to create a backup plan for the EC2 instances based on tag values. Specify the backup schedule to run twice daily. Copy on demand to us-west-2.

Correct Answer: BD

Community vote distribution

BD (100%)

- cloudenthusiast

Highly Voted

1 year, 5 months ago

Selected Answer: BD

Option B suggests using an EC2-backed Amazon Machine Image (AMI) lifecycle policy to automate the backup process. By configuring the policy to run twice daily and specifying the copy to the us-west-2 Region, the company can ensure regular backups are created and copied to the alternate region.

Option D proposes using AWS Backup, which provides a centralized backup management solution. By creating a backup vault and backup plan based on tag values, the company can automate the backup process for the EC2 instances. The backup schedule can be set to run twice daily, and the destination for the copy can be defined as the us-west-2 Region.

upvoted 10 times
- cloudenthusiast

1 year, 5 months ago

Both options automate the backup process and include copying the backups to the us-west-2 Region, ensuring data resilience in the event of a disaster. These solutions minimize administrative effort by leveraging automated backup and copy mechanisms provided by AWS services.

upvoted 7 times
- awsgeek75

Highly Voted

9 months, 2 weeks ago

Selected Answer: BD

LEAST admin overhead:  
A: On demand so wrong  
C: Lambda is overhead  
E: On-demand is wrong

BD is the only choice. Although D seems to cover for B also, happy to be corrected.

upvoted 7 times
- pmlabs

Most Recent

1 year, 1 month ago

B D seems to meet the requiremnts fully

upvoted 2 times
- Guru4Cloud

1 year, 2 months ago

Selected Answer: BD

B and D are the options that meet the requirements with the least administrative effort.

B uses EC2 image lifecycle policies to automatically create AMIs of the instances twice daily and copy them to the us-west-2 region. This automates regional backups.

D leverages AWS Backup to define a backup plan that runs twice daily and copies backups to us-west-2. AWS Backup automates EC2 instance backups.

Together, these options provide automated, regional EC2 backup capabilities with minimal administrative overhead.  
upvoted 2 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: BD**

options B and D will provide least administrative effort.  
upvoted 2 times

  **antropaws** 1 year, 5 months ago

**Selected Answer: BD**

I also vote B and D.  
upvoted 2 times

  **nosense** 1 year, 5 months ago

**Selected Answer: BD**

solutions are both automated and require no manual intervention to create or copy backups  
upvoted 5 times

A company operates a two-tier application for image processing. The application uses two Availability Zones, each with one public subnet and one private subnet. An Application Load Balancer (ALB) for the web tier uses the public subnets. Amazon EC2 instances for the application tier use the private subnets.

Users report that the application is running more slowly than expected. A security audit of the web server log files shows that the application is receiving millions of illegitimate requests from a small number of IP addresses. A solutions architect needs to resolve the immediate performance problem while the company investigates a more permanent solution.

What should the solutions architect recommend to meet this requirement?

- A. Modify the inbound security group for the web tier. Add a deny rule for the IP addresses that are consuming resources.
- B. Modify the network ACL for the web tier subnets. Add an inbound deny rule for the IP addresses that are consuming resources.
- C. Modify the inbound security group for the application tier. Add a deny rule for the IP addresses that are consuming resources.
- D. Modify the network ACL for the application tier subnets. Add an inbound deny rule for the IP addresses that are consuming resources.

**Correct Answer:** B

*Community vote distribution*

B (90%)

10%

- lucdt4**

Highly Voted

1 year, 5 months ago

Selected Answer: B

A wrong because security group can't deny (only allow)

upvoted 26 times
- cloudenthusiast**

Highly Voted

1 year, 5 months ago

Selected Answer: B

In this scenario, the security audit reveals that the application is receiving millions of illegitimate requests from a small number of IP addresses. To address this issue, it is recommended to modify the network ACL (Access Control List) for the web tier subnets.

By adding an inbound deny rule specifically targeting the IP addresses that are consuming resources, the network ACL can block the illegitimate traffic at the subnet level before it reaches the web servers. This will help alleviate the excessive load on the web tier and improve the application's performance.

upvoted 10 times
- awsgeek75**

Most Recent

9 months, 2 weeks ago

Selected Answer: B

A: Wrong as SG cannot deny. By default everything is deny in SG and you allow stuff

CD: App tier is not under attack so these are irrelevant options

B: Correct as NACL is exactly for this access control list to define rules for CIDR or IP addresses

upvoted 3 times
- Besisco**

2 months, 4 weeks ago

agree, even if an app tier was under attack, the block rule should be applied at a web tier ACL

upvoted 1 times
- TariqKipkemei**

12 months ago

Selected Answer: B

Modify the network ACL for the web tier subnets. Add an inbound deny rule for the IP addresses that are consuming resources.

upvoted 3 times
- potomac**

12 months ago

Selected Answer: B

A is wrong

Security groups act at the network interface level, not the subnet level, and they support Allow rules only.

upvoted 3 times
- Devsin2000**

1 year, 1 month ago

Selected Answer: A

The security Group can be applied to an ALB at web tier.

upvoted 1 times

  **OSHOAIB** 10 months ago

Security group rules are always permissive; you can't create rules that deny access.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group-rules.html>  
upvoted 3 times

  **Goutham4981** 11 months, 3 weeks ago

Security group can't deny.  
upvoted 3 times



  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

Since the bad requests are targeting the web tier, adding ACL deny rules for those IP addresses on the web subnets will block the traffic before it reaches the instances.



Security group changes (Options A and C) would not be effective since the requests are not even reaching those resources.

Modifying the application tier ACL (Option D) would not stop the bad traffic from hitting the web tier.  
upvoted 3 times

  **fakrap** 1 year, 5 months ago

**Selected Answer: B**

A is wrong because you cannot put any deny in security group  
upvoted 3 times

  **Rob1L** 1 year, 5 months ago


**Selected Answer: B**

You cannot Deny on SG, so it's B  
upvoted 6 times

  **nosense** 1 year, 5 months ago

**Selected Answer: A**

Option B is not as effective as option A  
upvoted 5 times

  **y0** 1 year, 5 months ago

Security group only have allow rules  
upvoted 3 times

  **nosense** 1 year, 5 months ago

yeah, my mistake. B should be  
upvoted 2 times

  **cloudenthusiast** 1 year, 5 months ago

A and C out due to the fact that SG does not have deny on allow rules.  
upvoted 3 times

A global marketing company has applications that run in the ap-southeast-2 Region and the eu-west-1 Region. Applications that run in a VPC in eu-west-1 need to communicate securely with databases that run in a VPC in ap-southeast-2.

Which network design will meet these requirements?

- A. Create a VPC peering connection between the eu-west-1 VPC and the ap-southeast-2 VPC. Create an inbound rule in the eu-west-1 application security group that allows traffic from the database server IP addresses in the ap-southeast-2 security group.
- B. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west-1 VPC. Update the subnet route tables. Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1.
- C. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west-1 VPC. Update the subnet route tables. Create an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses.
- D. Create a transit gateway with a peering attachment between the eu-west-1 VPC and the ap-southeast-2 VPC. After the transit gateways are properly peered and routing is configured, create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1.

**Correct Answer: C**

*Community vote distribution*

C (89%)

11%

- VellaDevil**

Highly Voted

1 year, 10 months ago

Selected Answer: C

Answer: C --> "You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC."  
<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>  
upvoted 49 times
- MatAlves**

7 months, 3 weeks ago

Wow, big thanks!

upvoted 2 times
- hsinchang**

1 year, 9 months ago

Thanks for this clarification!

upvoted 4 times
- Axeashes**

Highly Voted

1 year, 10 months ago

Selected Answer: C

"You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC."  
<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>  
upvoted 12 times
- potomac**

Most Recent

1 year, 6 months ago

Selected Answer: C

After establishing the VPC peering connection, the subnet route tables need to be updated in both VPCs to route traffic to the other VPC's CIDR blocks through the peering connection.  
upvoted 3 times
- Bennyboy789**

1 year, 8 months ago

Selected Answer: C

VPC Peering Connection: This allows communication between instances in different VPCs as if they are on the same network. It's a straightforward approach to connect the two VPCs.  
  
Subnet Route Tables: After establishing the VPC peering connection, the subnet route tables need to be updated in both VPCs to route traffic to the other VPC's CIDR blocks through the peering connection.  
  
Inbound Rule in Database Security Group: By creating an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses, you ensure that only the specified application servers from the eu-west-1 VPC can access the database servers in the ap-southeast-2 VPC.  
upvoted 4 times
- Guru4Cloud**

1 year, 8 months ago



Selected Answer: B



B) Configure VPC peering between ap-southeast-2 and eu-west-1 VPCs. Update routes. Allow traffic in ap-southeast-2 database SG from eu-west-1 application server SG.



This option establishes the correct network connectivity for the applications in eu-west-1 to reach the databases in ap-southeast-2:



VPC peering connects the two VPCs across regions - [https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html#:~:text=You%20can%20create%20a%20VPC,%2DRegion%20VPC%20peering%20connection\).](https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html#:~:text=You%20can%20create%20a%20VPC,%2DRegion%20VPC%20peering%20connection).)

Updating route tables enables routing between the VPCs  
Security group rule allowing traffic from eu-west-1 application server SG to ap-southeast-2 database SG secures connectivity  
upvoted 2 times



  **awsgeek75** 1 year, 3 months ago  
No, you cannot use a SG reference from another region so last part "Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1" cannot be setup. This is why B is wrong.  
upvoted 3 times



  **Guru4Cloud** 1 year, 8 months ago  
Options A, C, D have flaws:  
Option A peer direction is wrong  
Option C opens databases to application server IP addresses rather than SG  
Option D uses transit gateway which is unnecessary for just two VPCs  
upvoted 2 times



  **TariqKipkemei** 1 year, 9 months ago  
**Selected Answer: C**  
Selected C but B can also work  
upvoted 1 times



  **TariqKipkemei** 1 year, 9 months ago  
I just tried from the the console, You can specify the name or ID of another security group in the same region. To specify a security group in another AWS account (EC2-Classic only), prefix it with the account ID and a forward slash, for example: 111122223333/OtherSecurityGroup. You can Specify a single IP address, or an IP address range in CIDR notation in the same/other region.



In the exam both option B and C would be a pass. In the real world both option will work.  
upvoted 3 times



  **TariqKipkemei** 1 year, 5 months ago  
Correction, You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC. The C is the only option here.  
  
<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html#:~:text=You%20cannot-,reference,-the%20security%20group>  
upvoted 4 times



  **awsgeek75** 1 year, 3 months ago  
This is why B is wrong. You can never access cross region security group id  
upvoted 1 times



  **Chris22usa** 1 year, 10 months ago  
I realize D is right as ChatGpt indicates.Because here is not a problem just one application in a VPC connection to another in different region. Actually there many applications in different VPCs in a region which need to connect any other application crossingly in other region. So two transit gateway need to installed in two regions for multiple to multiple VPCs connections.  
upvoted 2 times

  **Iragmt** 1 year, 10 months ago  
However, there was also a part of "create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1"  
  
therefore, still C because we cannot reference SG ID of diff VPC, we should use the CIDR block  
upvoted 1 times

  **Chris22usa** 1 year, 10 months ago  
post it on ChaptGpt and it give me answer D. what heck with this?  
upvoted 2 times

  **haoAWS** 1 year, 10 months ago  
**Selected Answer: C**  
B is wrong because It is in a different region, so reference to the security group ID will not work. A is wrong because you need to update the route table. The answer should be C.  
upvoted 1 times

  **mattcl** 1 year, 10 months ago  
is B. what happens if application server IP addresses changes (Option C). You must change manually the IP in the security group again.  
upvoted 1 times



  **antropaws** 1 year, 10 months ago  
**Selected Answer: C**

I thought B, but I vote C after checking Axeashes response.  
upvoted 1 times

  **HelioNeto** 1 year, 11 months ago

**Selected Answer: C**



I think the answer is C because the security groups are in different VPCs. When the question wants to allow traffic from app vpc to database vpc i think using peering connection you will be able to add the security groups rules using private ip addresses of app servers. I don't think the database VPC will identify the security group id of another VPC.  
upvoted 2 times

  **REzirezi** 1 year, 11 months ago



D You cannot create a VPC peering connection between VPCs in different regions.  
upvoted 3 times

  **[Removed]** 1 year, 11 months ago

You can peer any two VPCs in different Regions, as long as they have distinct, non-overlapping CIDR blocks  
<https://docs.aws.amazon.com/devicefarm/latest/developerguide/amazon-vpc-cross-region.html>  
upvoted 3 times

  **fakrap** 1 year, 11 months ago

You can peer any two VPCs in different Regions, as long as they have distinct, non-overlapping CIDR blocks. This ensures that all of the private IP addresses are unique, and it allows all of the resources in the VPCs to address each other without the need for any form of network address translation (NAT).  
upvoted 2 times

  **nosense** 1 year, 11 months ago

**Selected Answer: B**

b for me. bcs correct inbound rule, and not overhead  
upvoted 2 times

  **cloudenthusiast** 1 year, 11 months ago

**Selected Answer: B**

Option B suggests configuring a VPC peering connection between the ap-southeast-2 VPC and the eu-west-1 VPC. By establishing this peering connection, the VPCs can communicate with each other over their private IP addresses.

Additionally, updating the subnet route tables is necessary to ensure that the traffic destined for the remote VPC is correctly routed through the VPC peering connection.

To secure the communication, an inbound rule is created in the ap-southeast-2 database security group. This rule references the security group ID of the application servers in the eu-west-1 VPC, allowing traffic only from those instances. This approach ensures that only the authorized application servers can access the databases in the ap-southeast-2 VPC.  
upvoted 5 times

A company is developing software that uses a PostgreSQL database schema. The company needs to configure multiple development environments and databases for the company's developers. On average, each development environment is used for half of the 8-hour workday.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure each development environment with its own Amazon Aurora PostgreSQL database
- B. Configure each development environment with its own Amazon RDS for PostgreSQL Single-AZ DB instances
- C. Configure each development environment with its own Amazon Aurora On-Demand PostgreSQL-Compatible database
- D. Configure each development environment with its own Amazon S3 bucket by using Amazon S3 Object Select

**Correct Answer: C**

*Community vote distribution*

C (67%)

B (33%)

  **cloudenthusiast** Highly Voted 1 year, 11 months ago

**Selected Answer: C**



Option C suggests using Amazon Aurora On-Demand PostgreSQL-Compatible databases for each development environment. This option provides the benefits of Amazon Aurora, which is a high-performance and scalable database engine, while allowing you to pay for usage on an on-demand basis. Amazon Aurora On-Demand instances are typically more cost-effective for individual development environments compared to the provisioned capacity options.

upvoted 12 times

  **cloudenthusiast** 1 year, 11 months ago

Option B suggests using Amazon RDS for PostgreSQL Single-AZ DB instances for each development environment. While Amazon RDS is a reliable and cost-effective option, it may have slightly higher costs compared to Amazon Aurora On-Demand instances.

upvoted 6 times

  **lragmt** 1 year, 9 months ago

I'm thinking that it should be B, since question does not mention any requirement only cost effective and this is just an development environment I guess we can leverage the use of RDS free tier also

upvoted 3 times

  **Stranko** Highly Voted 1 year, 2 months ago

**Selected Answer: C**

Guys, when you use the pricing calculator the cost between option B and C is really close. I doubt anyone wants to test on your knowledge of exact pricings in your region. I think that "On Demand" being explicitly specified in option C and not being specified in option B is the main difference here the exam wants to test. In that case I'd assume that option B means a constantly running instance and not "On Demand" which would make the choice pretty obvious. Again, I don't think AWS exam will test you on knowing that a single AZ is cheaper by 0,005 cents than Aurora :D

upvoted 11 times

  **tch** Most Recent 1 month, 3 weeks ago

**Selected Answer: C**

There is no upfront commitment with Aurora. You pay an hourly charge for each instance that you launch, and when you're finished with an Aurora DB instance, you can delete it. You do not need to overprovision storage as a safety margin, and you only pay for the storage you actually consume- from AWS official link

upvoted 1 times

  **zdi561** 3 months, 1 week ago

**Selected Answer: B**

B RDS single az is the cheapest. Aurora needs at least two azs.

upvoted 1 times

  **tonybuivannghia** 6 months, 2 weeks ago

**Selected Answer: C**

I choose C because Aurora on-Demand is Aurora Serverless:

The Aurora Serverless is cost effective. Scale out fine-grained increments to provide just the right number of database resources and pay only for capacity consumed.

upvoted 2 times

  **a7md0** 10 months ago

**Selected Answer: B**

Single-AZ DB instances cheaper

upvoted 1 times

🗨️ 👤 **trinh\_le** 1 year ago

**Selected Answer: B**

Single AZ more cost effective  
upvoted 2 times

🗨️ 👤 **chasingsummer** 1 year, 3 months ago

**Selected Answer: B**

1 instance(s) x 0.245 USD hourly x (4 / 24 hours in a day) x 730 hours in a month = 29.8083 USD ---> Amazon RDS PostgreSQL instances cost (monthly)  
1 instance(s) x 0.26 USD hourly x (4 / 24 hours in a day) x 730 hours in a month = 31.6333 USD ---> Amazon Aurora PostgreSQL-Compatible DB instances cost (monthly)  
upvoted 2 times

🗨️ 👤 **upliftinghut** 1 year, 3 months ago

**Selected Answer: C**

C is correct because B is cheaper but they don't mention to stop the DB when not in use  
upvoted 3 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago

**Selected Answer: C**

On-Demand is cheaper than Aurora or RDS because of low weekly usage  
upvoted 2 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: C**

We have environments that are used on average 4 hours per workday = 20 hours per week. So with option C (Aurora on-demand aka serverless) we pay for 20 hours per week. With option B (RDS) we pay for 168 hours per week (the answer does not mention anything about automating shutdown etc.).

So even if Aurora Serverless is slightly more expensive than RDS, C is cheaper because we pay only 20 (not 168) hours per week.  
upvoted 4 times

🗨️ 👤 **Salilgen** 4 months, 1 week ago

Question doesn't talk about serverless. Moreover, you can select on-demand pricing model for RDS for PostgreSQL Single-AZ too.  
IMO answer is B  
upvoted 1 times

🗨️ 👤 **Mikado211** 1 year, 4 months ago

**Selected Answer: B**

Aurora on demand is (a little) more expensive than Aurora  
Aurora is more expensive than RDS single instance

So cost effectiveness == RDS.

(B)

upvoted 1 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

But if you use the database only 20 hours per week (5 x 4), wouldn't you pay way less with Aurora serverless than with RDS?  
upvoted 3 times

🗨️ 👤 **Murtadhaceit** 1 year, 4 months ago

**Selected Answer: B**

AWS Services Calculator is showing B cheaper by less than a dollar for the same settings for both. I used "db.r6g.large" for RDS (Single-AZ) and Aurora and put 4 hours/day.  
upvoted 7 times

🗨️ 👤 **Stranko** 1 year, 2 months ago

I used the calculator, single AZ is cheaper for the exact same usage duration, if you pick On-Demand option for it too. In Aurora case (option C) you have "On Demand" explicitly specified, so if it has to be specified then I suppose that B option is about a constantly running instance. If B had an "On Demand" added, I'd vote B too.  
upvoted 2 times

🗨️ 👤 **JoseVincent68** 1 year, 4 months ago

**Selected Answer: B**

Amazon RDS Single AZ is cheaper than Aurora Multi-AZ  
upvoted 1 times

🗨️ 👤 **Wayne23Fang** 1 year, 6 months ago

**Selected Answer: B**

Aurora instances will cost you ~20% more than RDS MySQL Given the running hours the same.  
Also Aurora is HA.  
upvoted 1 times

  **baba365** 1 year, 7 months ago

... just trying to trick you. Aurora on demand is Aurora Serverless.  
upvoted 6 times

  **Anmol\_1010** 1 year, 6 months ago

that is good piece of infoamtion  
upvoted 1 times

  **deechean** 1 year, 8 months ago

**Selected Answer: C**

Aurora allows you to pay for the hours used. 4 hour every day, you only need 1/6 cost of 24 hours per day. You can check the Aurora pricing calculator.

upvoted 4 times

A company uses AWS Organizations with resources tagged by account. The company also uses AWS Backup to back up its AWS infrastructure resources. The company needs to back up all AWS resources.



Which solution will meet these requirements with the LEAST operational overhead?


- A. Use AWS Config to identify all untagged resources. Tag the identified resources programmatically. Use tags in the backup plan.
- B. Use AWS Config to identify all resources that are not running. Add those resources to the backup vault.
- C. Require all AWS account owners to review their resources to identify the resources that need to be backed up.
- D. Use Amazon Inspector to identify all noncompliant resources.

**Correct Answer: A**

*Community vote distribution*

A (100%)



-   **cloudenthusiast**


Highly Voted 

 1 year, 11 months ago

Selected Answer: A

This solution allows you to leverage AWS Config to identify any untagged resources within your AWS Organizations accounts. Once identified, you can programmatically apply the necessary tags to indicate the backup requirements for each resource. By using tags in the backup plan configuration, you can ensure that only the tagged resources are included in the backup process, reducing operational overhead and ensuring all necessary resources are backed up.



upvoted 5 times
-   **Gape4**

Most Recent 

 10 months, 2 weeks ago



Selected Answer: A

I will go for A. C and D doesn't make sense. B- resources not running? No

upvoted 2 times
-   **TariqKipkemei** 1 year, 5 months ago

Selected Answer: A



Use AWS config to deploy the tag rule and remediate resources that are not compliant.

upvoted 4 times
-   **Guru4Cloud** 1 year, 8 months ago

Selected Answer: A



This option has the least operational overhead:

AWS Config continuously evaluates resource configurations and can identify untagged resources  
Resources can be programmatically tagged via the AWS SDK based on Config data  
Backup plans can use tag criteria to automatically back up newly tagged resources  
No manual review or resource discovery needed

upvoted 3 times
-   **Bill1000** 1 year, 11 months ago

Selected Answer: A

Vote A

upvoted 3 times
-   **nosense** 1 year, 11 months ago

Selected Answer: A

a valid for me

upvoted 4 times

A social media company wants to allow its users to upload images in an application that is hosted in the AWS Cloud. The company needs a solution that automatically resizes the images so that the images can be displayed on multiple device types. The application experiences unpredictable traffic patterns throughout the day. The company is seeking a highly available solution that maximizes scalability.

What should a solutions architect do to meet these requirements?

- A. Create a static website hosted in Amazon S3 that invokes AWS Lambda functions to resize the images and store the images in an Amazon S3 bucket.
- B. Create a static website hosted in Amazon CloudFront that invokes AWS Step Functions to resize the images and store the images in an Amazon RDS database.
- C. Create a dynamic website hosted on a web server that runs on an Amazon EC2 instance. Configure a process that runs on the EC2 instance to resize the images and store the images in an Amazon S3 bucket.
- D. Create a dynamic website hosted on an automatically scaling Amazon Elastic Container Service (Amazon ECS) cluster that creates a resize job in Amazon Simple Queue Service (Amazon SQS). Set up an image-resizing program that runs on an Amazon EC2 instance to process the resize jobs.

**Correct Answer: A**

*Community vote distribution*

A (100%)

🗲️ 👤 **cloudenthusiast** Highly Voted 👍 1 year, 11 months ago

**Selected Answer: A**

By using Amazon S3 and AWS Lambda together, you can create a serverless architecture that provides highly scalable and available image resizing capabilities. Here's how the solution would work:

Set up an Amazon S3 bucket to store the original images uploaded by users.

Configure an event trigger on the S3 bucket to invoke an AWS Lambda function whenever a new image is uploaded.

The Lambda function can be designed to retrieve the uploaded image, perform the necessary resizing operations based on device requirements, and store the resized images back in the S3 bucket or a different bucket designated for resized images.

Configure the Amazon S3 bucket to make the resized images publicly accessible for serving to users.

upvoted 18 times

🗲️ 👤 **surajkrishnamurthy** Most Recent ⌚ 2 months, 4 weeks ago

**Selected Answer: A**

Is Static Website Read only ?

Eventhough I have selected Answer as A

upvoted 1 times

🗲️ 👤 **LeonSauveterre** 4 months, 4 weeks ago

**Selected Answer: A**

OK. A is better, although I chose D originally.

I guess it shouldn't be D simply because it overkills for this scenario and introduces higher operational overhead (even though option D can handle more complex workloads than Lambda if needed).

upvoted 1 times

🗲️ 👤 **cnureddy** 11 months, 1 week ago

How can end user upload an image to S3 bucket with static hosting. I believe it should be dynamic website (Answer D)

upvoted 2 times

🗲️ 👤 **mr123dd** 1 year, 4 months ago

image = static = S3 or cloudfront

but image is unstructured data so you dont store it in a relational database like RDS

and Step Function is not for processing

So A

upvoted 3 times

🗲️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

This meets all the key requirements:

S3 static website provides high availability and auto scaling to handle unpredictable traffic

Lambda functions invoked from the S3 site can resize images on the fly

Storing images in S3 buckets provides durability, scalability and high throughput  
Serverless approach with S3 and Lambda maximizes scalability and availability  
upvoted 2 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: A**

Scalability = S3, Lamda  
automatically resize images = Lambda  
upvoted 3 times

A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance. The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.

Which solution will allow the node to join the cluster?

- A. Grant the required permission in AWS Identity and Access Management (IAM) to the AmazonEKSNodeRole IAM role.
- B. Create interface VPC endpoints to allow nodes to access the control plane.
- C. Recreate nodes in the public subnet. Restrict security groups for EC2 nodes.
- D. Allow outbound traffic in the security group of the nodes.

**Correct Answer: B**

Community vote distribution

B (59%)

A (41%)

  **cloudenthusiast** Highly Voted  1 year, 11 months ago

**Selected Answer: B**

By creating interface VPC endpoints, you can enable the necessary communication between the Amazon EKS control plane and the nodes in private subnets. This solution ensures that the control plane maintains endpoint private access (set to true) and endpoint public access (set to false) for security compliance.

upvoted 21 times

  **y0** Highly Voted  1 year, 11 months ago

**Selected Answer: A**

Check this : <https://docs.aws.amazon.com/eks/latest/userguide/create-node-role.html>

Also, EKS does not require VPC endpoints. This is not the right use case for EKS

upvoted 21 times

  **TwinSpark** 12 months ago

correct i was going for B, but A looks better.

<https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html>

"When you enable endpoint private access for your cluster, Amazon EKS creates a Route 53 private hosted zone on your behalf and associates it with your cluster's VPC. This private hosted zone is managed by Amazon EKS, and it doesn't appear in your account's Route 53 resources."


upvoted 1 times

  **h0ng97\_spare\_002** 1 year, 1 month ago

<https://docs.aws.amazon.com/eks/latest/userguide/create-node-role.html#:~:text=Before,launched>

"Before you can launch nodes and register them into a cluster, you must create an IAM role for those nodes to use when they are launched."

upvoted 4 times

  **JA2018** 5 months, 1 week ago

which implies AmazonEKSNodeRole IAM role had already been configured..... that leaves answer B as the only viable choice

upvoted 1 times

  **bora4motion** Most Recent  2 days, 22 hours ago

**Selected Answer: B**

to me it looks like a network problem, not a "rights" problem.

upvoted 1 times

  **Dantecito** 2 months, 3 weeks ago

**Selected Answer: B**

B. "The company has received error notifications because the node cannot join the cluster" it's a network connectivity issue, not permissions.

upvoted 1 times



  **4729e6c** 3 months ago

**Selected Answer: B**

Since the company has configured the Amazon EKS control plane with private access enabled (true) and public access disabled (false), the EKS API server is only accessible through private network routes. The worker nodes, which are in private subnets, must be able to communicate with the EKS control plane for proper functionality.

When the nodes cannot join the cluster, the most common cause is that they are unable to reach the EKS API server. In this case, since the control plane is private-only, AWS recommends using interface VPC endpoints (AWS PrivateLink) to enable private communication between the nodes and the control plane.

upvoted 1 times

  **Salilgen** 4 months, 1 week ago

**Selected Answer: A**

IMO answer is A.

In private cluster you need interface VPC endpoints to necessary connect to some AWS services

(<https://docs.aws.amazon.com/eks/latest/userguide/private-clusters.html>) but EKS node always connect to control plane by EKS owned ENI.

See this: it is very clear.

<https://keetmalin.medium.com/eks-cluster-network-architecture-for-worker-nodes-635e067c8c2a>

upvoted 2 times

  **LeonSauveterre** 4 months, 4 weeks ago

**Selected Answer: B**

The question is trying to tell us:

1. "private access = true" and "public access = false". So the control plane endpoint is private and only accessible from within the VPC.
2. Nodes (data plane) are in private subnets.

Option A is NECESSARY but lack of permissions would generally cause authorization errors, not connectivity errors (that lead to failure of joining the cluster). Only after you got this right, you would likely receive errors about joining. So apparently we have already configured auth correctly, meaning A is not the answer.

Option C exposes the nodes to the internet. Wrong.

Option D is important for nodes to communicate with AWS services (because of allowing outbound traffic), but it's not sufficient if the required VPC interface endpoints are not even there.

upvoted 1 times



  **a7md0** 10 months ago

**Selected Answer: A**

AmazonEKSNodeRole IAM role

<https://docs.aws.amazon.com/eks/latest/userguide/create-node-role.html>

upvoted 2 times

  **emakid** 10 months, 1 week ago

**Selected Answer: B**

When Amazon EKS nodes cannot join the cluster, especially when the control plane is set to private access only, the issue typically revolves around networking and connectivity. When the EKS control plane is configured with private access only, the nodes must communicate with the control plane over private IP addresses. Creating VPC endpoints (specifically, `com.amazonaws.<region>.eks`) allows traffic between the EKS nodes and the control plane to be routed privately within the VPC, which resolves the connectivity issue.



upvoted 3 times

  **Gape4** 10 months, 2 weeks ago

**Selected Answer: B**

I think is B.


upvoted 2 times

  **MandAsh** 10 months, 3 weeks ago

**Selected Answer: B**

Error they have mentioned is at network level. They are not saying authorisation is failed rather node is unable to connect to cluster aka connectivity issue. So answer it must be B

upvoted 2 times

  **Rocconno** 10 months, 3 weeks ago

**Selected Answer: B**

<https://docs.aws.amazon.com/eks/latest/userguide/private-clusters.html>

"Any self-managed nodes must be deployed to subnets that have the VPC interface endpoints that you require. If you create a managed node group, the VPC interface endpoint security group must allow the CIDR for the subnets, or you must add the created node security group to the VPC interface endpoint security group."

upvoted 2 times

  **stalk98** 12 months ago

I Think is A

upvoted 1 times

  **trinh\_le** 1 year ago

**Selected Answer: B**



B is good to go

upvoted 3 times

  **JackyCCK** 1 year ago

S3/DynamoDB - VPC endpoint, other service should use interface endpoint so B is incorrect

upvoted 2 times

  **bujuman** 1 year, 1 month ago

**Selected Answer: B**

Because of these two assertions:

- Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance.  
( The company must also put the data plane in private subnets.

The best answer is related to Networking, Private Subnets (EKS Ctr Plane is strictly private and Data Plane stick under private subnets) and not related to EKS autodeployment that sure need an IAM policy. So according to me, answer B is the best answer.

upvoted 3 times

  **potomac** 1 year, 6 months ago

**Selected Answer: A**

Before can launch nodes and register nodes into a EKS cluster, must create an IAM role for those nodes to use when they are launched.

upvoted 2 times

A company is migrating an on-premises application to AWS. The company wants to use Amazon Redshift as a solution.

Which use cases are suitable for Amazon Redshift in this scenario? (Choose three.)

- A. Supporting data APIs to access data with traditional, containerized, and event-driven applications
- B. Supporting client-side and server-side encryption
- C. Building analytics workloads during specified hours and when the application is not active
- D. Caching data to reduce the pressure on the backend database
- E. Scaling globally to support petabytes of data and tens of millions of requests per minute
- F. Creating a secondary replica of the cluster by using the AWS Management Console

**Correct Answer:** BCE

Community vote distribution

BCE (60%)

ACE (18%)

7%

Other

  **elmogy** Highly Voted 1 year, 11 months ago

**Selected Answer:** BCE

Amazon Redshift is a data warehouse solution, so it is suitable for:

- Supporting encryption (client-side and server-side)
- Handling analytics workloads, especially during off-peak hours when the application is less active
- Scaling to large amounts of data and high query volumes for analytics purposes

The following options are incorrect because:

A) Data APIs are not typically used with Redshift. It is more for running SQL queries and analytics.

D) Redshift is not typically used for caching data. It is for analytics and data warehouse purposes.

F) Redshift clusters do not create replicas in the management console. They are standalone clusters. you could create DR cluster from snapshot and restore to another region (automated or manual) but I do not think this what is meant in this option.

upvoted 21 times

  **pentium75** 1 year, 4 months ago

"Data APIs are not typically used with Redshift" -> "With the Data API, you can programmatically access data in your Amazon Redshift cluster from different AWS services such as AWS Lambda, Amazon SageMaker notebooks, AWS Cloud9, and also your on-premises applications using the AWS SDK. This allows you to build cloud-native, containerized, serverless, web-based, and event-driven applications on the AWS Cloud."

upvoted 2 times

  **Rob1L** Highly Voted 1 year, 11 months ago

**Selected Answer:** BCE

B. Supporting client-side and server-side encryption: Amazon Redshift supports both client-side and server-side encryption for improved data security.

C. Building analytics workloads during specified hours and when the application is not active: Amazon Redshift is optimized for running complex analytic queries against very large datasets, making it a good choice for this use case.

E. Scaling globally to support petabytes of data and tens of millions of requests per minute: Amazon Redshift is designed to handle petabytes of data, and to deliver fast query and I/O performance for virtually any size dataset.

upvoted 6 times

  **Dantecito** Most Recent 2 months, 3 weeks ago

**Selected Answer:** ABC

Chatgpt says ABC.

D. Not for caching.

E. It is not optimized for high-volume like millions of requests per minute.

F. Does not natively support creating a secondary replica of the cluster.

upvoted 1 times

  **LeonSauveterre** 4 months, 4 weeks ago

**Selected Answer:** BCE

YOU MUST KNOW: Amazon Redshift is a cloud-based data warehouse service designed for analytics and reporting on large datasets, optimized for complex queries and analytical workloads rather than transactional or real-time applications.



After you memorized that, option A is incorrect. If you want to work on operational workloads, then Aurora, Amazon RDS, DynamoDB, and such, are better choices.

Option D: Caching is not a core use case for Redshift. You can choose DAX (Amazon DynamoDB Accelerator) or ElastiCache to mitigate database

pressure.

Option E: I choose that. BUT, Redshift is not designed for \*high-concurrency\* transactional workloads like "10s of mils of reqs/min", although it CAN handle massive datasets. Well... It seems partially true to me compared to other definitely wrong options.

upvoted 2 times

  **rpmaws** 7 months, 3 weeks ago

**Selected Answer: ACE**

B is not correct, how it can do encryption at client side ?

upvoted 2 times

  **3bdf1cc** 10 months ago

Found this related to A -- but specific to Redshift Serverless - but should qualify as a Redshift use case  
The Data API enables you to seamlessly access data from Redshift Serverless with all types of traditional, cloud-native, and containerized serverless web service-based applications and event-driven applications.

<https://www.amazonaws.cn/en/blog-selection/use-the-amazon-redshift-data-api-to-interact-with-amazon-redshift-serverless/>

upvoted 1 times

  **NSA\_Poker** 11 months ago

**Selected Answer: BCE**

The following are obviously incorrect:  
(D) Redshift is not as suitable as ElastiCache for caching.  
(F) A secondary replica of the cluster is not supported.

The debate is between BCE & ACE or simplified, between A & C.  
(A) is incorrect bc there is a difference btw Amazon Redshift Data API & API Gateway. API Gateway supports containerized and serverless workloads, as well as web applications. Amazon Redshift Data API is a built in API to access Redshift data with web services-based applications, including AWS Lambda, Amazon SageMaker notebooks, and AWS Cloud9.  
<https://aws.amazon.com/blogs/big-data/build-a-serverless-analytics-application-with-amazon-redshift-and-amazon-api-gateway/>

(B) is correct. You have the following options of protecting data at rest in Amazon Redshift. Use server-side encryption OR use client-side encryption

<https://docs.aws.amazon.com/redshift/latest/mgmt/security-encryption.html>

upvoted 2 times

  **JackyCCK** 1 year ago

Redshift is OLAP(online analytical processing) so D is wrong, "when the application is not active"

upvoted 2 times

  **JackyCCK** 1 year ago

\*C is wrong

upvoted 1 times


  **awsgeek75** 1 year, 3 months ago

**Selected Answer: ACE**

A, C, E are for data and Redshift is data warehouse.  
B is too generic of a choice  
D caching is not the main purpose of Redshift  
F replication is not main use of Redshift

CE are easy  
Between AB, I chose A because Redshift supports data API and client-side encryption is not Redshift specific

upvoted 3 times

  **1rob** 1 year, 3 months ago

**Selected Answer: ABD**

A: source <https://aws.amazon.com/blogs/big-data/using-the-amazon-redshift-data-api-to-interact-with-amazon-redshift-clusters/>  
B: source: <https://docs.aws.amazon.com/redshift/latest/mgmt/security-encryption.html>  
C: not sure, but you can configure scheduled queries, but the remark " and when the application is not active " , that is not relevant.  
D: source [https://docs.aws.amazon.com/redshift/latest/dg/c\\_challenges\\_achieving\\_high\\_performance\\_queries.html](https://docs.aws.amazon.com/redshift/latest/dg/c_challenges_achieving_high_performance_queries.html)  
E: Scaling globally is not supported; redshift is only a regional service.  
F: only read replica is supported. So not a secondary replica of the cluster.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: ABD**

A: <https://aws.amazon.com/de/blogs/big-data/get-started-with-the-amazon-redshift-data-api/>  
B: <https://docs.aws.amazon.com/redshift/latest/mgmt/security-encryption.html>  
D: [https://docs.aws.amazon.com/redshift/latest/dg/c\\_challenges\\_achieving\\_high\\_performance\\_queries.html#result-caching](https://docs.aws.amazon.com/redshift/latest/dg/c_challenges_achieving_high_performance_queries.html#result-caching)

Not C: Redshift is a Data Warehouse; you can use that for analytics, but it is not directly related to an "application"  
Not E: "Petabytes of data" yes, but "tens of millions of requests per minute" is not a typical feature of Redshift  
Nor F: Replicas are not a Redshift feature

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: ACE**

Technically both options A and B apply, this is from the links below:

A. You can access your Amazon Redshift database using the built-in Amazon Redshift Data API.

<https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html#:~:text=in%20Amazon%20Redshift-,Data%20API,-.%20Using%20this%20API>

B. You can encrypt data client-side and upload the encrypted data to Amazon Redshift. In this case, you manage the encryption process, the encryption keys, and related tools.

<https://docs.aws.amazon.com/redshift/latest/mgmt/security-encryption.html#:~:text=Use-,client%2Dside,-encryption%20%E2%80%93%20You%20can>

upvoted 2 times

  **potomac** 1 year, 6 months ago

**Selected Answer: ABC**

Amazon Redshift provides a Data API that you can use to painlessly access data from Amazon Redshift with all types of traditional, cloud-native, and containerized, serverless web services-based and event-driven applications.

Amazon Redshift supports up to 500 concurrent queries per cluster, which may be expanded by adding more nodes to the cluster.

upvoted 3 times

  **potomac** 1 year, 6 months ago

change to ABD

To reduce query runtime and improve system performance, Amazon Redshift caches the results of certain types of queries in memory on the leader node. When a user submits a query, Amazon Redshift checks the results cache for a valid, cached copy of the query results. If a match is found in the result cache, Amazon Redshift uses the cached results and doesn't run the query. Result caching is transparent to the user.

upvoted 1 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: BCE**



The key use cases for Amazon Redshift that fit this scenario are:

B) Redshift supports both client-side and server-side encryption to protect sensitive data.

C) Redshift is well suited for running batch analytics workloads during off-peak times without affecting OLTP systems.

E) Redshift can scale to massive datasets and concurrent users to support large analytics workloads.

upvoted 3 times

  **cd93** 1 year, 8 months ago



**Selected Answer: BCD**

Why E lol? It's a data warehouse! it has no need to support millions of requests, it is not mentioned anywhere (<https://aws.amazon.com/redshift/features>)

In fact Redshift editor supports max 500 connections and workgroup support max 2000 connections at once, see it's quota page

Redshift has a cache layer, D is correct

upvoted 3 times

  **mrsoa** 1 year, 9 months ago

**Selected Answer: BCE**

BCE, For B this is why

<https://docs.aws.amazon.com/redshift/latest/mgmt/security-encryption.html>

upvoted 2 times

  **james2033** 1 year, 9 months ago

**Selected Answer: ACE**

Quote: "The Data API enables you to seamlessly access data from Redshift Serverless with all types of traditional, cloud-native, and containerized serverless web service-based applications and event-driven applications." at <https://aws.amazon.com/blogs/big-data/use-the-amazon-redshift-data-api-to-interact-with-amazon-redshift-serverless/> (28/4/2023). Choose A. B and C are next chosen correct answers.

upvoted 2 times

  **james2033** 1 year, 9 months ago

Typo, I want said "C and E are next chosen correct answers."

upvoted 2 times

A company provides an API interface to customers so the customers can retrieve their financial information. The company expects a larger number of requests during peak usage times of the year.

The company requires the API to respond consistently with low latency to ensure customer satisfaction. The company needs to provide a compute host for the API.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use an Application Load Balancer and Amazon Elastic Container Service (Amazon ECS).
- B. Use Amazon API Gateway and AWS Lambda functions with provisioned concurrency.
- C. Use an Application Load Balancer and an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- D. Use Amazon API Gateway and AWS Lambda functions with reserved concurrency.

**Correct Answer: B**

Community vote distribution

B (78%)

A (23%)

  **cloudenthusiast** Highly Voted 1 year, 11 months ago

**Selected Answer: B**

In the context of the given scenario, where the company wants low latency and consistent performance for their API during peak usage times, it would be more suitable to use provisioned concurrency. By allocating a specific number of concurrent executions, the company can ensure that there are enough function instances available to handle the expected load and minimize the impact of cold starts. This will result in lower latency and improved performance for the API.

upvoted 13 times

  **Bennyboy789** Highly Voted 1 year, 8 months ago

**Selected Answer: B**

Provisioned - minimizing cold starts and providing low latency.

upvoted 7 times

  **LeonSauveterre** Most Recent 4 months, 4 weeks ago

**Selected Answer: B**

Concurrency of lambda function, provisioned vs reserved:


Reserved Concurrency:

1. Desired instances can start without interference from other functions
2. Guarantees available execution capacity for critical functions

Provisioned Concurrency:

1. Also guarantees consistent performance
2. [!!] Allows you to pre-warm a specific number of function instances to reduce cold start latency

upvoted 2 times



  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/lambda/latest/dg/lambda-concurrency.html#reserved-and-provisioned>

Consistency decreases if you exceed your provisioned instance. Let's say you have 1000 (default) provisioned instances and the load is 1500. The new 500 will have to wait until the first 1000 concurrent calls finish. This is solved by increasing the provisioned concurrency to 1500.

upvoted 2 times

  **1rob** 1 year, 5 months ago

**Selected Answer: A**

So I have my doubts here. The question also states ; "The company needs to provide a compute host for the API." Imho this implies to have some sort of physical host which has to be provided by the customer. Translating this further to aws this would mean an EC2 instance. And then when I would go for ECS instead of EKS.

Please share your opinion.

upvoted 7 times

  **pdragon1981** 1 year, 4 months ago

Exactly, initially I was thinking on B but if company must provide a host I would say that only option A is feasible

upvoted 3 times

🗨️ 👤 **pdragon1981** 1 year, 4 months ago

Sorry I understand bad the text, coorect answer is B, as for my understanding now the host is the device that the costumer needs to connect with API Gateway, bellow explains well the logic

<https://aws.amazon.com/api-gateway/>

upvoted 4 times

🗨️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

This option provides the least operational overhead:

API Gateway handles the API requests and integration with Lambda

Lambda automatically scales compute without managing servers

Provisioned concurrency ensures consistent low latency by keeping functions initialized

No need to manage containers or orchestration platforms as with ECS/EKS

upvoted 3 times

🗨️ 👤 **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: B**

The company requires the API to respond consistently with low latency to ensure customer satisfaction especially during high peak periods, there is no mention of cost efficient. Hence provisioned concurrency is the best option.

Provisioned concurrency is the number of pre-initialized execution environments you want to allocate to your function. These execution environments are prepared to respond immediately to incoming function requests. Configuring provisioned concurrency incurs charges to your AWS account.

<https://docs.aws.amazon.com/lambda/latest/dg/provisioned-concurrency.html#:~:text=for%20a%20function.-,Provisioned%20concurrency,-%E2%80%93Provisioned%20concurrency%20is>

upvoted 2 times

🗨️ 👤 **MirKhobaeb** 1 year, 11 months ago

**Selected Answer: B**

AWS Lambda provides a highly scalable and distributed infrastructure that automatically manages the underlying compute resources. It automatically scales your API based on the incoming request load, allowing it to respond consistently with low latency, even during peak times. AWS Lambda takes care of infrastructure provisioning, scaling, and resource management, allowing you to focus on writing the code for your API logic.

upvoted 4 times

A company wants to send all AWS Systems Manager Session Manager logs to an Amazon S3 bucket for archival purposes.

Which solution will meet this requirement with the MOST operational efficiency?

- A. Enable S3 logging in the Systems Manager console. Choose an S3 bucket to send the session data to.
- B. Install the Amazon CloudWatch agent. Push all logs to a CloudWatch log group. Export the logs to an S3 bucket from the group for archival purposes.
- C. Create a Systems Manager document to upload all server logs to a central S3 bucket. Use Amazon EventBridge to run the Systems Manager document against all servers that are in the account daily.
- D. Install an Amazon CloudWatch agent. Push all logs to a CloudWatch log group. Create a CloudWatch logs subscription that pushes any incoming log events to an Amazon Kinesis Data Firehose delivery stream. Set Amazon S3 as the destination.

**Correct Answer: A**

*Community vote distribution*

A (95%)

5%

  **master9** Highly Voted 1 year, 4 months ago

**Selected Answer: A**

send logs to Amazon S3 from AWS Systems Manager Session Manager. Here are the steps to do so:

Enable S3 Logging: Open the AWS Systems Manager console. In the navigation pane, choose Session Manager. Choose the Preferences tab, and then choose Edit. Select the check box next to Enable under S3 logging.

Create an S3 Bucket: To store the Session Manager logs, create an S3 bucket to hold the audit logs from the Session Manager interactive shell usage.

Configure IAM Role: AWS Systems Manager Agent (SSM Agent) uses the same AWS Identity and Access Management (IAM) role to activate itself and upload logs to Amazon S3. You can use either an IAM instance profile that's attached to an Amazon Elastic Compute Cloud (Amazon EC2) instance or the IAM role that's configured for the Default Host Management Configuration.

upvoted 8 times

  **cloudenthusiast** Highly Voted 1 year, 11 months ago

**Selected Answer: A**

option A does not involve CloudWatch, while option D does. Therefore, in terms of operational overhead, option A would generally have less complexity and operational overhead compared to option D.

Option A simply enables S3 logging in the Systems Manager console, allowing you to directly send session logs to an S3 bucket. This approach is straightforward and requires minimal configuration.

On the other hand, option D involves installing and configuring the Amazon CloudWatch agent, creating a CloudWatch log group, setting up a CloudWatch Logs subscription, and configuring an Amazon Kinesis Data Firehose delivery stream to store logs in an S3 bucket. This requires additional setup and management compared to option A.

So, if minimizing operational overhead is a priority, option A would be a simpler and more straightforward choice.

upvoted 5 times

  **pujithacg8** Most Recent 9 months ago

A, You can choose to store session log data in a specified Amazon Simple Storage Service (Amazon S3) bucket for debugging and troubleshooting purposes.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html#session-manager-logging-s3>

upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

Most efficient is A because it is a direct option in SM logging.

B can work but is more operational overhead as you end up using CloudWatch (not sure how but making assumption based on language of option)

C is definitely too much work

D Way too many moving parts

upvoted 3 times

  **potomac** 1 year, 6 months ago

**Selected Answer: A**

You can choose to store session log data in a specified Amazon Simple Storage Service (Amazon S3) bucket for debugging and troubleshooting purposes.

upvoted 2 times

  **deechean** 1 year, 8 months ago

**Selected Answer: A**

You can config the log archived to S3 in the Session Manager - > preference tab. Another option is CloudWatch log.  
<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html#session-manager-logging-s3>

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

°Simplicity - Enabling S3 logging requires just a simple configuration in the Systems Manager console to specify the destination S3 bucket. No other services need to be configured.  
°Direct integration - Systems Manager has native support to send session logs to S3 through this feature. No need for intermediary services.  
°Automated flow - Once S3 logging is enabled, the session logs automatically flow to the S3 bucket without manual intervention.  
°Easy management - The S3 bucket can be managed independently for log storage and archival purposes without impacting Systems Manager.  
°Cost-effectiveness - No charges for intermediate CloudWatch or Kinesis services. Just basic S3 storage costs.  
°Minimal overhead - No ongoing management of complex pipeline of services. Direct logs to S3 minimizes overhead.

upvoted 3 times



  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: A**

With the MOST operational efficiency then option A is best.  
Otherwise B is also an option with a little bit more ops than option A.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html>

upvoted 2 times

  **Zox42** 1 year, 10 months ago

**Selected Answer: A**

Answer A. <https://aws-labs.net/winlab5-manageinfra/sessmgrlog.html>

upvoted 2 times

  **Zuit** 1 year, 10 months ago

**Selected Answer: A**

GPT argued for D.

B could be an option, by installing a logging package on alle managed systems/ECs etc. <https://docs.aws.amazon.com/systems-manager/latest/userguide/distributor-working-with-packages-deploy.html>

However, as it mentions the "Session manager logs" I would tend towards A.



upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: A**

It should be "A".  
<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html>

upvoted 2 times

  **secdgs** 1 year, 10 months ago

**Selected Answer: A**

It have menu to Enable S3 Logging.  
<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html#session-manager-logging-s3>

upvoted 2 times

  **Markie999** 1 year, 11 months ago

**Selected Answer: B**



BBBBBBBBBB

upvoted 1 times

  **pentium75** 1 year, 4 months ago

"Install the CloudWatch agent" where?

upvoted 2 times

  **Bill1000** 1 year, 11 months ago

**Selected Answer: B**

The option 'A' says "Enable S3 logging in the Systems Manager console." This means that you will enable the logs !! FOR !! S3 events and its is not what the question asks. My vote is for Option B, based on this article: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html>

upvoted 1 times


  **baba365** 1 year, 9 months ago

To log session data using Amazon S3 (console)

Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/>.

In the navigation pane, choose Session Manager.  
Choose the Preferences tab, and then choose Edit.  
Select the check box next to Enable under S3 logging.

upvoted 3 times

  **vrevkov** 1 year, 10 months ago

But where do you want to install the Amazon CloudWatch agent in case of B?

upvoted 1 times

  **omoakin** 1 year, 11 months ago



DDDDDD

upvoted 1 times

  **Anmol\_1010** 1 year, 11 months ago

Option D is definitely not right,  
Its option B

upvoted 2 times

  **omoakin** 1 year, 11 months ago

Chat GPT says option A is incorrect cos it requires enabling S3 logging in the system manager console only logs information about the systems manager service not the session logs

Says correct answer is B

upvoted 1 times

  **[Removed]** 1 year, 11 months ago

Question may not be very clear. A should be the answer. Below link is the documetation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html#session-manager-logging-s3>

upvoted 5 times

An application uses an Amazon RDS MySQL DB instance. The RDS database is becoming low on disk space. A solutions architect wants to increase the disk space without downtime.

Which solution meets these requirements with the LEAST amount of effort?

- A. Enable storage autoscaling in RDS
- B. Increase the RDS database instance size
- C. Change the RDS database instance storage type to Provisioned IOPS
- D. Back up the RDS database, increase the storage capacity, restore the database, and stop the previous instance

Correct Answer: A

Community vote distribution

A (100%)

  **cloudenthusiast** Highly Voted  1 year, 11 months ago

Selected Answer: A

Enabling storage autoscaling allows RDS to automatically adjust the storage capacity based on the application's needs. When the storage usage exceeds a predefined threshold, RDS will automatically increase the allocated storage without requiring manual intervention or causing downtime. This ensures that the RDS database has sufficient disk space to handle the increasing storage requirements.  
upvoted 12 times

  **Gape4** Most Recent  10 months, 2 weeks ago

Selected Answer: A

Autoscaling.... without downtime...  
upvoted 2 times

  **potomac** 1 year, 6 months ago

Selected Answer: A

Amazon RDS for MariaDB, Amazon RDS for MySQL, Amazon RDS for PostgreSQL, Amazon RDS for SQL Server and Amazon RDS for Oracle support RDS Storage Auto Scaling. RDS Storage Auto Scaling automatically scales storage capacity in response to growing database workloads, with zero downtime.  
upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

Selected Answer: A

This question is so obvious  
upvoted 4 times

  **TariqKipkemei** 1 year, 9 months ago

Selected Answer: A

RDS Storage Auto Scaling continuously monitors actual storage consumption, and scales capacity up automatically when actual utilization approaches provisioned storage capacity. Auto Scaling works with new and existing database instances. You can enable Auto Scaling with just a few clicks in the AWS Management Console. There is no additional cost for RDS Storage Auto Scaling. You pay only for the RDS resources needed to run your applications.  
  
<https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/#:~:text=of%20the%20rest.,RDS%20Storage%20Auto%20Scaling,-continuously%20monitors%20actual>  
upvoted 3 times

  **james2033** 1 year, 9 months ago

Selected Answer: A

Quote "Amazon RDS now supports Storage Auto Scaling" and "... with zero downtime." (Jun 20th 2019) at <https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>  
upvoted 3 times

  **james2033** 1 year, 9 months ago

Hello moderator, please help me delete this discussion, I already add content before this comment.  
upvoted 1 times

  **james2033** 1 year, 9 months ago

Selected Answer: A

See "Amazon RDS now supports Storage Auto Scaling. Posted On: Jun 20, 2019. Starting today, Amazon RDS for MariaDB, Amazon RDS for MySQL, Amazon RDS for PostgreSQL, Amazon RDS for SQL Server and Amazon RDS for Oracle support RDS Storage Auto Scaling. RDS Storage Auto

Scaling automatically scales storage capacity in response to growing database workloads, with zero downtime.” at <https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

upvoted 3 times

  **haoAWS** 1 year, 10 months ago

**Selected Answer: A**

A is the best answer.

B will not work for increasing disk space, it only improve the IO performance.

C will not work because it will cause downtime.

D is too complicated and need much operational effort.

upvoted 2 times

  **[Removed]** 1 year, 11 months ago

<https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

upvoted 2 times

  **Anmol\_1010** 1 year, 11 months ago

The key word is No Down time. A would be bewt option

upvoted 3 times

A consulting company provides professional services to customers worldwide. The company provides solutions and tools for customers to expedite gathering and analyzing data on AWS. The company needs to centrally manage and deploy a common set of solutions and tools for customers to use for self-service purposes.

Which solution will meet these requirements?

- A. Create AWS CloudFormation templates for the customers.
- B. Create AWS Service Catalog products for the customers.
- C. Create AWS Systems Manager templates for the customers.
- D. Create AWS Config items for the customers.

**Correct Answer:** *B*

*Community vote distribution*

B (100%)

  **cloudenthusiast** Highly Voted  1 year, 11 months ago

**Selected Answer: B**



AWS Service Catalog allows you to create and manage catalogs of IT services that can be deployed within your organization. With Service Catalog, you can define a standardized set of products (solutions and tools in this case) that customers can self-service provision. By creating Service Catalog products, you can control and enforce the deployment of approved and validated solutions and tools.

upvoted 10 times

  **Oblako** 1 year, 5 months ago

"within your organization" => not for customers

upvoted 1 times

  **JA2018** 5 months, 1 week ago

for this use case, in order to access the self-service functions, customers must have some form of AWS identities (e.g. accounts, users, user groups or roles) already setup in the service provider's AWS organisation.

upvoted 1 times

  **Guru4Cloud** Highly Voted  1 year, 8 months ago

**Selected Answer: B**

Some key advantages of using Service Catalog:

Centralized management - Products can be maintained in a single catalog for easy discovery and governance.

Self-service access - Customers can deploy the solutions on their own without manual intervention.

Standardization - Products provide pre-defined templates for consistent deployment.

Access control - Granular permissions can be applied to restrict product visibility and access.

Reporting - Service Catalog provides detailed analytics on product usage and deployments.

upvoted 5 times

  **hsinchang** Most Recent  1 year, 9 months ago

**Selected Answer: B**

CloudFormation: a code as infrastructure service

Systems Manager: management solution for resources

Config: assess, audit and evaluate configurations

Other options does not fit this scenario.

upvoted 2 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: B**

AWS Service Catalog lets you centrally manage your cloud resources to achieve governance at scale of your infrastructure as code (IaC) templates, written in CloudFormation or Terraform. With AWS Service Catalog, you can meet your compliance requirements while making sure your customers can quickly deploy the cloud resources they need.

<https://aws.amazon.com/servicecatalog/#:~:text=How%20it%20works-,AWS%20Service%20Catalog,-lets%20you%20centrally>

upvoted 2 times

  **Yadav\_Sanjay** 1 year, 11 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/introduction.html>

upvoted 3 times

A company is designing a new web application that will run on Amazon EC2 Instances. The application will use Amazon DynamoDB for backend data storage. The application traffic will be unpredictable. The company expects that the application read and write throughput to the database will be moderate to high. The company needs to scale in response to application traffic.

Which DynamoDB table configuration will meet these requirements MOST cost-effectively?

- A. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard table class. Set DynamoDB auto scaling to a maximum defined capacity.
- B. Configure DynamoDB in on-demand mode by using the DynamoDB Standard table class.
- C. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class. Set DynamoDB auto scaling to a maximum defined capacity.
- D. Configure DynamoDB in on-demand mode by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class.



**Correct Answer: B**

*Community vote distribution*



B (74%)

A (22%)



2%

-   **Efren** Highly Voted 1 year, 11 months ago

B for me. Provisioned if we know how much traffic will come, but its unpredictable, so we have to go for on-demand



upvoted 12 times
-   **VellaDevil** 1 year, 10 months ago

Spot On

upvoted 3 times
-   **tch** Most Recent 1 month, 3 weeks ago



Selected Answer: B

DynamoDB offers two table classes designed to help you optimize for cost.  
The DynamoDB Standard table class is the default, and is recommended for the vast majority of workloads.  
The DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class is optimized for tables where storage is the dominant cost. For example, tables that store infrequently accessed data, such as application logs, old social media posts, e-commerce order history, and past gaming achievements, are good candidates for the Standard-IA table class.



upvoted 1 times
-   **fdae619** 4 months, 1 week ago

Selected Answer: B

Can someone answer why the standard class is a better choice to standard IA?

upvoted 1 times
-   **tch** 1 month, 3 weeks ago



DynamoDB offers two table classes designed to help you optimize for cost.  
The DynamoDB Standard table class is the default, and is recommended for the vast majority of workloads.  
The DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class is optimized for tables where storage is the dominant cost. For example, tables that store infrequently accessed data, such as application logs, old social media posts, e-commerce order history, and past gaming achievements, are good candidates for the Standard-IA table class.

upvoted 1 times
-   **emakid** 10 months, 1 week ago

Selected Answer: B



Configure DynamoDB in on-demand mode by using the DynamoDB Standard table class.

This option allows DynamoDB to automatically adjust to varying traffic patterns, which is ideal for unpredictable workloads. The Standard table class is suitable for applications with moderate to high read and write throughput, and on-demand mode ensures that you are billed based on the actual usage, providing cost efficiency for variable traffic patterns.

upvoted 3 times
-   **Gape4** 10 months, 2 weeks ago

Selected Answer: B

Key word : On demand. So I think B.

upvoted 2 times
-   **awsgeek75** 1 year, 3 months ago

Selected Answer: B

On demand  
<https://docs.aws.amazon.com/wellarchitected/latest/serverless-applications-lens/capacity.html>

"With on-demand capacity mode, DynamoDB charges you for the data reads and writes your application performs on your tables. You do not need to specify how much read and write throughput you expect your application to perform because DynamoDB instantly accommodates your workloads as they ramp up or down."

upvoted 3 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: B**

Not A because of "unpredictable" traffic  
Not C and D because we are expecting "moderate to high" traffic  
upvoted 4 times

  **leonliu4** 1 year, 4 months ago

**Selected Answer: B**

Leaning towards B, it's hard to predict the capacity for A, and autoscaling doesn't respond fast  
upvoted 3 times

  **peekingpicker** 1 year, 5 months ago

**Selected Answer: A**

it's A.  
remember that :  
he company expects that the application read and write throughput to the database will be moderate to high  
  
provisioned throughput is cheaper than ondemand capacity right ?  
upvoted 2 times

  **pentium75** 1 year, 4 months ago

but "unpredictable" which usually hints to on-demand  
upvoted 2 times

  **dilaaziz** 1 year, 5 months ago

**Selected Answer: D**

Data storage: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithTables.tableclasses.html>  
upvoted 1 times

  **Mr\_Marcus** 11 months, 1 week ago

Actually, this link confirmed "B" for me...  
upvoted 2 times

  **potomac** 1 year, 6 months ago

**Selected Answer: B**

On-demand mode is great for unpredictable traffic  
upvoted 3 times

  **bsbs1234** 1 year, 6 months ago

I choose B  
I think the items stored in the table in this question has large size. So each read/write, a big chunk of data pass through. A capacity unit is used to describe data throughput. provision to the high capacity units will be a waste because unpredicted traffic pattern.  
upvoted 2 times

  **Bennyboy789** 1 year, 8 months ago

**Selected Answer: B**

Unpredictable= on demand  
upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

The key factors are:

With On-Demand mode, you only pay for what you use instead of over-provisioning capacity. This avoids idle capacity costs.  
DynamoDB Standard provides the fastest performance needed for moderate-high traffic apps vs Standard-IA which is for less frequent access.  
Auto scaling with provisioned capacity can also work but requires more administrative effort to tune the scaling thresholds.  
upvoted 2 times

  **msdnpro** 1 year, 9 months ago

**Selected Answer: B**

Support for B from AWS:

On-demand mode is a good option if any of the following are true:  
-You create new tables with unknown workloads.  
-You have unpredictable application traffic.  
-You prefer the ease of paying for only what you use.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html>

upvoted 2 times



  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: B**

Technically both options A and B will work. But this statement 'traffic will be unpredictable' rules out option A, because 'provisioned mode' was made for scenarios where traffic is predictable.

So I will stick with B, because 'on-demand mode' is made for unpredictable traffic and instantly accommodates workloads as they ramp up or down.

upvoted 3 times

  **0628atv** 1 year, 9 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

upvoted 4 times

A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account.

The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.

Which authentication option will meet these requirements MOST securely?

- A. Integrate DynamoDB with AWS Secrets Manager in the inventory application account. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB table. Schedule secret rotation for every 30 days.
- B. In every business account, create an IAM user that has programmatic access. Configure the application to use the correct IAM user access key ID and secret access key to authenticate and read the DynamoDB table. Manually rotate IAM access keys every 30 days.
- C. In every business account, create an IAM role named BU\_ROLE with a policy that gives the role access to the DynamoDB table and a trust policy to trust a specific role in the inventory application account. In the inventory account, create a role named APP\_ROLE that allows access to the STS AssumeRole API operation. Configure the application to use APP\_ROLE and assume the crossaccount role BU\_ROLE to read the DynamoDB table.
- D. Integrate DynamoDB with AWS Certificate Manager (ACM). Generate identity certificates to authenticate DynamoDB. Configure the application to use the correct certificate to authenticate and read the DynamoDB table.

Correct Answer: C



Community vote distribution



C (100%)



  **cloudenthusiast** Highly Voted 1 year, 11 months ago



Selected Answer: C

IAM Roles: IAM roles provide a secure way to grant permissions to entities within AWS. By creating an IAM role in each business account named BU\_ROLE with the necessary permissions to access the DynamoDB table, the access can be controlled at the IAM role level.  
Cross-Account Access: By configuring a trust policy in the BU\_ROLE that trusts a specific role in the inventory application account (APP\_ROLE), you establish a trusted relationship between the two accounts.  
Least Privilege: By creating a specific IAM role (BU\_ROLE) in each business account and granting it access only to the required DynamoDB table, you can ensure that each team's table is accessed with the least privilege principle.  
Security Token Service (STS): The use of STS AssumeRole API operation in the inventory application account allows the application to assume the cross-account role (BU\_ROLE) in each business account.  
upvoted 31 times



  **TariqKipkemei** 1 year, 9 months ago  
Well broken down..thank you :)  
upvoted 3 times

  **MandAsh** Highly Voted 10 months, 3 weeks ago  
C because they have taken effort to explain it in details.. lol  
upvoted 6 times



  **Bennyboy789** Most Recent 1 year, 8 months ago  
Selected Answer: C  
Keyword: IAM ROLES  
upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago  
Selected Answer: C  
C is the most secure option to meet the requirements.

Using cross-account IAM roles and role chaining allows the inventory application to securely access resources in other accounts. The roles provide temporary credentials and can be permissions controlled.  
upvoted 3 times

  **hsinchang** 1 year, 9 months ago  
Selected Answer: C  
Looks complex, but IAM role seems more probable, I go with C.

upvoted 4 times

  **mattcl** 1 year, 10 months ago

Why not A?

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

A is wrong because it is incomplete. Just integrating with secrets manager doesn't give any access to DynamoDB.

upvoted 2 times

  **antropaws** 1 year, 10 months ago

**Selected Answer: C**

It's complex, but looks C.

upvoted 2 times

  **eehhssaan** 1 year, 11 months ago

i'll go with C .. coming from two minds

upvoted 3 times

  **nosense** 1 year, 11 months ago

a or c. C looks like a more secure

upvoted 2 times

  **omoakin** 1 year, 11 months ago

CCCCCCCCCCCC

upvoted 2 times

A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS). The company's workload is not consistent throughout the day. The company wants Amazon EKS to scale in and out according to the workload.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Use an AWS Lambda function to resize the EKS cluster.
- B. Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.
- C. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
- D. Use Amazon API Gateway and connect it to Amazon EKS.
- E. Use AWS App Mesh to observe network activity.

**Correct Answer:** BC

Community vote distribution

BC (100%)

  **Guru4Cloud** Highly Voted 8 months, 2 weeks ago

**Selected Answer:** BC

B and C are the correct options.

Using the Kubernetes Metrics Server (B) enables horizontal pod autoscaling to dynamically scale pods based on CPU/memory usage. This allows scaling at the application tier level.

The Kubernetes Cluster Autoscaler (C) automatically adjusts the number of nodes in the EKS cluster in response to pod resource requirements and events. This allows scaling at the infrastructure level.


upvoted 7 times

  **cloudenthusiast** Highly Voted 11 months, 2 weeks ago

**Selected Answer:** BC

By combining the Kubernetes Cluster Autoscaler (option C) to manage the number of nodes in the cluster and enabling horizontal pod autoscaling (option B) with the Kubernetes Metrics Server, you can achieve automatic scaling of your EKS cluster and container applications based on workload demand. This approach minimizes operational overhead as it leverages built-in Kubernetes functionality and automation mechanisms.

upvoted 5 times

  **wsdasdasdqwdaw** Most Recent 6 months, 1 week ago

K8S Metrics Server and Autoscaler => B and C


upvoted 4 times

  **TariqKipkemei** 9 months, 2 weeks ago

**Selected Answer:** BC

This is pretty straight forward.  
Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.  
Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.

upvoted 5 times

  **james2033** 9 months, 3 weeks ago

**Selected Answer:** BC

Kubernetes Metrics Server <https://docs.aws.amazon.com/eks/latest/userguide/metrics-server.html>

AWS Autoscaler <https://docs.aws.amazon.com/eks/latest/userguide/autoscaling.html> and <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/cloudprovider/aws/README.md>

upvoted 4 times

  **nosense** 11 months, 2 weeks ago

**Selected Answer:** BC

b and c is right  
upvoted 2 times

A company runs a microservice-based serverless web application. The application must be able to retrieve data from multiple Amazon DynamoDB tables A solutions architect needs to give the application the ability to retrieve the data with no impact on the baseline performance of the application.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. AWS AppSync pipeline resolvers
- B. Amazon CloudFront with Lambda@Edge functions
- C. Edge-optimized Amazon API Gateway with AWS Lambda functions
- D. Amazon Athena Federated Query with a DynamoDB connector

**Correct Answer: D**

*Community vote distribution*

D (48%)

A (35%)

B (17%)

- elmogy**

Highly Voted

1 year, 11 months ago

just passed yesterday 30-05-23, around 75% of the exam came from here, some with light changes.

upvoted 36 times
- omoakin**

Highly Voted

1 year, 11 months ago

Great work made it to the last question. Goodluck to you all

upvoted 20 times
- cyber\_bedouin**

1 year, 6 months ago

Thanks. Do you think the questions after 500 after relevant, they seem to be above associate level (harder)

upvoted 10 times
- Drew3000**

1 year, 1 month ago

I am jealous that 10 months ago there were only 523 questions :(

upvoted 10 times
- Awsbeginner87**

1 year, 1 month ago

right...how to go through 840 questions :(

upvoted 3 times
- Rhydian25**

10 months, 2 weeks ago

Currently 904 :')

upvoted 5 times
- 4729e6c**

3 months ago

1019 now and hard to finish them all in a month :(

upvoted 3 times
- bora4motion**

2 days, 22 hours ago

how did you do?

upvoted 1 times
- Awsbeginner87**

1 year, 1 month ago

moderators answer is wrong for many questions.which one should we consider:(

upvoted 1 times
- MostofMichelle**

1 year, 11 months ago

good luck to you as well.

upvoted 6 times
- CTao**

Most Recent

2 months, 3 weeks ago

**Selected Answer: A**

AWS AppSync: It provides a managed GraphQL service that simplifies data access across multiple sources, including DynamoDB tables.

upvoted 2 times

**4729e6c**

3 months ago

**Selected Answer: A**

Deepseek and Qwen are just published and they both say A.

upvoted 1 times

  **FlyingHawk** 3 months, 1 week ago

**Selected Answer: A**

AWS AppSync is a managed GraphQL service that allows applications to query multiple data sources, such as DynamoDB tables, in a single API call. It is highly efficient and does not require maintaining backend infrastructure.

Pipeline resolvers in AppSync enable the combination of multiple resolvers to fetch data from different sources, aggregate the data, and return it in a single response. This approach is operationally efficient, has minimal impact on baseline performance, and is designed for applications interacting with multiple DynamoDB tables.

<https://docs.aws.amazon.com/appsync/latest/devguide/tutorial-pipeline-resolvers.html>

upvoted 1 times

  **FlyingHawk** 4 months, 3 weeks ago

Why the suggested answer is D? I read the comments, feel A is more suitable as the requirement is to give the application the ability to retrieve the data with no impact on the baseline performance of the application. D might have affect the baseline performance, A is more effective

upvoted 2 times

  **LeonSauveterre** 4 months, 4 weeks ago

**Selected Answer: A**

A - Correct. Retrieves data from multiple data sources.


B - Incorrect. Lambda@Edge extends serverless processing to CloudFront, but it is typically used for content manipulation, caching, or geographic distribution.

C - That works but requires a lot of things to do manually. Why C when you can do with A?

D - Amazon Athena Federated Query is designed for ad hoc querying of data across different sources BUT it has impact on performance, and not operationally efficient.

PS: "ad hoc" means "for this specific purpose". Such requests can lack efficiency for repetitive tasks because it's not optimized or automated, since they're naturally exploratory or tailored to specific immediate needs.

upvoted 2 times

  **kevindu** 8 months, 3 weeks ago

**Selected Answer: A**

Is there anyone who has recently passed the exam who can tell me approximately how many of the original questions are in the actual exam?

upvoted 3 times

  **use4u** 1 year ago

let me know



for most of questions that system answer are correct or comment answer are correct?

upvoted 2 times

  **Sergiuss95** 1 year ago

Comment answer

upvoted 2 times

  **osmk** 1 year, 2 months ago

**Selected Answer: D**

[https://docs.amazonaws.cn/en\\_us/athena/latest/ug/connect-to-a-data-source.html](https://docs.amazonaws.cn/en_us/athena/latest/ug/connect-to-a-data-source.html)

upvoted 6 times

  **FlyingHawk** 3 months, 1 week ago

Athena queries are serverless and scalable but are not optimized for low-latency, high-concurrency scenarios typically required for a microservice-based operational web application. Additionally, query execution time and costs can become a concern if frequent access to DynamoDB tables is required.

upvoted 2 times

  **upliftinghut** 1 year, 3 months ago

**Selected Answer: D**

key word is most operational effective => D requires no coding

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: D**

I'll go with D as ABC looks too much work or irrelevant. Although not sure how AFQ actually achieves the read without impacting performance.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: D**



Not A - Pipe Resolvers require coding, would not consider that 'operationally efficient'

Not B - CloudFront caches web content at the edge, not DynamoDB query results for apps

Not C - Neither API Gateway or Lambda have anything to do with DynamoDB performance

D - Can do exactly that

upvoted 8 times

  **aws94** 1 year, 4 months ago

**Selected Answer: A**

I am not an expert but I used Bing+Gemini+Chatgpt=AAA  
upvoted 2 times

  **ekisako** 1 year, 5 months ago

**Selected Answer: A**

multiple database tables = AppSync pipeline resolvers  
upvoted 7 times

  **hungta** 1 year, 5 months ago

**Selected Answer: B**

For an operationally efficient solution that minimizes impact on baseline performance in a microservice-based serverless web application retrieving data from multiple DynamoDB tables, Amazon CloudFront with Lambda@Edge functions (Option B) is often the most suitable choice  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

CloudFront to retrieve data from DynamoDB tables?  
upvoted 4 times

  **thanhnv142** 1 year, 6 months ago

D is correct. There is construction of how to retrieve data from DynamoDB with Athena  
<https://docs.aws.amazon.com/athena/latest/ug/connect-to-a-data-source.html>  
upvoted 2 times

A company wants to analyze and troubleshoot Access Denied errors and Unauthorized errors that are related to IAM permissions. The company has AWS CloudTrail turned on.

Which solution will meet these requirements with the LEAST effort?

- A. Use AWS Glue and write custom scripts to query CloudTrail logs for the errors.
- B. Use AWS Batch and write custom scripts to query CloudTrail logs for the errors.
- C. Search CloudTrail logs with Amazon Athena queries to identify the errors.
- D. Search CloudTrail logs with Amazon QuickSight. Create a dashboard to identify the errors.

**Correct Answer: C**

*Community vote distribution*

C (77%)

D (23%)

- Guru4Cloud

Highly Voted

1 year, 2 months ago

Selected Answer: C

Athena allows you to run SQL queries on data in Amazon S3, including CloudTrail logs. It is the easiest way to query the logs and identify specific errors without needing to write any custom code or scripts.

With Athena, you can write simple SQL queries to filter the CloudTrail logs for the "AccessDenied" and "UnauthorizedOperation" error codes. This will return the relevant log entries that you can then analyze.

upvoted 5 times
- awsgeek75

Highly Voted

9 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html>

When troubleshooting you will want to query specific things in the log and Athena provides query language for that.

Quick Sight is data analytics and visualisation tool. You can use it to aggregate data and maybe make a dashboard for number of errors by type etc but that doesn't help you troubleshoot anything.

C is correct

upvoted 5 times
- pentium75

Most Recent

10 months, 1 week ago

Selected Answer: C

"Search CloudTrail logs with Amazon QuickSight", that doesn't work. QuickSight can visualize Athena query results, so "search CloudTrail logs with Amazon Athena, then create a dashboard with Amazon QuickSight" would make sense. But QuickSight without Athena won't work.

upvoted 4 times
- [Removed]

10 months, 4 weeks ago

Selected Answer: C

Athena is for searching

upvoted 3 times
- bogobob

11 months, 3 weeks ago

Selected Answer: D

The question asks specifically to "analyze and troubleshoot". While Athena is easy to get the data, you then just have a list of logs. Not very useful to troubleshoot...

upvoted 1 times
- pentium75

10 months, 1 week ago

But without Athena, there is nothing you can visualize in QuickSight.

upvoted 3 times
- awsgeek75

9 months, 2 weeks ago

How will pretty pictures in QuickSight help with troubleshooting?

upvoted 2 times
- NickGordon

12 months ago

Selected Answer: D

Quick Sight is an analytics tool. Sounds like a LEAST effort option

upvoted 3 times

 **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: C**

C for me. Using Athena with CloudTrail logs is a powerful way to enhance your analysis of AWS service activity. For example, you can use queries to identify trends and further isolate activity by attributes, such as source IP address or user.

<https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html#:~:text=CloudTrail%20Lake%20documentation.-,Using%20Athena,-with%20CloudTrail%20logs>

upvoted 2 times

 **james2033** 1 year, 3 months ago

**Selected Answer: C**

IAM and CloudTrail <https://docs.aws.amazon.com/IAM/latest/UserGuide/cloudtrail-integration.html#stscldtralexample-assumerole> .  
Query CloudTrail logs by Athena <https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html#tips-for-querying-cloudtrail-logs#tips-for-querying-cloudtrail-logs>

upvoted 2 times

 **james2033** 1 year, 3 months ago

Choose C, not D, because need "analyze and troubleshoot", not just see on dashboard (in D).


upvoted 2 times

 **live\_reply\_developers** 1 year, 3 months ago

**Selected Answer: C**

Amazon Athena is an interactive query service provided by AWS that enables you to analyze data , is a little bit more suitable integrated with cloud trail that permit to verify WHO accessed the service.

upvoted 2 times

 **manuh** 1 year, 4 months ago

**Selected Answer: C**

Dashboard isnt requires. Also refer to this <https://repost.aws/knowledge-center/troubleshoot-iam-permission-errors>

upvoted 2 times

 **haoAWS** 1 year, 4 months ago

**Selected Answer: D**

I am struggling for the C and D for a long time, and ask the chatGPT. The chatGPT says D is better, since Athena requires more expertise on SQL.

upvoted 1 times

 **antropaws** 1 year, 4 months ago

**Selected Answer: D**


Both C and D are feasible. I vote for D:

Amazon QuickSight supports logging the following actions as events in CloudTrail log files:

- Whether the request was made with root or AWS Identity and Access Management user credentials
- Whether the request was made with temporary security credentials for an IAM role or federated user
- Whether the request was made by another AWS service

<https://docs.aws.amazon.com/quicksight/latest/user/logging-using-cloudtrail.html>

upvoted 1 times

 **PCWu** 1 year, 4 months ago

**Selected Answer: C**

The Answer will be C:  
Need to use Athena to query keywords and sort out the error logs.  
D: No need to use Amazon QuickSight to create the dashboard.

upvoted 2 times

 **Axeashes** 1 year, 4 months ago

**Selected Answer: C**

"Using Athena with CloudTrail logs is a powerful way to enhance your analysis of AWS service activity."  
<https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html>

upvoted 2 times

 **oras2023** 1 year, 4 months ago

**Selected Answer: C**

Analyse and TROUBLESHOOT, look like Athena

upvoted 1 times

 **oras2023** 1 year, 4 months ago

<https://docs.aws.amazon.com/athena/latest/ug/cloudtrail-logs.html>

upvoted 2 times

 **alexandercamachop** 1 year, 4 months ago

**Selected Answer: D**

It specifies analyze, not query logs.

Which is why option D is the best one as it provides dashboards to analyze the logs.

upvoted 3 times

A company wants to add its existing AWS usage cost to its operation cost dashboard. A solutions architect needs to recommend a solution that will give the company access to its usage cost programmatically. The company must be able to access cost data for the current year and forecast costs for the next 12 months.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Access usage cost-related data by using the AWS Cost Explorer API with pagination.
- B. Access usage cost-related data by using downloadable AWS Cost Explorer report .csv files.
- C. Configure AWS Budgets actions to send usage cost data to the company through FTP.
- D. Create AWS Budgets reports for usage cost data. Send the data to the company through SMTP.

**Correct Answer: A**

*Community vote distribution*

A (100%)

- BrijMohan08**

Highly Voted

1 year, 2 months ago

Selected Answer: A

Keyword  
12 months, API Support  
<https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html>  
upvoted 6 times
- alexandercamachop**

Highly Voted

1 year, 4 months ago

Selected Answer: A

Answer is: A  
says dashboard = Cost Explorer, therefor C & D are eliminated.  
also says programmatically, means non manual intervention therefor API.  
upvoted 5 times
- awsgeek75**

Most Recent

9 months, 4 weeks ago

Selected Answer: A

Programmatically + LEAST overhead = API  
upvoted 4 times
- TariqKipkemei**

11 months, 3 weeks ago

Selected Answer: A

access to its usage cost programmatically = AWS Cost Explorer API  
upvoted 3 times
- thanhnv142**

1 year ago

A: correct  
1. programatically = API  
2. In the next 12 months = cost explorer  
upvoted 4 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: A

Access usage cost-related data by using the AWS Cost Explorer API with pagination  
upvoted 3 times
- wendaz**

1 year ago

don't repeat the answer, it is useless... explain , okay? i have seen your replies many time just to copy the options.. it makes no sense...  
upvoted 9 times
- james2033**

1 year, 3 months ago

Selected Answer: A

AWS Cost Explorer API with paginated request: <https://docs.aws.amazon.com/cost-management/latest/userguide/ce-api-best-practices.html#ce-api-best-practices-optimize-costs>  
upvoted 3 times
- MrAWSAssociate**

1 year, 4 months ago

Selected Answer: A

From AWS Documentation\*:

"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."

\* Source:

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

<https://docs.aws.amazon.com/AWSJavaScriptSDK/v3/latest/clients/client-cost-explorer/interfaces/costexplorerpaginationconfiguration.html>

upvoted 5 times

  **oras2023** 1 year, 4 months ago

**Selected Answer: A**

least operational overhead = API access

upvoted 4 times

  **oras2023** 1 year, 4 months ago

least operational overhead = API access

upvoted 2 times

A solutions architect is reviewing the resilience of an application. The solutions architect notices that a database administrator recently failed over the application's Amazon Aurora PostgreSQL database writer instance as part of a scaling exercise. The failover resulted in 3 minutes of downtime for the application.

Which solution will reduce the downtime for scaling exercises with the LEAST operational overhead?

- A. Create more Aurora PostgreSQL read replicas in the cluster to handle the load during failover.
- B. Set up a secondary Aurora PostgreSQL cluster in the same AWS Region. During failover, update the application to use the secondary cluster's writer endpoint.
- C. Create an Amazon ElastiCache for Memcached cluster to handle the load during failover.
- D. Set up an Amazon RDS proxy for the database. Update the application to use the proxy endpoint.

**Correct Answer: D**

*Community vote distribution*

D (76%)

B (20%)

4%



  alexandercamachop 

Highly Voted 



 1 year, 11 months ago


**Selected Answer: D**

D is the correct answer.  
It is talking about the write database. Not reader.  
Amazon RDS proxy allows you to automatically route write request to the healthy writer, minimizing downtime.  
upvoted 12 times

  nilandd44gg 1 year, 9 months ago

One of the benefits of Amazon RDS Proxy is that it can improve application recovery time after database failovers. While RDS Proxy supports both MySQL as well as PostgreSQL engines, in this post, we will use a MySQL test workload to demonstrate how RDS Proxy reduces client recovery time after failover by up to 79% for Amazon Aurora MySQL and by up to 32% for Amazon RDS for MySQL.  
<https://aws.amazon.com/blogs/database/improving-application-availability-with-amazon-rds-proxy/>  
<https://aws.amazon.com/rds/proxy/faqs/>  
upvoted 6 times



  bora4motion 

Most Recent 



 2 days, 22 hours ago

**Selected Answer: D**

I'm going with D. It will reduce the failover time between DBs.  
upvoted 1 times



  ugaga 10 months, 2 weeks ago

AWS EXAM is also an AWS promotion...  
RDS Proxy really does not reduce recovery time if you don't have alternative RDS instance. You will see if you understand what proxy is. But this is exam (= promotion) of AWS, D could be an answer, because some AWS document says so.  
upvoted 2 times

  pentium75 1 year, 4 months ago



**Selected Answer: D**

"RDS Proxy reduces client recovery time after failover by up to 79% for Amazon Aurora MySQL "  
<https://aws.amazon.com/de/blogs/database/improving-application-availability-with-amazon-rds-proxy/>  
upvoted 3 times



  ftaws 1 year, 4 months ago

**Selected Answer: B**

RDS Proxy is used for DB timeout not downtime.  
How to reduce downtime with RDS Proxy?  
There is no change downtime if we use RDS Proxy.  
upvoted 3 times

  pentium75 1 year, 4 months ago

"How to reduce downtime with RDS Proxy", by eliminating the need for the application to retrieve the new DNS record after the old one times out.  
upvoted 4 times

  MatAlves 7 months, 3 weeks ago

I would never have thought about DNS records in this question. Nice catch!  
upvoted 2 times

  **pentium75** 1 year, 4 months ago

"RDS Proxy reduces client recovery time after failover by up to 79% for Amazon Aurora MySQL "  
<https://aws.amazon.com/de/blogs/database/improving-application-availability-with-amazon-rds-proxy/>  
upvoted 3 times

  **Cyberkayu** 1 year, 4 months ago

**Selected Answer: B**

they are using Aurora, RDS proxy dont work here  
Answer B  
upvoted 2 times

  **pentium75** 1 year, 4 months ago

Wrong: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>  
upvoted 2 times

  **pentium75** 1 year, 4 months ago

"RDS Proxy reduces client recovery time after failover by up to 79% for Amazon Aurora MySQL "  
<https://aws.amazon.com/de/blogs/database/improving-application-availability-with-amazon-rds-proxy/>  
upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

D. Set up an Amazon RDS proxy for the database. Update the application to use the proxy endpoint.  
upvoted 3 times

  **hachiri** 1 year, 8 months ago

point is Aurora Multi-Master  
Set up a secondary Aurora PostgreSQL cluster in the \*same\* AWS Region  
upvoted 2 times

  **hachiri** 1 year, 8 months ago

I mean correct is B  
upvoted 1 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: C**

Availability is the main requirement here. Even if RDS proxy is used, it will still find the writer instance unavailable during the scaling exercise.  
Best option is to create an Amazon ElastiCache for Memcached cluster to handle the load during the scaling operation.  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

"RDS Proxy reduces client recovery time after failover by up to 79% for Amazon Aurora MySQL "  
<https://aws.amazon.com/de/blogs/database/improving-application-availability-with-amazon-rds-proxy/>  
upvoted 2 times

  **wizcloudifa** 1 year ago

that will only account for the read operation part of it, we are concerned with the write operations here  
upvoted 2 times

  **pentium75** 1 year, 4 months ago

Failover is faster with RDS proxy  
upvoted 2 times

  **AshishRocks** 1 year, 11 months ago

Set up an Amazon RDS proxy for the database. Update the application to use the proxy endpoint.  
D is the answer  
upvoted 4 times

A company has a regional subscription-based streaming service that runs in a single AWS Region. The architecture consists of web servers and application servers on Amazon EC2 instances. The EC2 instances are in Auto Scaling groups behind Elastic Load Balancers. The architecture includes an Amazon Aurora global database cluster that extends across multiple Availability Zones.

The company wants to expand globally and to ensure that its application has minimal downtime.

Which solution will provide the MOST fault tolerance?



- A. Extend the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region. Use an Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- B. Deploy the web tier and the application tier to a second Region. Add an Aurora PostgreSQL cross-Region Aurora Replica in the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region. Promote the secondary to primary as needed.
- C. Deploy the web tier and the application tier to a second Region. Create an Aurora PostgreSQL database in the second Region. Use AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- D. Deploy the web tier and the application tier to a second Region. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region. Promote the secondary to primary as needed.


**Correct Answer: D**

*Community vote distribution*

D (95%)

2%



-   **TariqKipkemei**


Highly Voted 

 1 year, 9 months ago

Selected Answer: D

Auto Scaling groups can span Availability Zones, but not AWS regions.  
Hence the best option is to deploy the web tier and the application tier to a second Region. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region. Promote the secondary to primary as needed.



upvoted 21 times
-   **awsgeek75**

Most Recent 

 1 year, 3 months ago



Selected Answer: D

A: Not possible for autoscaling across regions  
BC: Using PostgreSQL, not sure why?  
D: MOST fault tolerant != MOST scalable. This gives least downtime.

upvoted 4 times
-   **potomac** 1 year, 6 months ago

Selected Answer: D



EC2 Auto Scaling groups are regional constructs. They can span Availability Zones, but not AWS regions

upvoted 3 times
-   **thanhnv142** 1 year, 6 months ago



527:

D is correct:

  - B & C is not correct because it mentions Aurora PostgreSQL which is not mentioned in the question
  - A is not correct because Auto scaling group can not span regions

upvoted 4 times
-   **wsdasdasdqwdaw** 1 year, 6 months ago

Simple as that.

upvoted 2 times
-   **Guru4Cloud** 1 year, 8 months ago

Selected Answer: D

Using an Aurora global database that spans both the primary and secondary regions provides automatic replication and failover capabilities for the database tier.

Deploying the web and application tiers to a second region provides fault tolerance for those components.

Using Route53 health checks and failover routing will route traffic to the secondary region if the primary region becomes unavailable.



This provides fault tolerance across all tiers of the architecture while minimizing downtime. Promoting the secondary database to primary ensures the second region can continue operating if needed.

A is close, but doesn't provide an automatic database failover capability.

B and C provide database replication, but not automatic failover.



So D is the most comprehensive and fault tolerant architecture.

upvoted 4 times

  **JA2018** 5 months, 1 week ago

PostgreSQL in options b & C could be a distractor.... PostgreSQL not mentioned in stem



upvoted 1 times

  **Zox42** 1 year, 10 months ago

**Selected Answer: D**

Answer D

upvoted 2 times

  **Zuit** 1 year, 10 months ago

**Selected Answer: D**

D seems fitting: Global Databbase and deploying it in the new region



upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: B**



B is correct!

upvoted 1 times

  **manuh** 1 year, 10 months ago

Replicated db doesnt mean they will act as a single db once the transfer is completed. Global db is the correct approach

upvoted 2 times

  **JA2018** 5 months, 1 week ago


PostgreSQL in options b & C could be a distractor.... PostgreSQL not mentioned in stem

upvoted 1 times

  **r3mo** 1 year, 10 months ago

"D" is the answer: because Aws Aurora Global Database allows you to read and write from any region in the global cluster. This enables you to distribute read and write workloads globally, improving performance and reducing latency. Data is replicated synchronously across regions, ensuring strong consistency.

upvoted 4 times

  **Henrytml** 1 year, 10 months ago

**Selected Answer: A**



A is the only answer remain using ELB, both Web/App/DB has been taking care with replicating in 2nd region, lastly route 53 for failover over multiple regions

upvoted 1 times

  **Henrytml** 1 year, 10 months ago

i will revoke my answer to standby web in 2nd region, instead of trigger to scale out

upvoted 1 times

  **manuh** 1 year, 10 months ago

also Asg cant span beyond a region

upvoted 2 times

  **alexandercamachop** 1 year, 11 months ago


**Selected Answer: D**

B&C are discarted.

The answer is between A and D.

I would go with D because it explicitley created this web / app tier in second region, instead A just autoscales into a secondary region, rather then always having resources in this second region.

upvoted 4 times

  **JA2018** 5 months, 1 week ago

as mentioned ASG cannot span across different AWS regions

upvoted 1 times

A data analytics company wants to migrate its batch processing system to AWS. The company receives thousands of small data files periodically during the day through FTP. An on-premises batch job processes the data files overnight. However, the batch job takes hours to finish running.

The company wants the AWS solution to process incoming data files as soon as possible with minimal changes to the FTP clients that send the files. The solution must delete the incoming data files after the files have been processed successfully. Processing for each file needs to take 3-8 minutes.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use an Amazon EC2 instance that runs an FTP server to store incoming files as objects in Amazon S3 Glacier Flexible Retrieval. Configure a job queue in AWS Batch. Use Amazon EventBridge rules to invoke the job to process the objects nightly from S3 Glacier Flexible Retrieval. Delete the objects after the job has processed the objects.
- B. Use an Amazon EC2 instance that runs an FTP server to store incoming files on an Amazon Elastic Block Store (Amazon EBS) volume. Configure a job queue in AWS Batch. Use Amazon EventBridge rules to invoke the job to process the files nightly from the EBS volume. Delete the files after the job has processed the files.
- C. Use AWS Transfer Family to create an FTP server to store incoming files on an Amazon Elastic Block Store (Amazon EBS) volume. Configure a job queue in AWS Batch. Use an Amazon S3 event notification when each file arrives to invoke the job in AWS Batch. Delete the files after the job has processed the files.
- D. Use AWS Transfer Family to create an FTP server to store incoming files in Amazon S3 Standard. Create an AWS Lambda function to process the files and to delete the files after they are processed. Use an S3 event notification to invoke the Lambda function when the files arrive.

Correct Answer: D

Community vote distribution

D (96%)

4%

  **pentium75** Highly Voted 10 months, 1 week ago

Selected Answer: D

Obviously we choose AWS Transfer Family over hosting the FTP server ourselves on an EC2 instance. And "process incoming data files as soon as possible" -> trigger Lambda when files arrive. Lambda functions can run up to 15 minutes, it takes "3-8 minutes" per file -> works.

AWS Batch just schedules jobs, but these still need to run somewhere (Lambda, Fargate, EC2).  
upvoted 5 times

  **wizcloudifa** 6 months ago

well if you opt for unmanaged compute option in AWS Batch in that case you dont have to worry about the compute, AWS batch automatically provisions it for you  
upvoted 2 times

  **wsdasdasdqwda** Most Recent 1 year ago

FTP => AWS Transfer Family, => C or D, but in C is used EBS not S3 which needs EC2 and in general is more complex => very clear D.  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

Selected Answer: D

The key points:

Use AWS Transfer Family for the FTP server to receive files directly into S3. This avoids managing FTP servers.  
Process each file as soon as it arrives using Lambda triggered by S3 events. Lambda provides fast processing time per file.  
Lambda can also delete files after processing succeeds.  
Options A, B, C involve more operational overhead of managing FTP servers and batch jobs. Processing latency would be higher waiting for batch windows.  
Storing files in Glacier (Option A) adds latency for retrieving files.  
upvoted 2 times

  **hsinchang** 1 year, 3 months ago

Selected Answer: D

Processing for each file needs to take 3-8 minutes clearly indicates Lambda functions.  
upvoted 2 times

🗄️ 👤 **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: D**

Process incoming data files with minimal changes to the FTP clients that send the files = AWS Transfer Family.  
Process incoming data files as soon as possible = S3 event notification.  
Processing for each file needs to take 3-8 minutes = AWS Lambda function.  
Delete file after processing = AWS Lambda function.  
upvoted 4 times

🗄️ 👤 **antropaws** 1 year, 4 months ago

**Selected Answer: D**

Most likely D.  
upvoted 2 times

🗄️ 👤 **r3mo** 1 year, 4 months ago

"D" Since each file takes 3-8 minutes to process the lambda function can process the data file without a problem.  
upvoted 2 times

🗄️ 👤 **maver144** 1 year, 4 months ago

**Selected Answer: D**

You cannot setup AWS Transfer Family to save files into EBS.  
upvoted 3 times

🗄️ 👤 **oras2023** 1 year, 4 months ago

<https://aws.amazon.com/aws-transfer-family/>  
upvoted 2 times

🗄️ 👤 **secdgs** 1 year, 4 months ago

**Selected Answer: D**

D. Because  
1. process immediate when file transfer to S3 not wait for process several file in one time.  
2. takes 3-8 can use Lambda.  
  
C. Wrong because AWS Batch is used for running large-scale or large amount of data in one time.  
upvoted 2 times

🗄️ 👤 **Aymanovitchy** 1 year, 4 months ago

To meet the requirements of processing incoming data files as soon as possible with minimal changes to the FTP clients, and deleting the files after successful processing, the most operationally efficient solution would be:

D. Use AWS Transfer Family to create an FTP server to store incoming files in Amazon S3 Standard. Create an AWS Lambda function to process the files and delete them after processing. Use an S3 event notification to invoke the Lambda function when the files arrive.  
upvoted 2 times

🗄️ 👤 **bajwa360** 1 year, 4 months ago

**Selected Answer: D**

It should be D as lambda is more operationally viable solution given the fact each processing takes 3-8 minutes that lambda can handle  
upvoted 2 times

🗄️ 👤 **alexandercamachop** 1 year, 4 months ago

**Selected Answer: C**

Answer has to be between C or D.  
Because Transfer Family is obvious due to FTP.  
Now I would go with C because it uses AWS Batch, which makes more sense for Batch processing rather than AWS Lambda.  
upvoted 1 times

🗄️ 👤 **pentium75** 10 months, 1 week ago

Also Batch just triggers jobs, they still need to run somewhere (like in Lambda).  
upvoted 2 times

🗄️ 👤 **pentium75** 10 months, 1 week ago

Why? "Process incoming data files as soon as possible", by triggering the Lambda function when files arrive. Batch is for scheduled jobs.  
upvoted 2 times

🗄️ 👤 **Bill1000** 1 year, 4 months ago

I am between C and D. My reason is:

"The company wants the AWS solution to process incoming data files **as soon as possible** with minimal changes to the FTP clients that send the files."  
upvoted 3 times

A company is migrating its workloads to AWS. The company has transactional and sensitive data in its databases. The company wants to use AWS Cloud solutions to increase security and reduce operational overhead for the databases.

Which solution will meet these requirements?

- A. Migrate the databases to Amazon EC2. Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.
- B. Migrate the databases to Amazon RDS Configure encryption at rest.
- C. Migrate the data to Amazon S3 Use Amazon Macie for data security and protection
- D. Migrate the database to Amazon RDS. Use Amazon CloudWatch Logs for data security and protection.

Correct Answer: B

Community vote distribution

B (100%)

  **AshishRocks** Highly Voted 1 year, 11 months ago

B is the answer  
Why not C - Option C suggests migrating the data to Amazon S3 and using Amazon Macie for data security and protection. While Amazon Macie provides advanced security features for data in S3, it may not be directly applicable or optimized for databases, especially for transactional and sensitive data. Amazon RDS provides a more suitable environment for managing databases.  
upvoted 11 times

  **oras2023** Highly Voted 1 year, 11 months ago

**Selected Answer: B**  
B. Migrate the databases to Amazon RDS Configure encryption at rest.  
Looks like best option  
upvoted 5 times

  **venomblade** Most Recent 5 months, 2 weeks ago



**Selected Answer: B**  
S3 and Macie works too. However, the question mentioned about transactional data. RDS is much better for running transactional workloads. Thus, RDS is the correct answer here.  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: B**  
A: Operational overhead of EC2 and whatever DB is running on it  
C: Macie is not for data security, it's for identifying PII and sensitive data  
D: CloudWatch is for cloud events and does not secure databases  
B: RDS is managed so least operational overhead. Encryption at rest means security  
upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**  
Migrate the databases to Amazon RDS Configure encryption at rest.  
upvoted 4 times

  **wendaz** 1 year, 6 months ago  
down voted.  
upvoted 1 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: B**  
Reduce Ops = Migrate the databases to Amazon RDS Configure encryption at rest  
upvoted 4 times

  **alexandercamachop** 1 year, 11 months ago

**Selected Answer: B**  
B for sure.  
First the correct is Amazon RDS, then encryption at rest makes the database secure.  
upvoted 4 times

A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities. The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions. The company needs to improve application performance and decrease latency for the online game in preparation for user growth.

Which solution will meet these requirements?

- A. Add an Amazon CloudFront distribution in front of the NLBs. Increase the Cache-Control max-age parameter.
- B. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.
- C. Add AWS Global Accelerator in front of the NLBs. Configure a Global Accelerator endpoint to use the correct listener ports.
- D. Add an Amazon API Gateway endpoint behind the NLBs. Enable API caching. Override method caching for the different stages.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Guru4Cloud** Highly Voted  1 year, 2 months ago

**Selected Answer: C**

The key considerations are:

The application uses TCP and UDP for multiplayer gaming, so Network Load Balancers (NLBs) are appropriate.

AWS Global Accelerator can be added in front of the NLBs to improve performance and reduce latency by intelligently routing traffic across AWS Regions and Availability Zones.

Global Accelerator provides static anycast IP addresses that act as a fixed entry point to application endpoints in the optimal AWS location. This improves availability and reduces latency.

The Global Accelerator endpoint can be configured with the correct NLB listener ports for TCP and UDP.

upvoted 6 times

  **TariqKipkemei** Highly Voted  1 year, 3 months ago

**Selected Answer: C**

TCP ,UDP, Gaming = global accelerator and Network Load Balancer

upvoted 5 times

  **awsgeek75** Most Recent  9 months, 4 weeks ago

**Selected Answer: C**

A: CloudFront is for caching. Not required

B: ALB is for HTTP layer, won't help with TCP UDP issues

D: API Gateway, API Caching total rubbish, ignore this option

C: Is correct as Global Accelerator uses unicast for reducing latency globally.



upvoted 2 times

  **Henrytml** 1 year, 4 months ago

**Selected Answer: C**

only b and c handle TCP/UDP, and C comes with accelerator to enhance performance

upvoted 2 times

  **manuh** 1 year, 4 months ago

Does alb handle udp? Can u share a source?

upvoted 1 times


  **alexandercamachop** 1 year, 4 months ago

**Selected Answer: C**

UDP and TCP is AWS Global accelerator as it works in the Transportation layer.

Now this with NLB is perfect.

upvoted 3 times

  **oras2023** 1 year, 4 months ago

**Selected Answer: C**

C is helping to reduce latency for end clients

upvoted 3 times

A company needs to integrate with a third-party data feed. The data feed sends a webhook to notify an external service when new data is ready for consumption. A developer wrote an AWS Lambda function to retrieve data when the company receives a webhook callback. The developer must make the Lambda function available for the third party to call.



Which solution will meet these requirements with the MOST operational efficiency?


- A. Create a function URL for the Lambda function. Provide the Lambda function URL to the third party for the webhook.
- B. Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lambda function. Provide the public hostname of the SQS queue to the third party for the webhook.

**Correct Answer: A**

*Community vote distribution*

A (100%)



-   **TariqKipkemei**


Highly Voted 

 1 year, 9 months ago

Selected Answer: A

A function URL is a dedicated HTTP(S) endpoint for your Lambda function. When you create a function URL, Lambda automatically generates a unique URL endpoint for you.



upvoted 9 times
-   **james2033**


Highly Voted 

 1 year, 9 months ago

Selected Answer: A

Keyword "Lambda function" and "webhook". See <https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-saas-furls.html#create-stripe-cfn-stack>

upvoted 6 times
-   **emakid**

Most Recent 

 10 months, 1 week ago

Selected Answer: A

AWS Lambda can provide a URL to call using Function URLs. This is a relatively new feature in AWS Lambda that allows you to create HTTPS endpoints for your Lambda functions, making it easy to invoke the function directly over the web.



Key Features of Lambda Function URLs:

Direct Access: Provides a simple and direct way to call a Lambda function via an HTTP(S) request.

Easy Configuration: You can create a function URL for a Lambda function using the AWS Management Console, AWS CLI, or AWS SDKs.

Managed Service: AWS manages the infrastructure for you, handling scaling, patching, and maintenance.



Security: You can configure authentication and authorization using AWS IAM or AWS Lambda function URL settings.

upvoted 2 times
-   **awsgeek75** 1 year, 3 months ago



Selected Answer: A

Apart from simplest and most operational, I think A is the only option that will work!

BCD cannot even be implemented in real world imho. Happy to be corrected



upvoted 3 times
-   **Orit** 1 year, 5 months ago

B is the answerThe best solution to make the Lambda function available for the third party to call with the MOST operational efficiency is to deploy an Application Load Balancer (ALB) in front of the Lambda function and provide the ALB URL to the third party for the webhook. This solution is the most efficient because it allows the third party to call the Lambda function without having to worry about managing the Lambda function's availability or scaling. The ALB will automatically distribute traffic across multiple Lambda functions, if necessary, and will also provide redundancy in case of a failure.

upvoted 1 times
-   **hro** 1 year ago

I believe you are correct. Lambda functions as targets - implementing ALBs

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/lambda-functions.html>

upvoted 1 times
-   **Guru4Cloud** 1 year, 8 months ago

Selected Answer: A

The key points:

A Lambda function needs to be invoked by a third party via a webhook.

Using a function URL provides a direct invoke endpoint for the Lambda function. This is simple and efficient.

Options B, C, and D insert unnecessary components like ALB, SNS, SQS between the webhook and the Lambda function. These add complexity without benefit.

A function URL can be generated and provided to the third party quickly without additional infrastructure.

upvoted 5 times

🗲️ 👤 **Abrar2022** 1 year, 10 months ago

**Selected Answer: A**

key word: Lambda function URLs

upvoted 3 times

🗲️ 👤 **maver144** 1 year, 10 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/lambda/latest/dg/lambda-urls.html>

upvoted 2 times

🗲️ 👤 **jkhan2405** 1 year, 10 months ago

**Selected Answer: A**

It's A

upvoted 2 times

🗲️ 👤 **alexandercamachop** 1 year, 11 months ago

**Selected Answer: A**

A would seem like the correct one but not sure.

upvoted 2 times

A company has a workload in an AWS Region. Customers connect to and access the workload by using an Amazon API Gateway REST API. The company uses Amazon Route 53 as its DNS provider. The company wants to provide individual and secure URLs for all customers.

Which combination of steps will meet these requirements with the MOST operational efficiency? (Choose three.)

- A. Register the required domain in a registrar. Create a wildcard custom domain name in a Route 53 hosted zone and record in the zone that points to the API Gateway endpoint.
- B. Request a wildcard certificate that matches the domains in AWS Certificate Manager (ACM) in a different Region.
- C. Create hosted zones for each customer as required in Route 53. Create zone records that point to the API Gateway endpoint.
- D. Request a wildcard certificate that matches the custom domain name in AWS Certificate Manager (ACM) in the same Region.
- E. Create multiple API endpoints for each customer in API Gateway.
- F. Create a custom domain name in API Gateway for the REST API. Import the certificate from AWS Certificate Manager (ACM).

Correct Answer: ADF

Community vote distribution

ADF (100%)

  **AshishRocks** Highly Voted 1 year, 10 months ago

Step A involves registering the required domain in a registrar and creating a wildcard custom domain name in a Route 53 hosted zone. This allows you to map individual and secure URLs for all customers to your API Gateway endpoints.

Step D is to request a wildcard certificate from AWS Certificate Manager (ACM) that matches the custom domain name you created in Step A. This wildcard certificate will cover all subdomains and ensure secure HTTPS communication.

Step F is to create a custom domain name in API Gateway for your REST API. This allows you to associate the custom domain name with your API Gateway endpoints and import the certificate from ACM for secure communication.

upvoted 9 times

  **Guru4Cloud** Highly Voted 1 year, 8 months ago

**Selected Answer: ADF**

The key points:

Using a wildcard domain and certificate avoids managing individual domains/certs per customer. This is more efficient. The domain, hosted zone, and certificate should all be in the same region as the API Gateway REST API for simplicity.

Creating multiple API endpoints per customer (Option E) adds complexity and is not required.

Option B and C add unnecessary complexity by separating domains, certificates, and hosted zones.

upvoted 8 times

  **awsgeek75** Most Recent 1 year, 3 months ago

ADF looks right but not sure why C is wrong:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-api-gateway.html#routing-to-api-gateway-config>

upvoted 3 times

  **LeonSauveterre** 4 months, 4 weeks ago

My guess on why C is not ideal: Creating individual hosted zones for each customer is too much. It definitely increases operational overhead, while using a wildcard domain eliminates this need (which is more ideal).

upvoted 1 times

  **ukivanlamlpi** 1 year, 9 months ago

**Selected Answer: ADF**

<https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-custom-domains.html>

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHZWorkingWith.html>

upvoted 5 times

  **jaydesai8** 1 year, 10 months ago

**Selected Answer: ADF**

ADF - makes sense

upvoted 3 times

  **jkhan2405** 1 year, 10 months ago

**Selected Answer: ADF**

It's ADF

upvoted 3 times

  **MAMADOUG** 1 year, 10 months ago

For me AFD

upvoted 3 times

  **alexandercamachop** 1 year, 11 months ago

**Selected Answer: ADF**

ADF - One to create the custom domain in Route 53 (Amazon DNS)

Second to request wildcard certificate from ADM

Thirds to import the certificate from ACM.

upvoted 3 times

  **AncaZalog** 1 year, 11 months ago

is ADF

upvoted 2 times

A company stores data in Amazon S3. According to regulations, the data must not contain personally identifiable information (PII). The company recently discovered that S3 buckets have some objects that contain PII. The company needs to automatically detect PII in S3 buckets and to notify the company’s security team.

Which solution will meet these requirements?

- A. Use Amazon Macie. Create an Amazon EventBridge rule to filter the SensitiveData event type from Macie findings and to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- B. Use Amazon GuardDuty. Create an Amazon EventBridge rule to filter the CRITICAL event type from GuardDuty findings and to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- C. Use Amazon Macie. Create an Amazon EventBridge rule to filter the SensitiveData:S3Object/Personal event type from Macie findings and to send an Amazon Simple Queue Service (Amazon SQS) notification to the security team.
- D. Use Amazon GuardDuty. Create an Amazon EventBridge rule to filter the CRITICAL event type from GuardDuty findings and to send an Amazon Simple Queue Service (Amazon SQS) notification to the security team.

**Correct Answer: A**

*Community vote distribution*

A (87%)

13%

- alexandercamachop

Highly Voted

1 year, 11 months ago

Selected Answer: A

B and D are discarded as Macie is to identify PII.  
Now that we have between A and C.  
SNS is more suitable for this option as a pub/sub service, we subscribe the security team and then they will receive the notifications.

upvoted 15 times
- LeonSauveterre

Most Recent

4 months, 4 weeks ago

Selected Answer: A

Macie + SNS = Identify PII + Notification

upvoted 1 times
- awsgeek75

1 year, 3 months ago

Selected Answer: A

BD: Wrong products  
AC: Uses Macie which is the right product but C uses SQS to notify security team which is an incomplete solution (what's listening to SQS?)

upvoted 2 times
- pentium75

1 year, 4 months ago

Selected Answer: A

Detect PII -> Macie, A or C  
Notify security team -> SNS, A or B

upvoted 4 times
- potomac

1 year, 5 months ago

Selected Answer: A

C is SQS, not SNS

upvoted 4 times
- Wayne23Fang

1 year, 8 months ago

SQS mentioned in C.

upvoted 2 times
- Ale1973

1 year, 8 months ago

Selected Answer: A

Amazon SQS is typically used for decoupling and managing messages between distributed application components. It's not typically used for sending notifications directly to humans. On my opinion C isn't a best practice

upvoted 2 times
- Kp88

1 year, 9 months ago

Those who say C , please read carefully (I made the same mistake lol). Teams can't be notified with SQS hence A.

upvoted 3 times

  **ukivanlamlpi** 1 year, 9 months ago

**Selected Answer: C**

there are different type of sensitive data: <https://docs.aws.amazon.com/macie/latest/user/findings-types.html>. if the question only focus on PII, then C is the answer. however, in reality, you will use A, because you will not want bank card, credential...etc all sensitive data , not only PII  
upvoted 3 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: A**

Automatically detect PII in S3 buckets = Amazon Macie  
Notify security team = Amazon SNS  
Trigger notification based on SensitiveData event type from Macie findings = EventBridge  
upvoted 2 times

  **NASHDBA** 1 year, 10 months ago

**Selected Answer: C**

There are different types of Sensitive Data. Here we are only referring to PII. Hence SensitiveData:S3Object/Personal. to use SNS, the security team must subscribe. SQS sends the information as designed  
upvoted 1 times

  **narddrer** 1 year, 10 months ago



**Selected Answer: C**

SensitiveData:S3Object/Personal  
upvoted 1 times

  **jaydesai8** 1 year, 10 months ago

**Selected Answer: A**

Sensitive = MACIE, and SNS to sent notification to the Security Team  
upvoted 3 times

  **Iragmt** 1 year, 10 months ago



C. Because the question mentioned PII only, there are other Sensitive Data aside from PII.  
reference: <https://docs.aws.amazon.com/macie/latest/user/findings-publish-event-schemas.html> look for Event example for a sensitive data finding  
upvoted 2 times

  **Ale1973** 1 year, 8 months ago

But Amazon SQS is typically used for decoupling and managing messages between distributed application components. It's not typically used for sending notifications directly to humans!  
upvoted 3 times

  **kapit** 1 year, 10 months ago

AAAAAAA  
upvoted 2 times

  **jack79** 1 year, 10 months ago

C <https://docs.aws.amazon.com/macie/latest/user/findings-types.html>  
and notice the ensitiveData:S3Object/Personal  
The object contains personally identifiable information (such as mailing addresses or driver's license identification numbers), personal health information (such as health insurance or medical identification numbers), or a combination of the two.  
upvoted 3 times

  **Ale1973** 1 year, 8 months ago

But Amazon SQS is typically used for decoupling and managing messages between distributed application components. It's not typically used for sending notifications directly to humans!  
upvoted 2 times

  **MAMADOUG** 1 year, 10 months ago

I vote for A, Sensitive = MACIE, and SNS to prevent Security Team  
upvoted 4 times

A company wants to build a logging solution for its multiple AWS accounts. The company currently stores the logs from all accounts in a centralized account. The company has created an Amazon S3 bucket in the centralized account to store the VPC flow logs and AWS CloudTrail logs. All logs must be highly available for 30 days for frequent analysis, retained for an additional 60 days for backup purposes, and deleted 90 days after creation.

Which solution will meet these requirements MOST cost-effectively?

- A. Transition objects to the S3 Standard storage class 30 days after creation. Write an expiration action that directs Amazon S3 to delete objects after 90 days.
- B. Transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class 30 days after creation. Move all objects to the S3 Glacier Flexible Retrieval storage class after 90 days. Write an expiration action that directs Amazon S3 to delete objects after 90 days.
- C. Transition objects to the S3 Glacier Flexible Retrieval storage class 30 days after creation. Write an expiration action that directs Amazon S3 to delete objects after 90 days.
- D. Transition objects to the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class 30 days after creation. Move all objects to the S3 Glacier Flexible Retrieval storage class after 90 days. Write an expiration action that directs Amazon S3 to delete objects after 90 days.

**Correct Answer:** C

*Community vote distribution*

C (69%)	A (25%)	6%
---------	---------	----

  alexandercamachop Highly Voted 1 year, 11 months ago

**Selected Answer: C**  
C seems the most sutiable.  
Is the lowest cost.  
After 30 days is backup only, doesn't specify frequent access.  
Therefor we must transition the items after 30 days to Glacier Flexible Retrieval.

Also it says deletion after 90 days, so all answers specifying a transition after 90 days makes no sense.  
upvoted 18 times

  MAMADOUG 1 year, 10 months ago

Agree with you  
upvoted 3 times

  deechean Highly Voted 1 year, 8 months ago

**Selected Answer: A**  
The Glacier min storage duration is 90 days. All the options using Glacier are wrong. Only A is feasible.  
upvoted 10 times

  daniel33 1 year, 7 months ago

S3 Standard is priced at \$0.023 per GB for the first 50 TB stored per month  
S3 Glacier Flexible Retrieval costs \$0.0036 per GB stored per month  
If you move or delete data in Glacier within 90-days since their creation, you will pay an additional charge, that is called an early deletion fee. In US East you will pay \$0.004/GB if you have deleted 1 GB in 2 months, \$0.008/GB if you have deleted 1 GB in 1 month and \$0.012 if you have deleted 1 GB within 3 months.  
  
Even with the early deletion fee, it appears to me that answer 'A' would still be cheaper.  
upvoted 3 times

  awsgeek75 1 year, 3 months ago

But the objects are deleted after 90 days so how is this charge applicable?  
upvoted 2 times

  pentium75 1 year, 4 months ago

But why 'transition to the S3 Standard storage class', aren't they there already by default?  
upvoted 4 times

  LeonSauveterre Most Recent 4 months, 4 weeks ago

**Selected Answer: C**  
A: When logs are created, they are initially automatically created to S3 Standard. What are you transitioning anyway?  
B & D: You could've just deleted logs after 90 days. No need to move them somewhere else then delete them. That's just money down the drain.  
upvoted 1 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago

**Selected Answer: C**

C: Lowest cost  
upvoted 2 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago

A: Standard storage is default so this is wrong.  
B: Looks wrong because it moves object to S3GFR after 90 days when they could just be deleted so extra cost  
D: Same problem as B  
upvoted 2 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: C**

Not A: Objects are created in S3 Standard, so it doesn't make sense to 'transition' them there "30 days after creation"  
Not B or C: No need to "move all objects to the S3 Glacier Flexible Retrieval storage class after 90 days" because we want to delete, not archive, them. Even if we would delete them right after moving, we would pay 90 days minimum storage duration. Plus, we are using "Infrequent Access" classes here, but we have no access at all.  
upvoted 3 times

🗨️ 👤 **Rhydian25** 10 months, 2 weeks ago

I guess you wanted to write "Not B or D"  
upvoted 3 times

🗨️ 👤 **ftaws** 1 year, 4 months ago

**Selected Answer: A**

requirement : frequently analysis  
search cost : S3 STD 0.0004 vs IA 0.001  
so IA is more expensive than STD(A)  
upvoted 1 times

🗨️ 👤 **EdenWang** 1 year, 5 months ago

**Selected Answer: C**

C is most cost-effective  
upvoted 3 times

🗨️ 👤 **Hades2231** 1 year, 8 months ago

**Selected Answer: C**

Things to note are: 30 days frequent access and 90 days after creation, so you only need to do 2 things, not 3. Objects in S3 will be stored by default for 30 days before you can move it to somewhere else, so C is the answer.  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html>  
upvoted 5 times

🗨️ 👤 **rjbihari** 1 year, 8 months ago

C is the correct one .  
As after 30 days it doesn't says about access / retrieval , only backup so move items after 30 days to Glacier Flexible Retrieval.  
And after it says deletion , so expiration action will ensure that the objects are deleted after 90 days, even if they are not accessed  
upvoted 2 times

🗨️ 👤 **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

I think - it is B  
The first 30 days, the logs need to be highly available for frequent analysis. The S3 Standard storage class is the most expensive storage class, but it also provides the highest availability.  
After 30 days, the logs still need to be retained for backup purposes, but they do not need to be accessed frequently. The S3 Standard-IA storage class is a good option for this, as it is less expensive than the S3 Standard storage class.  
After 90 days, the logs can be moved to the S3 Glacier Flexible Retrieval storage class. This is the most cost-effective storage class for long-term archiving.  
The expiration action will ensure that the objects are deleted after 90 days, even if they are not accessed  
upvoted 2 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

"After 90 days, the logs can be moved to the S3 Glacier Flexible Retrieval storage class. This is the most cost-effective storage class for long-term archiving." yeah but we don't need long-term archiving, we want to delete them after 90 days.  
upvoted 3 times

🗨️ 👤 **TariqKipkemei** 1 year, 9 months ago



**Selected Answer: C**

C is the most cost effective solution.  
upvoted 2 times

🗨️ 👤 **antropaws** 1 year, 10 months ago



**Selected Answer: C**

C most likely.  
upvoted 2 times

  **y0eri** 1 year, 10 months ago

**Selected Answer: A**

Question says "All logs must be highly available for 30 days for frequent analysis" I think the answer is A. Glacier is not made for frequent access.  
upvoted 2 times

  **y0eri** 1 year, 10 months ago

I take that back. Moderator, please delete my comment.  
upvoted 4 times

  **KMohsoe** 1 year, 10 months ago

**Selected Answer: B**

I think B  
upvoted 1 times

A company is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store.

Which solution will meet these requirements?

- A. Create a new AWS Key Management Service (AWS KMS) key. Use AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS.
- B. Create a new AWS Key Management Service (AWS KMS) key. Enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster.
- C. Create the Amazon EKS cluster with default options. Use the Amazon Elastic Block Store (Amazon EBS) Container Storage Interface (CSI) driver as an add-on.
- D. Create a new AWS Key Management Service (AWS KMS) key with the alias/aws/ebs alias. Enable default Amazon Elastic Block Store (Amazon EBS) volume encryption for the account.

**Correct Answer: B**

Community vote distribution

B (96%)

4%

  **Guru4Cloud** Highly Voted 8 months, 2 weeks ago

**Selected Answer: B**

B is the correct solution to meet the requirement of encrypting secrets in the etcd store for an Amazon EKS cluster.

The key points:

Create a new KMS key to use for encryption.  
Enable EKS secrets encryption using that KMS key on the EKS cluster. This will encrypt secrets in the Kubernetes etcd store.  
Option A uses Secrets Manager which does not encrypt the etcd store.  
Option C uses EBS CSI which is unrelated to etcd encryption.  
Option D enables EBS encryption but does not address etcd encryption.  
upvoted 7 times

  **TariqKipkemei** Highly Voted 9 months, 2 weeks ago

**Selected Answer: B**

EKS supports using AWS KMS keys to provide envelope encryption of Kubernetes secrets stored in EKS. Envelope encryption adds an additional, customer-managed layer of encryption for application secrets or user data that is stored within a Kubernetes cluster.  
  
<https://eksctl.io/usage/kms-encryption/>  
upvoted 6 times

  **FlyingHawk** Most Recent 3 months, 1 week ago



**Selected Answer: B**

The problem is to ensure all secrets in Amazon EKS be encrypted, so the correct solution is to enable the secret encryption on the EKS cluster with a new AWS Key:  
eksctl utils enable-secrets-encryption \ --cluster my-cluster \ --key-arn arn:aws:kms:region-code:account:key/key  
[https://docs.aws.amazon.com/en\\_en/eks/latest/userguide/enable-kms.html](https://docs.aws.amazon.com/en_en/eks/latest/userguide/enable-kms.html)  
upvoted 1 times

  **LeonSauveterre** 4 months, 4 weeks ago



**Selected Answer: B**

"must be encrypted in the Kubernetes etcd key-value store" - that leaves B only.  
upvoted 1 times

  **manuh** 10 months, 1 week ago

**Selected Answer: A**

Why not a  
upvoted 1 times

  **TariqKipkemei** 9 months, 2 weeks ago  
option A does not enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster  
upvoted 2 times

  **MrAWSAssociate** 10 months, 2 weeks ago

**Selected Answer: B**

B is the right option.

<https://docs.aws.amazon.com/eks/latest/userguide/enable-kms.html>

upvoted 5 times

  **alexandercamachop** 11 months ago

**Selected Answer: B**

It is B, because we need to encrypt inside of the EKS cluster, not outside.  
AWS KMS is to encrypt at rest.

upvoted 5 times

  **AncaZalog** 11 months ago

is B, not D

upvoted 2 times

A company wants to provide data scientists with near real-time read-only access to the company's production Amazon RDS for PostgreSQL database. The database is currently configured as a Single-AZ database. The data scientists use complex queries that will not affect the production database. The company needs a solution that is highly available.

Which solution will meet these requirements MOST cost-effectively?

- A. Scale the existing production database in a maintenance window to provide enough power for the data scientists.
- B. Change the setup from a Single-AZ to a Multi-AZ instance deployment with a larger secondary standby instance. Provide the data scientists access to the secondary instance.
- C. Change the setup from a Single-AZ to a Multi-AZ instance deployment. Provide two additional read replicas for the data scientists.
- D. Change the setup from a Single-AZ to a Multi-AZ cluster deployment with two readable standby instances. Provide read endpoints to the data scientists.

**Correct Answer: D**

*Community vote distribution*

D (83%)

Other

  **NASHDBA** Highly Voted 1 year, 10 months ago

**Selected Answer: D**

Highly Available = Multi-AZ Cluster

Read-only + Near Real time = readable standby.

Read replicas are async whereas readable standby is synchronous.



<https://stackoverflow.com/questions/70663036/differences-b-w-aws-read-replica-and-the-standby-instances>

upvoted 23 times

  **chickenmf** 1 year, 1 month ago

a Multi-AZ instance deployment is also highly available for a lower cost, no?

upvoted 2 times

  **Smart** 1 year, 8 months ago


This^ is the reason.

upvoted 3 times

  **maver144** Highly Voted 1 year, 10 months ago

It's either C or D. To be honest, I find the newest questions to be ridiculously hard (roughly 500+). I agree with @alexandercamachop that Multi Az in Instance mode is cheaper than Cluster. However, with Cluster we have reader endpoint available to use out-of-box, so there is no need to provide read-replicas, which also has its own costs. The ridiculous part is that I'm pretty sure even the AWS support would have troubles to answer which configuration is MOST cost-effective.

upvoted 13 times

  **maver144** 1 year, 10 months ago

Near real-time is clue for C, since read replicas are async, but still its not obvious question.

upvoted 2 times

  **manuh** 1 year, 10 months ago

Absolutely true that 500+ questions are damn difficult to answer. I still dont know why is B incorrect. Shouldn't 1 extra be better than 2 ?

upvoted 1 times

  **cyber\_bedouin** 1 year, 4 months ago

they are not all hard, most are normal. its just this one and that one about EKS encryption control plane (earlier than this page).

upvoted 3 times

  **AdamVigas** Most Recent 2 months, 2 weeks ago

**Selected Answer: C**

two readdable standby instances...if they arre standby they are not supposed to read...am i getting something wrong ? i would go for C

upvoted 1 times



  **emakid** 10 months, 1 week ago

**Selected Answer: C**

Option D: Multi-AZ cluster deployment with two readable standby instances would be more costly and is not necessary if read replicas are sufficient for the data scientists' needs.

Thus, Option C is the most cost-effective and operationally efficient solution to meet the company's requirements.

upvoted 2 times

  **osmk** 1 year, 2 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/multi-az-db-clusters-concepts.html>

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: D**

Not A - Not "highly available"

Not B - "Access to the secondary instance" is not possible in Multi-AZ

Not C - Multi-AZ + two (!) read replicas is more expensive than cluster

D - Provides "readable standby instances"

upvoted 3 times

  **SHAAHIBHUSHANAWS** 1 year, 5 months ago

D

<https://aws.amazon.com/about-aws/whats-new/2023/01/amazon-rds-multi-az-readable-standbys-rds-postgresql-inbound-replication/>

upvoted 3 times

  **bogobob** 1 year, 5 months ago

**Selected Answer: D**

<https://aws.amazon.com/blogs/database/choose-the-right-amazon-rds-deployment-option-single-az-instance-multi-az-instance-or-multi-az-database-cluster/>

C would mean you are paying for 4 instances (primary, backup, and 2 read instances). D would be 3 (primary, and 2 backup). Difficult to be sure, pricing calculator doesn't even include clusters yet.

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: D**

Option D is the most cost-effective solution that meets the requirements for this scenario.

The key considerations are:

Data scientists need read-only access to near real-time production data without affecting performance.

High availability is required.

Cost should be minimized.

upvoted 2 times


  **ukivanlamlpi** 1 year, 9 months ago

**Selected Answer: D**

<https://aws.amazon.com/blogs/database/choose-the-right-amazon-rds-deployment-option-single-az-instance-multi-az-instance-or-multi-az-database-cluster/>

only multi AZ cluster have reader endpoint. multi AZ instance secondary replicate is not allow to access

upvoted 2 times

  **msdnpro** 1 year, 9 months ago

**Selected Answer: D**

Support for D:

Amazon RDS now offers Multi-AZ deployments with readable standby instances (also called Multi-AZ DB cluster deployments) in preview. You should consider using Multi-AZ DB cluster deployments with two readable DB instances if you need additional read capacity in your Amazon RDS Multi-AZ deployment and if your application workload has strict transaction latency requirements such as single-digit milliseconds transactions.

<https://aws.amazon.com/blogs/database/readable-standby-instances-in-amazon-rds-multi-az-deployments-a-new-high-availability-option/>



upvoted 2 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: D**

Unlike Multi-AZ instance deployment, where the secondary instance can't be accessed for read or writes, Multi-AZ DB cluster deployment consists of primary instance running in one AZ serving read-write traffic and two other standby running in two different AZs serving read traffic.

upvoted 4 times

  **Iragmt** 1 year, 9 months ago



**Selected Answer: D**

D. using Multi-AZ DB cluster deployments with two readable DB instances if you need additional read capacity in your Amazon RDS Multi-AZ deployment and if your application workload has strict transaction latency requirements such as single-digit milliseconds transactions.

<https://aws.amazon.com/blogs/database/readable-standby-instances-in-amazon-rds-multi-az-deployments-a-new-high-availability-option/>

while on read replicas, Amazon RDS then uses the asynchronous replication method for the DB engine to update the read replica whenever there is a change to the primary DB instance. [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_ReadRepl.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html)

upvoted 2 times

  **manuh** 1 year, 10 months ago

**Selected Answer: B**

Why not b. Shouldnt it have less number of instances than both c and d?

upvoted 2 times

  **baba365** 1 year, 9 months ago

Complex queries on single db will affect performance of db

upvoted 2 times

  **pentium75** 1 year, 4 months ago

You can't 'access the secondary instance' as suggested by B

upvoted 2 times

  **pentium75** 1 year, 4 months ago

"The Multi-AZ instance is suitable for business/mission critical applications that require high availability with low RTO/RPO and resilience to availability zone outage. However, this high availability option isn't a scaling solution for read-only scenarios. You can't use a standby replica to serve read traffic. To serve read-only traffic, use a Multi-AZ DB cluster or a read replica instead."

upvoted 2 times

  **baba365** 1 year, 9 months ago

Multi-AZ is about twice the price of Single-AZ. For example:

db.t2.micro single - \$0.017/hour

db.t2.micro multi - \$0.034/hour

option C: 1 primary + 1 standby + 2 replica = 4Db

option D: 1 primary + 2 standby = 3Db

D. appears to be most cost effective

upvoted 3 times

  **wsdasdasdqwda** 1 year, 6 months ago

I think the best explanation I've read so far.



upvoted 2 times

  **0628atv** 1 year, 10 months ago

D:

<https://aws.amazon.com/tw/blogs/database/readable-standby-instances-in-amazon-rds-multi-az-deployments-a-new-high-availability-option/>



upvoted 2 times

  **vrevkov** 1 year, 10 months ago

**Selected Answer: D**

Forgot to vote

upvoted 3 times

  **vrevkov** 1 year, 10 months ago

I think it's D.

C: Multi-AZ instance = active + standby + two read replicas = 4 RDS instances

D: Multi-AZ cluster = Active + two standby = 3 RDS instances

Single-AZ and Multi-AZ deployments: Pricing is billed per DB instance-hour consumed from the time a DB instance is launched until it is stopped or deleted.

<https://aws.amazon.com/rds/postgresql/pricing/?pg=pr&loc=3>

In the case of a cluster, you will pay less.

upvoted 3 times

A company runs a three-tier web application in the AWS Cloud that operates across three Availability Zones. The application architecture has an Application Load Balancer, an Amazon EC2 web server that hosts user session states, and a MySQL database that runs on an EC2 instance. The company expects sudden increases in application traffic. The company wants to be able to scale to meet future application capacity demands and to ensure high availability across all three Availability Zones.

Which solution will meet these requirements?

- A. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment. Use Amazon ElastiCache for Redis with high availability to store session data and to cache reads. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
- B. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment. Use Amazon ElastiCache for Memcached with high availability to store session data and to cache reads. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
- C. Migrate the MySQL database to Amazon DynamoDB Use DynamoDB Accelerator (DAX) to cache reads. Store the session data in DynamoDB. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
- D. Migrate the MySQL database to Amazon RDS for MySQL in a single Availability Zone. Use Amazon ElastiCache for Redis with high availability to store session data and to cache reads. Migrate the web server to an Auto Scaling group that is in three Availability Zones.

**Correct Answer: A**

*Community vote distribution*

A (80%)B (20%)

- alexandercamachop

Highly Voted

1 year, 4 months ago

Selected Answer: A

Memcached is best suited for caching data, while Redis is better for storing data that needs to be persisted. If you need to store data that needs to be accessed frequently, such as user profiles, session data, and application settings, then Redis is the better choice

upvoted 16 times
- nonameforyou

1 year, 4 months ago

and for high availability, it's better than memcached

upvoted 2 times
- nonameforyou

1 year, 4 months ago

but does rds multi-az provide the needed scalability?

upvoted 3 times
- wsdasdasdqwdaw

1 year ago

it is multi-az cluster deployment, same as B, so yes, it is providing the needed scalability. Great explanation.

upvoted 2 times
- AncaZalog

Highly Voted

1 year, 4 months ago

is A not B

upvoted 5 times
- zdi561

Most Recent

2 months, 3 weeks ago

Selected Answer: A

AI answer, For managing user sessions on AWS ElastiCache, Redis is generally the better choice compared to Memcached because Redis offers more advanced data structures, built-in persistence features, and better support for complex session data, while Memcached is simpler but only suitable for basic key-value caching with limited data types

upvoted 1 times
- osmk

8 months, 2 weeks ago

Selected Answer: A

Replication: Redis supports creating multiple replicas for read scalability and high availability.<https://aws.amazon.com/elasticache/redis-vs-memcached/>

upvoted 2 times
- awsgeek75

9 months, 2 weeks ago

Selected Answer: A

A because of "Amazon EC2 web server that hosts user session states"

C: RDS to DynamoDB doesn't make total sense

D: Single zone is not HA

Between A and B, A is suitable because of session state and Elasticache with Redis is more HA than option B  
upvoted 2 times

  **mr123dd** 10 months ago

**Selected Answer: A**

B: from what I know, Memcached provide better performance and simplicity but lower availability than redis.  
C: mysql is relational database, dynamodb is nosql  
D: single AZ  
upvoted 2 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: A**

ElastiCache for Redis supports HA, ElastiCache for Memcached does not:  
<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/SelectEngine.html>

C could in theory work, but session data is typically stored in ElastiCache, not in DynamoDB.

D is not HA.  
upvoted 3 times

  **Cyberkayu** 10 months, 2 weeks ago

**Selected Answer: B**

'hosts user session states' in question, thus redis  
upvoted 2 times

  **pentium75** 10 months, 1 week ago

Right, but Redis is A

upvoted 2 times

  **potomac** 12 months ago

**Selected Answer: A**

Redis is a widely adopted in-memory data store for use as a database, cache, message broker, queue, session store, and leaderboard.  
<https://aws.amazon.com/elasticache/redis/>  
upvoted 5 times

  **thanhnv142** 1 year ago

B is correct.

We are left with 2 options: A and B. But it requires that the system be able to scale to meet future application capacity demands. Redis is very good.  
But its drawback is not scalable. Thats why they implement memcached.

upvoted 1 times

  **ErnShm** 1 year, 1 month ago

A

Redis as an in-memory data store with high availability and persistence is a popular choice among application developers to store and manage session data for internet-scale applications. Redis provides the sub-millisecond latency, scale, and resiliency required to manage session data such as user profiles, credentials, session state, and user-specific personalization.

upvoted 2 times

  **Gajendr** 10 months, 1 week ago

Redis provides replication while memcached doesnt.

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

The key reasons why option A is preferable:

RDS Multi-AZ provides high availability for MySQL by synchronously replicating data across AZs. Automatic failover handles AZ outages.  
ElastiCache for Redis is better suited for session data caching than Memcached. Redis offers more advanced data structures and flexibility.  
Auto scaling across 3 AZs provides high availability for the web tier

upvoted 2 times

  **ukivanlamlpi** 1 year, 3 months ago

**Selected Answer: B**

the different between Redis and Memcache is that Memcache support multithread process to handle the increase of application traffic.  
<https://aws.amazon.com/elasticache/redis-vs-memcached/>

upvoted 2 times

  **pentium75** 10 months, 1 week ago

ElastiCache for Memcached says "No" for "High Availability"

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/SelectEngine.html>

upvoted 2 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: B**

This requirement wins for me: "be able to scale to meet future application capacity demands".  
Memcached implements a multi-threaded architecture, it can make use of multiple processing cores. This means that you can handle more operations by scaling up compute capacity.

<https://aws.amazon.com/elasticache/redis-vs-memcached/#:~:text=by%20their%20rank.-,Multithreaded%20architecture,-Since%20Memcached%20is>  
upvoted 1 times

  **plndmns** 1 year, 3 months ago

cache reads is memcached right?  
upvoted 1 times

  **MrAWSAssociate** 1 year, 4 months ago

**Selected Answer: B**

B is correct!  
upvoted 3 times

A global video streaming company uses Amazon CloudFront as a content distribution network (CDN). The company wants to roll out content in a phased manner across multiple countries. The company needs to ensure that viewers who are outside the countries to which the company rolls out content are not able to view the content.



Which solution will meet these requirements?

- A. Add geographic restrictions to the content in CloudFront by using an allow list. Set up a custom error message.
- B. Set up a new URL for restricted content. Authorize access by using a signed URL and cookies. Set up a custom error message.
- C. Encrypt the data for the content that the company distributes. Set up a custom error message.
- D. Create a new URL for restricted content. Set up a time-restricted access policy for signed URLs.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **potomac** Highly Voted  12 months ago

**Selected Answer: A**

The CloudFront geographic restrictions feature lets you control distribution of your content at the country level for all files that you're distributing with a given web distribution.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/georestrictions.html>

upvoted 5 times

  **awsgeek75** Highly Voted  9 months, 4 weeks ago

**Selected Answer: A**

This question asks us to guess Netflix subscription model in 2 mins! lol!

BCD are impractical for geo restrictions as you cannot restrict URL by region and you cannot encrypt by geo region (country etc)

upvoted 5 times

  **Guru4Cloud** Most Recent  1 year, 2 months ago

**Selected Answer: A**

Add geographic restrictions to the content in CloudFront by using an allow list. Set up a custom error message

upvoted 3 times

  **TariqKipkemei** 1 year, 3 months ago

**Selected Answer: A**

Add geographic restrictions to the content in CloudFront by using an allow list. Set up a custom error message.

upvoted 2 times

  **jaydesai8** 1 year, 3 months ago

**Selected Answer: A**

A makes sense - cloudfront has the capabilities of georestriction

upvoted 2 times

  **antropaws** 1 year, 4 months ago

**Selected Answer: A**

Pretty sure it's A.

upvoted 2 times

  **alexandercamachop** 1 year, 4 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/georestrictions.html>

upvoted 5 times

  **AncaZalog** 1 year, 4 months ago

is B not A

upvoted 1 times

  **manuh** 1 year, 4 months ago

Signed url or cookies can be used for the banner country as well?

upvoted 1 times

  **antropaws** 1 year, 4 months ago

Why's that?

upvoted 2 times

  **awsgeek75** 9 months, 2 weeks ago

How is signed URL going to be geo restricted? Anyone with signed url can access the content on that url regardless of their location so B is wrong.

upvoted 2 times

A company wants to use the AWS Cloud to improve its on-premises disaster recovery (DR) configuration. The company's core production business application uses Microsoft SQL Server Standard, which runs on a virtual machine (VM). The application has a recovery point objective (RPO) of 30 seconds or fewer and a recovery time objective (RTO) of 60 minutes. The DR solution needs to minimize costs wherever possible.

Which solution will meet these requirements?

- A. Configure a multi-site active/active setup between the on-premises server and AWS by using Microsoft SQL Server Enterprise with Always On availability groups.
- B. Configure a warm standby Amazon RDS for SQL Server database on AWS. Configure AWS Database Migration Service (AWS DMS) to use change data capture (CDC).
- C. Use AWS Elastic Disaster Recovery configured to replicate disk changes to AWS as a pilot light.
- D. Use third-party backup software to capture backups every night. Store a secondary set of backups in Amazon S3.

Correct Answer: B

Community vote distribution



B (59%)

C (41%)

  **1Alpha1** Highly Voted 1 year, 3 months ago

Selected Answer: B



Backup & Restore (RPO in hours, RTO in 24 hours or less)  
Pilot Light (RPO in minutes, RTO in hours)  
Warm Standby (RPO in seconds, RTO in minutes) \*\*\* Right Answer \*\*\*  
Active-Active (RPO is none or possibly seconds, RTO in seconds)  
upvoted 10 times

  **1Alpha1** 1 year, 3 months ago  
[https://disaster-recovery.workshop.aws/en/intro/disaster-recovery.html#:~:text=Pilot%20Light%20\(RPO%20in%20minutes,that%20includes%20that%20critical%20core.](https://disaster-recovery.workshop.aws/en/intro/disaster-recovery.html#:~:text=Pilot%20Light%20(RPO%20in%20minutes,that%20includes%20that%20critical%20core.)  
upvoted 4 times

  **pentium75** Highly Voted 1 year, 4 months ago

Selected Answer: C

Not A - too expensive and not using AWS services  
Not B - "RDS for SQL Server" does not support everything that "SQL Server Standard which runs on a VM" does; CDC supports even less ([https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Source.SQLServer.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SQLServer.html)). Also it would be more expensive than C.  
Not D - "Every night" would not meet the RPO requirement  
upvoted 7 times

  **awsgeek75** 1 year, 3 months ago  
Thanks I was confused between B and C. This makes perfect sense!  
upvoted 1 times

  **zdi561** Most Recent 2 months, 3 weeks ago

Selected Answer: C

AWS Elastic Disaster Recovery (AWS DRS):  
A service that can replicate entire virtual machines, including your SQL Server instance, to a secondary region for disaster recovery  
upvoted 1 times



  **FlyingHawk** 3 months, 1 week ago



Selected Answer: B

AWS DMS with CDC can replicate changes from the on-premises SQL Server to the RDS instance in near real-time, meeting the 30-second RPO requirement.

A warm standby RDS instance ensures a faster recovery, meeting the 60-minute RTO requirement.

This solution is more cost-effective than upgrading to SQL Server Enterprise and provides a managed, scalable database service.  
upvoted 1 times

  **FlyingHawk** 3 months, 1 week ago  
C - Not optimized for real-time replication, making it challenging to meet the 30-second RPO requirement.  
upvoted 1 times

  **Salilgen** 4 months, 1 week ago

**Selected Answer: B**

IMO answer is B.

AWS Elastic Disaster Recovery is cheaper and may meet the RTO/RPO requirements of the question, but option C states that it is configured as a pilot light (RTO/RPO 10s minute)

upvoted 2 times

  **LeonSauveterre** 4 months, 4 weeks ago

**Selected Answer: C**

RPO: the maximum acceptable amount of \*data loss\* during an outage or disaster.

RTO: the maximum acceptable amount of \*downtime\* after a failure.

Why B (RDS with DMS) might be better:

1. RTO Guarantee: A warm standby database provides faster recovery since the database is preconfigured and ready to accept traffic.
2. Simplicity: For workloads compatible with RDS and DMS, this approach simplifies management.

Why C (Elastic DR) might be better:

1. Workload Compatibility: If the on-premises SQL Server workload uses features that RDS or DMS cannot support, Elastic DR ensures full replication by mirroring the entire VM or disk.
2. Cost: Elastic DR replicates data but does not run a live database, making it much cheaper than a warm standby RDS.


"The DR solution needs to minimize costs wherever possible" - That makes me think C is the answer.

upvoted 2 times

  **FlyingHawk** 3 months, 1 week ago

Elastic DRS is better suited for applications with less stringent RPO requirements. Achieving a 30-second RPO with this solution could be challenging, as it is not optimized for real-time replication.

upvoted 1 times

  **example\_** 9 months, 1 week ago

**Selected Answer: B**

Pilot light (RPO in minutes, RTO in tens of minutes)

Warm standby (RPO in seconds, RTO in minutes)

[https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/rel\\_planning\\_for\\_recovery\\_disaster\\_recovery.html](https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/rel_planning_for_recovery_disaster_recovery.html)

upvoted 2 times

  **MrAliMohsan** 8 months, 2 weeks ago

When cost is a concern, and you wish to achieve a similar RPO and RTO objectives as defined in the warm standby strategy, you could consider cloud native solutions, like AWS Elastic Disaster Recovery, that take the pilot light approach and offer improved RPO and RTO targets.

upvoted 3 times

  **abhiarns** 1 year ago

**Selected Answer: C**

AWS DRS(AWS Elastic Disaster Recovery) enables RPOs of seconds and RTOs of minutes.

upvoted 2 times

  **osmk** 1 year, 2 months ago

**Selected Answer: B**

<https://docs.aws.amazon.com/whitepapers/latest/disaster-recovery-workloads-on-aws/disaster-recovery-options-in-the-cloud.html#warm-standby>

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: C**

A: Not possible

B: With RDS it means your failover will launch a different database engine. This is wrong in general

D: No comments

C: It is a disk based replication so it will be similar DB server and this is the product managed by AWS for the DR of on-prem setups.

<https://aws.amazon.com/blogs/modernizing-with-aws/how-to-set-up-disaster-recovery-for-sql-server-always-on-availability-groups-using-aws-elastic-disaster-recovery/>

upvoted 2 times

  **1rob** 1 year, 5 months ago

**Selected Answer: C**

AWS Elastic Disaster Recovery

If you are considering the pilot light or warm standby strategy for disaster recovery, AWS Elastic Disaster Recovery could provide an alternative approach with improved benefits. Elastic Disaster Recovery can offer an RPO and RTO target similar to warm standby, but maintain the low-cost approach of pilot light

From <[https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/rel\\_planning\\_for\\_recovery\\_disaster\\_recovery.html](https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/rel_planning_for_recovery_disaster_recovery.html)>

upvoted 3 times

  **potomac** 1 year, 5 months ago

**Selected Answer: B**

With the pilot light approach, you replicate your data from one environment to another and provision a copy of your core workload infrastructure, not the fully functional copy of your production environment in a recovery environment.

upvoted 2 times

  **saymolet** 1 year, 4 months ago

<https://aws.amazon.com/blogs/architecture/disaster-recovery-dr-architecture-on-aws-part-iii-pilot-light-and-warm-standby/>

upvoted 3 times

  **pentium75** 1 year, 4 months ago

We have no idea if they are using SQL Server features that require OS customization etc., so we can't assume that the app would run on RDS for SQL Server at all. We need a replica of the VM that SQL Server is currently running on, thus C.

upvoted 2 times

  **thanhnv142** 1 year, 6 months ago

C: Pilot light

- In pilot light, databases are always on, thus minimize RPO (can satisfy the 30s requirement)
- Only apps are turn off. But it can satisfy the 60 minutes requirement
- Warm standby, of course, can satisfy all the RPO and RTO requirements, but it is more expensive than pilot light

upvoted 4 times



  **richguo** 1 year, 7 months ago

**Selected Answer: C**

B(warm standby) is doable, but C (pilot light) is most cost effectively.

<https://aws.amazon.com/tw/blogs/architecture/disaster-recovery-dr-architecture-on-aws-part-iii-pilot-light-and-warm-standby/>

upvoted 3 times

  **LazyTs** 1 year, 8 months ago

**Selected Answer: B**

The company wants to improve... so needs something guaranteed to be better than 60 mins RTO

upvoted 2 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Configure a warm standby Amazon RDS for SQL Server database on AWS. Configure AWS Database Migration Service (AWS DMS) to use change data capture (CDC).

upvoted 3 times

  **Eminenza22** 1 year, 8 months ago

Warm standby is costlier than Pilot Light

upvoted 3 times

  **PantryRaid** 1 year, 8 months ago

**Selected Answer: C**

AWS DRS enables RPOs of seconds and RTOs of minutes. Pilot light is also cheaper than warm standby.

<https://aws.amazon.com/disaster-recovery/>

upvoted 3 times

A company has an on-premises server that uses an Oracle database to process and store customer information. The company wants to use an AWS database service to achieve higher availability and to improve application performance. The company also wants to offload reporting from its primary database system.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions. Point the reporting functions toward a separate DB instance from the primary DB instance.
- B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database. Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica.
- C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database. Direct the reporting functions to use the reader instance in the cluster deployment.
- D. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora database. Direct the reporting functions to the reader instances.

Correct Answer: D

Community vote distribution

D (72%)

C (28%)

  **mrsoa** Highly Voted 1 year, 9 months ago

**Selected Answer: D**

Its D  
Multi-AZ DB clusters aren't available with the following engines:  
RDS for MariaDB  
RDS for Oracle  
RDS for SQL Server

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS\\_Fea\\_Regions\\_DB-eng.Feature.MultiAZDBClusters.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS_Fea_Regions_DB-eng.Feature.MultiAZDBClusters.html)  
upvoted 34 times



  **alexandercamachop** Highly Voted 1 year, 11 months ago

**Selected Answer: C**

C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database. Direct the reporting functions to use the reader instance in the cluster deployment.

A and B discarded.  
The answer is between C and D  
D says use an Amazon RDS to build an Amazon Aurora, makes no sense.  
C is the correct one, high availability in multi az deployment.  
Also point the reporting to the reader replica.

upvoted 13 times

  **1rob** 1 year, 5 months ago

Multi-AZ DB clusters aren't available with the following engines:  
RDS for MariaDB  
RDS for Oracle  
RDS for SQL Server

upvoted 6 times

  **bogobob** 1 year, 5 months ago

using RDS to build Aurora from an Oracle DB <https://aws.amazon.com/tutorials/break-free-from-legacy-databases/migrate-oracle-to-amazon-aurora/>

upvoted 3 times

  **LeonSauveterre** Most Recent 4 months, 4 weeks ago

**Selected Answer: D**

In the MOST operationally efficient way, D is better, and A is too much. HOWEVER, Aurora is not Oracle, so this would require migrating from Oracle to Aurora (potentially requiring significant changes to the application, schema, and queries). The question did not say anything about not wanting to migrate.

B - 100% wrong. Single-AZ deployment does not meet the higher availability requirement.  
C - 100% wrong. There's no way to create an Oracle db using RDS.

upvoted 1 times

🗨️ 👤 **chest\_jd** 6 months ago

Unfortunately ALL ANSWERS ARE NOT CORRECT.

A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions.  
How can I create RDS using DMS ?

B. Single-Zone is not suitable for high available solution.

C. RDS cluster deployment are not available for DBs: RDS for Db2, RDS for MariaDB, RDS for Oracle, RDS for SQL Server  
[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS\\_Fea\\_Regions\\_DB-eng.Feature.MultiAZDBClusters.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS_Fea_Regions_DB-eng.Feature.MultiAZDBClusters.html)

D. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora database.  
Aurora DOES NOT support Oracle DB. Only MySQL and PostgreSQL

upvoted 1 times

🗨️ 👤 **Ravan** 1 year, 1 month ago

**Selected Answer: D**

Multi-AZ (Availability Zone) deployments are not available for the following Amazon RDS database engines:

1. Amazon Aurora with MySQL compatibility
2. Amazon Aurora with PostgreSQL compatibility
3. Amazon RDS for SQL Server Express Edition
4. Amazon RDS for Oracle Standard Edition One
5. Amazon RDS for Oracle Standard Edition
6. Amazon RDS for Oracle SE2 (Standard Edition 2)

For these database engines, Amazon RDS provides high availability using other mechanisms specific to each engine, such as Read Replicas or different standby configurations. However, Multi-AZ deployments, which automatically provision and maintain a synchronous standby replica in a different Availability Zone for failover support, are not supported for these engines.

upvoted 4 times

🗨️ 👤 **noircesar25** 1 year, 2 months ago

this link expalins why the answer is C and confirms that rds for oracle supports multi-AZ  
<https://aws.amazon.com/blogs/aws/multi-az-option-for-amazon-rds-oracle/>

upvoted 2 times

🗨️ 👤 **osmk** 1 year, 2 months ago

**Selected Answer: D**

requiring high availability and performance.<https://aws.amazon.com/rds/aurora/>

upvoted 2 times

🗨️ 👤 **awsgeek75** 1 year, 3 months ago

**Selected Answer: D**

Between C&D, D is correct as C is not possible:

<https://aws.amazon.com/blogs/aws/multi-az-option-for-amazon-rds-oracle/>

upvoted 3 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: D**

Not A - Creating multiple instances and keeping them in sync in DMS is surely not "operationally efficient"

Not B - "replica in the same zone" -> does not provide "higher availability"

Not C - "Multi-AZ cluster" does not support Oracle engine

Thus D. Question does not mention that the app would use Oracle-specific features; we're also not asked to minimize application changes. Ideal solution from AWS point of view is to move from Oracle to Aurora.

upvoted 2 times

🗨️ 👤 **XXXXXXINN** 7 months, 1 week ago

but Aurora doesnt support oracle db though

upvoted 1 times

🗨️ 👤 **aws94** 1 year, 4 months ago

**Selected Answer: C**

i am sure just look here

<https://aws.amazon.com/ar/blogs/aws/amazon-rds-multi-az-db-cluster/>

upvoted 1 times

🗨️ 👤 **aws94** 1 year, 4 months ago

sorry this is the right link:



<https://aws.amazon.com/ar/blogs/aws/multi-az-option-for-amazon-rds-oracle/>

upvoted 2 times

🗨️ 👤 **pentium75** 1 year, 4 months ago

"Multi-AZ cluster (!)" does not support Oracle. Multi-AZ instance would.



upvoted 2 times

  **EEK2k** 1 year, 5 months ago

**Selected Answer: C**

It should be C. Oracle DB is supported in RDS Multi-AZ with one standby for HA. <https://aws.amazon.com/rds/features/multi-az/>. Additionally, a reader instance/replica could be added to RDS Multi-AZ with one standby setup to offload the read requests. Aurora is only supported MySQL and Postgres compatible DB so "D" is out.

upvoted 2 times

  **1rob** 1 year, 5 months ago

<https://aws.amazon.com/rds/features/multi-az/> gives: Amazon RDS Multi-AZ is available for RDS for PostgreSQL, RDS for MySQL, RDS for MariaDB, RDS for SQL Server, RDS for Oracle, and RDS for Db2. Amazon RDS Multi-AZ with two readable standbys is available for RDS for PostgreSQL and RDS for MySQL.  
So no reader instance.

upvoted 1 times

  **potomac** 1 year, 5 months ago

**Selected Answer: D**

Multi-AZ DB clusters are NOT available with the following engines:

RDS for MariaDB

RDS for Oracle

RDS for SQL Server


[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS\\_Fea\\_Regions\\_DB-eng.Feature.MultiAZDBClusters.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS_Fea_Regions_DB-eng.Feature.MultiAZDBClusters.html)

upvoted 2 times

  **danielmakita** 1 year, 6 months ago

It is C. Aurora database doesn't support Oracle.

upvoted 1 times

  **JA2018** 5 months, 1 week ago

as pentium75 had pointed out, the stem does not mention that the app would use Oracle-specific features; we're also not asked to minimize application changes. Ideal solution from AWS point of view is to move from Oracle to Aurora.

upvoted 1 times

  **wsdasdasdqwdaw** 1 year, 6 months ago



You can use Aurora instead of Oracle. There are tutorials how to migrate Oracle to Aurora. On top C is not supported. The is not Multi-AZ DB CLUSTER for Oracle.

upvoted 3 times

  **wsdasdasdqwdaw** 1 year, 6 months ago

It is D

upvoted 1 times

  **thanhnv142** 1 year, 6 months ago

None options seems valid. Not C because it is not supported. But not D as well. RDS is not Aurora. They are two separate services. Additionally, In multi AZ instance deployment, it only provides fault tolerance, not High avai.

upvoted 2 times

  **Nikki013** 1 year, 8 months ago

**Selected Answer: D**

Multi-AZ Cluster does not support Oracle as engine:

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS\\_Fea\\_Regions\\_DB-eng.Feature.MultiAZDBClusters.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RDS_Fea_Regions_DB-eng.Feature.MultiAZDBClusters.html)

upvoted 2 times

  **Bennyboy789** 1 year, 8 months ago

**Selected Answer: D**

D is my choice.

Multi-AZ DB cluster does not support Oracle DB.

upvoted 3 times

  **rjbihari** 1 year, 8 months ago

Option C is correct one .

As there is no option for 'Aurora(Oracle Compatible)'.so this kick out D from race.

upvoted 1 times

A company wants to build a web application on AWS. Client access requests to the website are not predictable and can be idle for a long time. Only customers who have paid a subscription fee can have the ability to sign in and use the web application.



Which combination of steps will meet these requirements MOST cost-effectively? (Choose three.)



- A. Create an AWS Lambda function to retrieve user information from Amazon DynamoDB. Create an Amazon API Gateway endpoint to accept RESTful APIs. Send the API calls to the Lambda function.
- B. Create an Amazon Elastic Container Service (Amazon ECS) service behind an Application Load Balancer to retrieve user information from Amazon RDS. Create an Amazon API Gateway endpoint to accept RESTful APIs. Send the API calls to the Lambda function.
- C. Create an Amazon Cognito user pool to authenticate users.
- D. Create an Amazon Cognito identity pool to authenticate users.
- E. Use AWS Amplify to serve the frontend web content with HTML, CSS, and JS. Use an integrated Amazon CloudFront configuration.
- F. Use Amazon S3 static web hosting with PHP, CSS, and JS. Use Amazon CloudFront to serve the frontend web content.



**Correct Answer: ACE**



*Community vote distribution*



ACE (86%)9%



-   **manOfThePeople** Highly Voted 1 year, 8 months ago



If in doubt between E or F. S3 doesn't support server-side scripts, PHP is a server-side script. The answer is ACE.  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>  
upvoted 21 times
-   **msdnpro** Highly Voted 1 year, 10 months ago

**Selected Answer: ACE**  
Option B (Amazon ECS) is not the best option since the website "can be idle for a long time", so Lambda (Option A) is a more cost-effective choice.  
  
Option D is incorrect because User pools are for authentication (identity verification) while Identity pools are for authorization (access control).  
  
Option F is wrong because S3 web hosting only supports static web files like HTML/CSS, and does not support PHP or JavaScript.  
upvoted 10 times
-   **0628atv** 1 year, 10 months ago

[https://aws.amazon.com/getting-started/projects/build-serverless-web-app-lambda-apigateway-s3-dynamodb-cognito/module-1/?nc1=h\\_ls](https://aws.amazon.com/getting-started/projects/build-serverless-web-app-lambda-apigateway-s3-dynamodb-cognito/module-1/?nc1=h_ls)  
upvoted 3 times
-   **Gape4** Most Recent 10 months, 2 weeks ago

**Selected Answer: ACE**  
I will go for A C E  
upvoted 2 times
-   **awsgeek75** 1 year, 3 months ago

**Selected Answer: ACE**  
A: App may be idle for long time so Lambda is perfect (charge per invocation)  
C: Cognito user pool for user auth  
E: Amplify is low code web dev tool  
  
B: Wrong, too much cost when idle  
D: Identity pool is session management/identification. Does not help with auth.  
F: S3 + PHP doesn't work also no security  
upvoted 3 times
-   **rcptryk** 1 year, 5 months ago

**Selected Answer: ACE**  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>  
S3 doesn't support server-side scripting  
upvoted 2 times
-   **potomac** 1 year, 5 months ago

**Selected Answer: ACE**

User Pool = authentication  
Identity Pool = authorization  
upvoted 7 times

  **thanhnv142** 1 year, 6 months ago

A D F:  
A: for hosting the dynamic content of the app. Pay as execution  
D: for granting temporary privilege access to users who has paid a fee.  
F: for hosting the static content of the app  
upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

There is no static content in this web application so F is wrong. You cannot host PHP on S3 also so it is just wrong.  
upvoted 2 times

  **kwang312** 1 year, 7 months ago

**Selected Answer: ACE**  
ACE is correct answer  
upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: CEF**  
C) Create an Amazon Cognito user pool to authenticate users.  
  
E) Use AWS Amplify to serve the frontend web content with HTML, CSS, and JS. Use an integrated CloudFront configuration.  
  
F) Use Amazon S3 static web hosting with PHP, CSS, and JS. Use Amazon CloudFront to serve the frontend web content.  
upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

There is no static content in this web application so F is wrong. You cannot host PHP on S3 also so it is just wrong.  
upvoted 2 times

  **TariqKipkemei** 1 year, 9 months ago



**Selected Answer: ACE**  
Build a web application = AWS Amplify  
Sign in users = Amazon Cognito user pool  
Traffic can be idle for a long time = AWS Lambda  
  
Amazon S3 does not support server-side scripting such as PHP, JSP, or ASP.NET.  
[https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html?icmpid=docs\\_amazons3\\_console#:~:text=website%20relies%20on-,server%2Dside,-processing%2C%20including%20server](https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html?icmpid=docs_amazons3_console#:~:text=website%20relies%20on-,server%2Dside,-processing%2C%20including%20server)  
Traffic can be idle for a long time = AWS Lambda  
upvoted 2 times

  **james2033** 1 year, 9 months ago



**Selected Answer: ACE**  
Use exclusion method: No need for Container (no need run all time), remove B. PHP cannot run with static Amazon S3, remove F.  
Use selection method: Idle for sometime, choose AWS Lambda, choose A. "Amazon Cognito is an identity platform for web and mobile apps."  
(<https://docs.aws.amazon.com/cognito/latest/developerguide/what-is-amazon-cognito.html> ), choose C. Create an identity pool  
<https://docs.aws.amazon.com/cognito/latest/developerguide/tutorial-create-identity-pool.html> . AWS Amplify <https://aws.amazon.com/amplify/> for build full-stack web-app in hours.  
upvoted 6 times

  **baba365** 1 year, 9 months ago

Ans: ACF  
  
use AWS SDK for PHP/JS with S3  
  
[https://docs.aws.amazon.com/sdk-for-php/v3/developer-guide/php\\_s3\\_code\\_examples.html](https://docs.aws.amazon.com/sdk-for-php/v3/developer-guide/php_s3_code_examples.html)  
upvoted 1 times

  **unbendable** 1 year, 6 months ago

did you actually read the link or just copy the first link from google here? the sdk is intended for usage in a php application. it does not say anything about php support in a s3 bucket  
upvoted 2 times

  **Zox42** 1 year, 9 months ago

**Selected Answer: ACE**  
Answer is ACE  
upvoted 2 times

  **jaydesai8** 1 year, 10 months ago

**Selected Answer: ACE**  
Lambda =serverless  
User Pool = For user authentication



Amplify = hosting web/mobile apps  
upvoted 3 times

  **live\_reply\_developers** 1 year, 10 months ago

**Selected Answer: ACE**

S3 doesn't support PHP as stated in answer F.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>  
upvoted 4 times

  **wRhIH** 1 year, 10 months ago

**Selected Answer: ACE**

I don't think S3 can handle anything dynamic such as PHP. So I go for ACE  
upvoted 2 times

  **antropaws** 1 year, 10 months ago

**Selected Answer: ACF**

ACF no doubt. Check the difference between user pools and identity pools.  
upvoted 2 times

A media company uses an Amazon CloudFront distribution to deliver content over the internet. The company wants only premium customers to have access to the media streams and file content. The company stores all content in an Amazon S3 bucket. The company also delivers content on demand to customers for a specific purpose, such as movie rentals or music downloads.

Which solution will meet these requirements?

- A. Generate and provide S3 signed cookies to premium customers.
- B. Generate and provide CloudFront signed URLs to premium customers.
- C. Use origin access control (OAC) to limit the access of non-premium customers.
- D. Generate and activate field-level encryption to block non-premium customers.

**Correct Answer: B**

*Community vote distribution*

B (100%)

- NayeraB**

Highly Voted

8 months, 2 weeks ago

This question page is filled with premium customers I just can't upvoted 8 times
- awsgeek75**

Most Recent

9 months, 4 weeks ago

Selected Answer: B

CloudFront Signed URL with Custom Policy are exactly for this.  
A: Nope, cookies don't help as they don't restrict URL  
C: Wrong. OAC for non-premium customers, how is that even possible without any details here?  
D: Field encryption, while good idea, does not help restricting the content by customer

upvoted 3 times
- pentium75**

10 months, 1 week ago

Selected Answer: B

Authentication is done by Cloudfront, thus B

upvoted 4 times
- ferdzcruz**

10 months, 1 week ago

Content on demand = CloudFront. B

upvoted 3 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: B

Generate and provide CloudFront signed URLs to premium customers.

upvoted 3 times
- TariqKipkemei**

1 year, 3 months ago

Selected Answer: B

Use CloudFront signed URLs or signed cookies to restrict access to documents, business data, media streams, or content that is intended for selected users, for example, users who have paid a fee.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/PrivateContent.html#:~:text=CloudFront%20signed%20URLs>

upvoted 3 times
- james2033**

1 year, 3 months ago

Selected Answer: B

See <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-signed-urls.html#private-content-how-signed-urls-work>

upvoted 2 times
- haoAWS**

1 year, 4 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/PrivateContent.html>  
Notice that A is not correct because it should be CloudFront signed URL, not S3.

upvoted 3 times
- antropaws**

1 year, 4 months ago

Why not C?

upvoted 1 times

  **pentium75** 10 months, 1 week ago

OAC requires the consumers to have an IAM role with access to the S3 content, this is not what we're after here.

upvoted 2 times

  **antropaws** 1 year, 4 months ago

<https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-cloudfront-introduces-origin-access-control-oac/>

upvoted 2 times

  **alexandercamachop** 1 year, 4 months ago

**Selected Answer: B**

Signed URLs

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/PrivateContent.html>

upvoted 3 times

  **haoAWS** 1 year, 4 months ago

Then why A is incorrect?

upvoted 1 times

  **pentium75** 10 months, 1 week ago

Because Authentication is done by Cloudfront, not S3.

upvoted 2 times

A company runs Amazon EC2 instances in multiple AWS accounts that are individually billed. The company recently purchased a Savings Plan. Because of changes in the company's business requirements, the company has decommissioned a large number of EC2 instances. The company wants to use its Savings Plan discounts on its other AWS accounts.

Which combination of steps will meet these requirements? (Choose two.)

- A. From the AWS Account Management Console of the management account, turn on discount sharing from the billing preferences section.
- B. From the AWS Account Management Console of the account that purchased the existing Savings Plan, turn on discount sharing from the billing preferences section. Include all accounts.
- C. From the AWS Organizations management account, use AWS Resource Access Manager (AWS RAM) to share the Savings Plan with other accounts.
- D. Create an organization in AWS Organizations in a new payer account. Invite the other AWS accounts to join the organization from the management account.
- E. Create an organization in AWS Organizations in the existing AWS account with the existing EC2 instances and Savings Plan. Invite the other AWS accounts to join the organization from the management account.

**Correct Answer:** AD

*Community vote distribution*

AD (50%)

AE (39%)

5%



  **Aigerim2010** Highly Voted 1 year, 9 months ago

i had this question today  
upvoted 20 times

  **Nikki013** Highly Voted 1 year, 8 months ago

**Selected Answer: AD**

It is not recommended to have workload on the management account.  
upvoted 6 times

  **JA2018** 5 months, 1 week ago

confirm will fail the AWS WAR process  
upvoted 1 times

  **zdi561** Most Recent 3 months, 1 week ago

**Selected Answer: AD**

google 'aws how to share ec2 shaving plans among accounts of an organization'  
upvoted 1 times

  **FlyingHawk** 3 months, 1 week ago

**Selected Answer: AE**

A E, firstly we create an organization in the existing AWS account with savings plan, so this account becomes a management account. Secondly, from the console of this account, turn on discount sharing.  
upvoted 1 times

  **MaxMingxing** 4 months ago

**Selected Answer: AE**

Answer from ChatGPT  
upvoted 1 times

  **FlyingHawk** 4 months, 3 weeks ago

**Selected Answer: AB**

If the organization is already created, I am feeling A and B are the right one as "To share an Reserved Instance or Savings Plans discount with an account, both accounts must have sharing activated." <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html>. If no AWS organization was there, then you need to create an organization, then invite others join, D. then A and B.  
upvoted 1 times

  **FlyingHawk** 3 months, 1 week ago

B is not correct, Discount sharing can only be enabled from the management account of an AWS Organization, not from the account where the Savings Plan was purchased.  
The correct answer is A E.  
<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html#ri-turn-on-process>

upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: AD**

For me, E makes no sense as the discount is with a new payer and cannot be transferred to an existing account unless customer service is involved.



upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

Also, "A company runs Amazon EC2 instances in multiple AWS accounts that are individually bled"

It's not bled, it is "billed"

upvoted 6 times

  **JA2018** 5 months, 1 week ago

ROTFL :P

upvoted 1 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: AD**

Organization should be created by a new account that is reserved for management. Thus D, followed by A (discount sharing must be enabled in the management account).

upvoted 5 times

  **ErnShm** 1 year, 8 months ago

AE

<https://repost.aws/questions/QUQoJuQLNOTDiyEuCLARIBFQ/transfer-savings-plan-across-organizations#:~:text=AWS%20Support%20can%20transfer%20Savings%20Plans%20from%20the%20management%20account%20to%20a%20member%20account%20or%20from%20a%20member%20account%20to%20the%20management%20account%20within%20a%20single%20Organization%20with%20an%20AWS%20Support%20Case.>

upvoted 1 times

  **lemur88** 1 year, 8 months ago

**Selected Answer: AD**

Not E - it mentions using an account with existing EC2s as the management account, which goes against the best practice for a management account

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_best-practices\\_mgmt-acct.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_best-practices_mgmt-acct.html)

upvoted 4 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: AE**

AE is best

upvoted 1 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: AE**

AE is best

upvoted 1 times

  **james2033** 1 year, 9 months ago

**Selected Answer: AE**

- B is not accepted, because "include all accounts", remove B.  
- D has "Create an organization in AWS Organization in a new payer acount", it is wrong, remove D.  
- at C: AWS Resource Access Manager (AWS RAM) <https://aws.amazon.com/ram/> it is for security, not for billing. Remove C.  
Has A, E remain, and choosed.

A. "turn on discount sharing" is ok. This case: Has discount for many EC2 instances in one account, then want to share with other user. At E, create Organization, then share.

upvoted 2 times

  **pentium75** 1 year, 4 months ago

What is the problem with "include all accounts"?

upvoted 2 times

  **antropaws** 1 year, 10 months ago

**Selected Answer: AE**

I vote AE.

upvoted 2 times

  **MrAWSAssociate** 1 year, 10 months ago

**Selected Answer: AE**

AE are correct !

upvoted 2 times

  **oras2023** 1 year, 10 months ago

**Selected Answer: CD**

It's not good practice to create a payer account with any workload so it must be D.



By the reason that we need Organizations for sharing, then we need to turn on its from our PAYER account. (all sub-accounts start share discounts)

upvoted 1 times

  **oras2023** 1 year, 10 months ago

changed to AD

upvoted 4 times

  **maver144** 1 year, 10 months ago

**Selected Answer: AE**

@alexandercamachop it is AE. I believe its just typo. RAM is not needed anyhow.

upvoted 5 times

  **oras2023** 1 year, 10 months ago

You are right

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html>

upvoted 5 times

A retail company uses a regional Amazon API Gateway API for its public REST APIs. The API Gateway endpoint is a custom domain name that points to an Amazon Route 53 alias record. A solutions architect needs to create a solution that has minimal effects on customers and minimal data loss to release the new version of APIs.



Which solution will meet these requirements?


- A. Create a canary release deployment stage for API Gateway. Deploy the latest API version. Point an appropriate percentage of traffic to the canary stage. After API verification, promote the canary stage to the production stage.
- B. Create a new API Gateway endpoint with a new version of the API in OpenAPI YAML file format. Use the import-to-update operation in merge mode into the API in API Gateway. Deploy the new version of the API to the production stage.
- C. Create a new API Gateway endpoint with a new version of the API in OpenAPI JSON file format. Use the import-to-update operation in overwrite mode into the API in API Gateway. Deploy the new version of the API to the production stage.
- D. Create a new API Gateway endpoint with new versions of the API definitions. Create a custom domain name for the new API Gateway API. Point the Route 53 alias record to the new API Gateway API custom domain name.

**Correct Answer: A**

*Community vote distribution*



A (100%)


-   **AudreyNguyenHN**

Highly Voted 

 1 year, 9 months ago



We made it all the way here. Good luck everyone!

upvoted 18 times
-   **dddddddddddw12**



Highly Voted 

 1 year, 9 months ago

what are the total number of questions this package has as on 14 July 2023 , is it 544 or 551 ?



upvoted 8 times
-   **MrAliMohsan** 8 months, 2 weeks ago

August 2024 - 981 questions :'(



upvoted 3 times
-   **Faridtnx** 1 year, 1 month ago


March 2024 its 825 questions. Constantly adding.

Doe ur question, ExamTopic always shows a few more question in listing compared to actual number

upvoted 5 times
-   **NayeraB** 1 year, 2 months ago

It's 20th of Feb 2024, and it's 798 (it says 804 at the top I donno why tho)

upvoted 3 times
-   **rosanna**



Most Recent 

 5 months ago

**Selected Answer: A**

where do you guys get to know the answer to these things? I study from Stephan's content but lots of topics are uncovered, where am I supposed to find all the exam content in one place?



upvoted 1 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>

upvoted 4 times



  **potomac** 1 year, 5 months ago

**Selected Answer: A**

In a canary release deployment, total API traffic is separated at random into a production release and a canary release with a pre-configured ratio. Typically, the canary release receives a small percentage of API traffic and the production release takes up the rest. The updated API features are only visible to API traffic through the canary. You can adjust the canary traffic percentage to optimize test coverage or performance.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>

upvoted 7 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: A**

Using a canary release deployment allows incremental rollout of the new API version to a percentage of traffic. This minimizes impact on customers and potential data loss during the release.

upvoted 3 times

  **TariqKipkemei** 1 year, 9 months ago

**Selected Answer: A**

Minimal effects on customers and minimal data loss = Canary deployment

upvoted 5 times

  **james2033** 1 year, 9 months ago

**Selected Answer: A**

Key word "canary release". See this term in See: <https://www.jetbrains.com/teamcity/ci-cd-guide/concepts/canary-release/> and/or <https://martinfowler.com/bliki/CanaryRelease.html>

upvoted 2 times

  **Abrar2022** 1 year, 10 months ago

**Selected Answer: A**

keyword: "latest versions on an api"

Canary release is a software development strategy in which a "new version of an API" (as well as other software) is deployed for testing purposes.

upvoted 4 times

  **jkhan2405** 1 year, 10 months ago

**Selected Answer: A**

It's A

upvoted 2 times

  **alexandercamachop** 1 year, 11 months ago

**Selected Answer: A**

A. Create a canary release deployment stage for API Gateway. Deploy the latest API version. Point an appropriate percentage of traffic to the canary stage. After API verification, promote the canary stage to the production stage.

Canary release meaning only certain percentage of the users.

upvoted 4 times

A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.

Which solution will meet these requirements?



- A. Update the Route 53 records to use a latency routing policy. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
- B. Set up a Route 53 active-passive failover configuration. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
- C. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoints. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
- D. Update the Route 53 records to use a multivalue answer routing policy. Create a health check. Direct traffic to the website if the health check passes. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.


**Correct Answer: B**

*Community vote distribution*

B (83%)

D (17%)



-   **TariqKipkemei**


Highly Voted 

 1 year, 5 months ago

Selected Answer: B

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



upvoted 6 times
-   **ssa03**


Highly Voted 

 1 year, 8 months ago

Selected Answer: B



B is correct

upvoted 5 times
-   **Danilus**

Most Recent 

 5 months, 3 weeks ago



key -minimizes changes and infrastructure overhead  
key- backup static error page  
Option B is the appropriate solution because it sets up automatic failover to an alternative endpoint (the static error page in S3) when the ALB becomes unavailable. Option D is not viable as it does not implement a true failover mechanism and may continue to direct traffic to an unhealthy ALB.

upvoted 1 times
-   **Guru4Cloud** 1 year, 8 months ago



Selected Answer: D

Setting up a Route 53 active-passive failover configuration with the ALB as the primary endpoint and an Amazon S3 static website as the passive endpoint meets the requirements with minimal overhead.



Route 53 health checks can monitor the ALB health. If the ALB becomes unhealthy, traffic will automatically failover to the S3 static website. This provides automatic failover with minimal configuration changes

upvoted 3 times
-   **Guru4Cloud** 1 year, 8 months ago

Sorry. I mean B



upvoted 6 times
-   **Nirav1112** 1 year, 9 months ago

B is correct

upvoted 3 times
-   **mrsoa** 1 year, 9 months ago

Selected Answer: B

B seems correct

upvoted 4 times
-   **Bmaster** 1 year, 9 months ago

B is correct..

<https://repost.aws/knowledge-center/fail-over-s3-r53>  
upvoted 4 times

  **awsgeek75** 1 year, 3 months ago

Nice link find!  
upvoted 2 times

A recent analysis of a company's IT expenses highlights the need to reduce backup costs. The company's chief information officer wants to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. The company must preserve the existing investment in the on-premises backup applications and workflows.

What should a solutions architect recommend?

- A. Set up AWS Storage Gateway to connect with the backup applications using the NFS interface.
- B. Set up an Amazon EFS file system that connects with the backup applications using the NFS interface.
- C. Set up an Amazon EFS file system that connects with the backup applications using the iSCSI interface.
- D. Set up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface.

**Correct Answer:** D

Community vote distribution

D (100%)

  **awsgeek75** Highly Voted 9 months, 4 weeks ago

**Selected Answer:** D

Tape... lol

The company must preserve it's existing investment so they want to keep using existing applications. This means EFS won't work. and NFS may not be compatible. VTL is the only thing that may be compatible with an application workflow that backups to tapes.

Who the hell comes up with these questions!

upvoted 6 times

  **TariqKipkemei** Highly Voted 11 months, 3 weeks ago

**Selected Answer:** D

Use Tape Gateway to replace physical tapes on premises with virtual tapes on AWS—reducing your data storage costs without changing your tape-based backup workflows. Tape Gateway supports all leading backup applications and caches virtual tapes on premises for low-latency data access. It compresses your tape data, encrypts it, and stores it in a virtual tape library in Amazon Simple Storage Service (Amazon S3). From there, you can transfer it to either Amazon S3 Glacier Flexible Retrieval or Amazon S3 Glacier Deep Archive to help minimize your long-term storage costs.

<https://aws.amazon.com/storagegateway/vtl/#:~:text=Use-,Tape%20Gateway,-to%20replace%20physical>

upvoted 5 times

  **Nisarg2121** Most Recent 1 year ago

**Selected Answer:** D

Tape Gateway is use for attache with app.

upvoted 4 times

  **gouranga45** 1 year ago

**Selected Answer:** D

Option says it all

upvoted 4 times

  **Po\_chih** 1 year ago

**Selected Answer:** D

Tape Gateway enables you to replace using physical tapes on premises with virtual tapes in AWS without changing existing backup workflows. Tape Gateway supports all leading backup applications and caches virtual tapes on premises for low-latency data access. Tape Gateway encrypts data between the gateway and AWS for secure data transfer, and compresses data and transitions virtual tapes between Amazon S3 and Amazon S3 Glacier Flexible Retrieval, or Amazon S3 Glacier Deep Archive, to minimize storage costs.

upvoted 3 times

  **ssa03** 1 year, 2 months ago

**Selected Answer:** D

[https://aws.amazon.com/storagegateway/vtl/?nc1=h\\_ls](https://aws.amazon.com/storagegateway/vtl/?nc1=h_ls)

upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer:** D

Set up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface.

upvoted 2 times



**Bmaster** 1 year, 3 months ago

D is correct

[https://aws.amazon.com/storagegateway/vtl/?nc1=h\\_ls](https://aws.amazon.com/storagegateway/vtl/?nc1=h_ls)

upvoted 2 times

A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming data. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.
- B. Use AWS Glue to deliver streaming data to Amazon S3.
- C. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.
- D. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.

**Correct Answer: A**

*Community vote distribution*

A (83%)

D (17%)

- emakid** 10 months, 1 week ago

**Selected Answer: A**

Amazon Kinesis Data Firehose is a fully managed service for delivering real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. It requires minimal setup and maintenance, automatically scales to match the throughput of your data, and offers near real-time data delivery with minimal operational overhead.

upvoted 3 times
- awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

High volume streaming data = Kinesis  
B: Glue is for ETL (to S3 is ok) but not for streaming  
C: Lambda more overhead  
D: Streaming != Data migration

upvoted 3 times
- pentium75** 1 year, 4 months ago

**Selected Answer: A**

sensor data = Kinesis

upvoted 4 times
- TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

Amazon Kinesis Data Firehose: Capture, transform, and load data streams into AWS data stores (S3) in near real-time.

[https://aws.amazon.com/pm/kinesis/?gclid=CjwKCAiAu9yqBhBmEiwAHTx5px9z182o0HBEX0BGXU7VeOCOdNpkJMxgbSfvcHINKN4NHVnbEa0Y1xoCuU0QAvD\\_BwE&trk=239a97c0-9c5d-42a5-ac65-7381b62f3756&sc\\_channel=ps&ef\\_id=CjwKCAiAu9yqBhBmEiwAHTx5px9z182o0HBEX0BGXU7VeOCOdNpkJMxgbSfvcHINKN4NHVnbEa0Y1xoCuU0QAvD\\_BwE:G:s&s\\_kwcid=AL!4422!3!651612444428!e!!g!!kinesis%20firehose!19836376048!149982297311#:~:text=Kinesis%20Data%20Firehose-,Capture%2C,-transform%2C%20and%20load](https://aws.amazon.com/pm/kinesis/?gclid=CjwKCAiAu9yqBhBmEiwAHTx5px9z182o0HBEX0BGXU7VeOCOdNpkJMxgbSfvcHINKN4NHVnbEa0Y1xoCuU0QAvD_BwE&trk=239a97c0-9c5d-42a5-ac65-7381b62f3756&sc_channel=ps&ef_id=CjwKCAiAu9yqBhBmEiwAHTx5px9z182o0HBEX0BGXU7VeOCOdNpkJMxgbSfvcHINKN4NHVnbEa0Y1xoCuU0QAvD_BwE:G:s&s_kwcid=AL!4422!3!651612444428!e!!g!!kinesis%20firehose!19836376048!149982297311#:~:text=Kinesis%20Data%20Firehose-,Capture%2C,-transform%2C%20and%20load)

upvoted 3 times
- potomac** 1 year, 5 months ago

**Selected Answer: A**

A for sure

upvoted 3 times
- ssa03** 1 year, 8 months ago

**Selected Answer: A**

Correct Answer: A



upvoted 3 times
- manOfThePeople** 1 year, 8 months ago



A is the answer, near real-time = Kinesis Data Firehose.



upvoted 4 times
- Guru4Cloud** 1 year, 8 months ago



**Selected Answer: D**



Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3  
upvoted 3 times



  **pentium75** 1 year, 4 months ago  
That is A  
upvoted 2 times



  **bjexamprep** 1 year, 8 months ago  
**Selected Answer: D**  
Kinesis Data Firehose is only real-time answer  
upvoted 3 times

  **pentium75** 1 year, 4 months ago  
That is A  
upvoted 3 times

  **mrsoa** 1 year, 9 months ago  
**Selected Answer: A**  
A is the correct answer  
upvoted 3 times

  **Deepakin96** 1 year, 9 months ago  
**Selected Answer: A**  
Kinesis = Near Real Time  
upvoted 4 times

  **Kaiden123** 1 year, 9 months ago  
**Selected Answer: A**  
Data collection in near real time = Amazon Kinesis Data Firehose  
upvoted 4 times

  **Bmaster** 1 year, 9 months ago  
A is correct..  
upvoted 2 times

A company has separate AWS accounts for its finance, data analytics, and development departments. Because of costs and security concerns, the company wants to control which services each AWS account can use.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Systems Manager templates to control which AWS services each department can use.
- B. Create organization units (OUs) for each department in AWS Organizations. Attach service control policies (SCPs) to the OUs.
- C. Use AWS CloudFormation to automatically provision only the AWS services that each department can use.
- D. Set up a list of products in AWS Service Catalog in the AWS accounts to manage and control the usage of specific AWS services.

**Correct Answer: B**

*Community vote distribution*

B (89%)

11%

🗳️ 👤 **zdi561** 2 months, 3 weeks ago

**Selected Answer: D**

SCP set a guardrail (what you can not do, permission on IAM user and role), ASC provides templates services that you can only use.  
upvoted 1 times

🗳️ 👤 **awsgeek75** 9 months, 4 weeks ago

**Selected Answer: B**

Departments = Organizational Units  
upvoted 3 times

🗳️ 👤 **TariqKipkemei** 11 months, 3 weeks ago

**Selected Answer: B**

Create organization units (OUs) for each department in AWS Organizations. Attach service control policies (SCPs) to the OUs  
upvoted 2 times

🗳️ 👤 **ssa03** 1 year, 2 months ago

**Selected Answer: B**

Correct Answer: B  
upvoted 2 times

🗳️ 👤 **lemur88** 1 year, 2 months ago

**Selected Answer: B**

SCPs to centralize permissioning  
upvoted 2 times

🗳️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

Create organization units (OUs) for each department in AWS Organizations. Attach service control policies (SCPs) to the OUs.  
upvoted 2 times

🗳️ 👤 **xyb** 1 year, 2 months ago

**Selected Answer: B**

control services --> SCP  
upvoted 2 times

🗳️ 👤 **Ale1973** 1 year, 2 months ago

**Selected Answer: D**

My rational: Scenary is "A company has separate AWS accounts", it is not mentioning anything about use of Organizations or needs related to centralized managment of these accounts.  
Then, set up a list of products in AWS Service Catalog in the AWS accounts (on each AWS account) is the best way to manage and control the usage of specific AWS services.  
upvoted 1 times



🗳️ 👤 **pentium75** 10 months, 1 week ago

"Separate AWS accounts" just says that it's multiple accounts, it does not indicate that they are NOT connected into a organization.

Service Catalog alone does not restrict anything. You'd need to create a service in Service Catalog for everything you're allowing to use, then

grant permissions on those services, and you'd need to remove other permissions from everyone. All of which is not mentioned in D. Just "setting up a list of products in AWS Service Catalog in the AWS accounts" will not restrict anyone from doing what he could do before.

upvoted 3 times

  **mrsoa** 1 year, 3 months ago

**Selected Answer: B**

BBBBBBBBBB

upvoted 2 times

  **Deepakin96** 1 year, 3 months ago

**Selected Answer: B**

To control different AWS account you required AWS Organisation

upvoted 2 times

  **Bmaster** 1 year, 3 months ago

B is correct!!!!

upvoted 2 times

A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a third-party provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead.

What should the solutions architect do to meet these requirements?

- A. Deploy a NAT instance in the VPC. Route all the internet-based traffic through the NAT instance.
- B. Deploy a NAT gateway in the public subnets. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
- C. Configure an internet gateway and attach it to the VPC. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
- D. Configure a virtual private gateway and attach it to the VPC. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

**Correct Answer: B**

*Community vote distribution*

B (100%)

🗲️ 👤 **awsgeek75** 9 months, 4 weeks ago

**Selected Answer: B**

A: Probably an old question so this option is here but NAT instance is overhead  
C: Not secure as IG opens up a lot of things  
D: VPG connects to a service  
B: NG is managed solution. Secure by config  
upvoted 3 times

🗲️ 👤 **TariqKipkemei** 11 months, 3 weeks ago

**Selected Answer: B**

Deploy a NAT gateway in the public subnets. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway  
upvoted 4 times

🗲️ 👤 **ssa03** 1 year, 2 months ago

**Selected Answer: B**

Correct Answer: B  
upvoted 4 times

🗲️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

Deploy a NAT gateway in the public subnets. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.  
upvoted 3 times

🗲️ 👤 **Deepakin96** 1 year, 3 months ago

**Selected Answer: B**

NAT Gateway is safe  
upvoted 3 times

🗲️ 👤 **Bmaster** 1 year, 3 months ago

B is correct  
upvoted 2 times

A company is using AWS Key Management Service (AWS KMS) keys to encrypt AWS Lambda environment variables. A solutions architect needs to ensure that the required permissions are in place to decrypt and use the environment variables.

Which steps must the solutions architect take to implement the correct permissions? (Choose two.)

- A. Add AWS KMS permissions in the Lambda resource policy.
- B. Add AWS KMS permissions in the Lambda execution role.
- C. Add AWS KMS permissions in the Lambda function policy.
- D. Allow the Lambda execution role in the AWS KMS key policy.
- E. Allow the Lambda resource policy in the AWS KMS key policy.

Correct Answer: *BD*

Community vote distribution

BD (100%)

- Guru4Cloud

Highly Voted

1 year, 2 months ago

Selected Answer: BD

To decrypt environment variables encrypted with AWS KMS, Lambda needs to be granted permissions to call KMS APIs. This is done in two places:  
  
The Lambda execution role needs kms:Decrypt and kms:GenerateDataKey permissions added. The execution role governs what AWS services the function code can access.  
The KMS key policy needs to allow the Lambda execution role to have kms:Decrypt and kms:GenerateDataKey permissions for that specific key.  
This allows the execution role to use that particular key.  
upvoted 7 times
- wizcloudifa

Highly Voted

6 months ago

Selected Answer: BD

As per the principle of least privilege, granting permissions = role level  
upvoted 5 times
- TariqKipkemei

Most Recent

11 months, 3 weeks ago

Selected Answer: BD

Allow the Lambda execution role in the AWS KMS key policy then add AWS KMS permissions in the role.  
upvoted 3 times
- ssa03

1 year, 2 months ago

Selected Answer: BD

Correct Answer: BD  
upvoted 3 times
- Nirav1112

1 year, 2 months ago

its B & D  
upvoted 2 times
- mrsoa

1 year, 3 months ago

Selected Answer: BD

BD BD BD BD  
upvoted 2 times
- Deepakin96

1 year, 3 months ago

Selected Answer: BD

Its B and D  
upvoted 2 times
- Bmaster

1 year, 3 months ago

My choice is B,D  
upvoted 2 times

A company has a financial application that produces reports. The reports average 50 KB in size and are stored in Amazon S3. The reports are frequently accessed during the first week after production and must be stored for several years. The reports must be retrievable within 6 hours.

Which solution meets these requirements MOST cost-effectively?

- A. Use S3 Standard. Use an S3 Lifecycle rule to transition the reports to S3 Glacier after 7 days.
- B. Use S3 Standard. Use an S3 Lifecycle rule to transition the reports to S3 Standard-Infrequent Access (S3 Standard-IA) after 7 days.
- C. Use S3 Intelligent-Tiering. Configure S3 Intelligent-Tiering to transition the reports to S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Glacier.
- D. Use S3 Standard. Use an S3 Lifecycle rule to transition the reports to S3 Glacier Deep Archive after 7 days.

**Correct Answer: A**

*Community vote distribution*

A (74%)C (24%)

- zjcorpuz**

Highly Voted

1 year, 9 months ago

Answer is A  
Amazon S3 Glacier:  
Expedited Retrieval: Provides access to data within 1-5 minutes.  
Standard Retrieval: Provides access to data within 3-5 hours.  
Bulk Retrieval: Provides access to data within 5-12 hours.  
Amazon S3 Glacier Deep Archive:  
Standard Retrieval: Provides access to data within 12 hours.  
Bulk Retrieval: Provides access to data within 48 hours.  
upvoted 25 times
- oayoade**

Highly Voted

1 year, 8 months ago

**Selected Answer: C**

All the "....after 7 days" options are wrong.  
Before you transition objects to S3 Standard-IA or S3 One Zone-IA, you must store them for at least 30 days in Amazon S3  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html#:~:text=Minimum%20Days%20for%20Transition%20to%20S3%20Standard%2DIA%20or%20S3%20One%20Zone%2DIA>  
upvoted 11 times

**LeonSauveterre**

4 months, 4 weeks ago

Objects must be stored for at least 30 days before transitioning to S3 Standard-IA or S3 One Zone-IA.  
  
Option A is to S3 Glacier. A is correct.  
upvoted 2 times

**MatAlves**

7 months, 3 weeks ago

After reading the document from your link, it's clear that NONE of the restrictions apply to S3 standard, but to S3 Standard IA.  
  
A is the correct answer.  
upvoted 2 times

**Hades2231**

1 year, 8 months ago

This is worth noticing! Glad I came across your comment 1 day before my test.  
upvoted 4 times

**Marco\_St**

1 year, 3 months ago

so Could I ask is A or C for this question? I voted for A but it seems you had the same question in the exam and it was C? Thanks! I will attend the exam soon.  
upvoted 2 times

**franbarberan**

1 year, 7 months ago

the 7 days limitation is only if you want to move from s3 standart to S3 Standard-IA or S3 One Zone-IA, if you move to s3 glacier dont have this limitation, correct answer is A  
upvoted 14 times

**TEC65**

Most Recent

1 week, 3 days ago

**Selected Answer: A**

a is the answer  
upvoted 1 times

🗄️ 👤 **rosanna** 5 months ago

**Selected Answer: A**

there's a period before you can transition to any storage class (down the hierarchy) "except for glacier of all types" a minimum of 30 days. However, when transitioning to Glacier directly there's no minimum period to wait so you can choose "A" safely

upvoted 1 times

🗄️ 👤 **1e22522** 9 months ago

**Selected Answer: A**

Its A ya bunch of nerds

upvoted 3 times

🗄️ 👤 **Gape4** 10 months, 1 week ago

**Selected Answer: C**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html>

upvoted 1 times

🗄️ 👤 **Linuslin** 12 months ago

**Selected Answer: A**

C is incorrect.

Unsupported lifecycle transitions

Amazon S3 does not support any of the following lifecycle transitions.

You can't transition from the following:

Any storage class to the S3 Standard storage class.

Any storage class to the Reduced Redundancy Storage (RRS) class.

The S3 Intelligent-Tiering storage class to the S3 Standard-IA storage class.

The S3 One Zone-IA storage class to the S3 Intelligent-Tiering, S3 Standard-IA, or S3 Glacier Instant Retrieval storage classes.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html>

upvoted 4 times

🗄️ 👤 **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

BC are lifecycle with tiering and infrequent access which are not required here.

D is deep archive and can take hours to retrieve so it is not suitable

A is cheapest workable option

upvoted 4 times

🗄️ 👤 **Marco\_St** 1 year, 3 months ago

**Selected Answer: A**

frequent access pattern- Standard.

upvoted 2 times

🗄️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: A**

Not B - More expensive than A

Not C - Intelligent-Tiering moves only objects of at least 128 KB

Not D - Glacier Deep Archive takes more than 6 hours to retrieve

upvoted 4 times

🗄️ 👤 **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

Any option with S3 Intelligent-Tiering is out, this is only required when the access patterns are unknown.

From the question the access patterns are well known, enough to tie the frequently accessed reports to S3 standard and transition them to S3 glacier after 7days.

upvoted 4 times

🗄️ 👤 **iwannabeawsgod** 1 year, 6 months ago

**Selected Answer: A**

its A for me

upvoted 3 times

🗄️ 👤 **Carlos\_0** 1 year, 6 months ago

**Selected Answer: A**

Tiene mas sentido

upvoted 2 times

🗄️ 👤 **sl2man** 1 year, 6 months ago

**Selected Answer: A**

Option A

Amazon S3 Glacier Standard Retrieval: Provides access to data within 3-5 hours.

upvoted 4 times

🗨️ 👤 **Ramdi1** 1 year, 7 months ago

**Selected Answer: A**

most cost effective has to be glacier so A

With C it is using intelligence tiering which is 30 days minimum from what I have read, I may be wrong on how I read that.

upvoted 2 times

🗨️ 👤 **tabbyDolly** 1 year, 7 months ago

answer A

frequent access during the first week -> keeps data in s3 standard for 7 days

stored for several year and retrievable within 6 hours -> can be moved to s3 glacier for data archive purpose

upvoted 2 times

🗨️ 👤 **anikety123** 1 year, 7 months ago

**Selected Answer: A**

Its A. Data cannot be transitioned from Intelligent Tiering to Standard IA

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-transition-general-considerations.html>

upvoted 4 times

A company needs to optimize the cost of its Amazon EC2 instances. The company also needs to change the type and family of its EC2 instances every 2-3 months.

What should the company do to meet these requirements?

- A. Purchase Partial Upfront Reserved Instances for a 3-year term.
- B. Purchase a No Upfront Compute Savings Plan for a 1-year term.
- C. Purchase All Upfront Reserved Instances for a 1-year term.
- D. Purchase an All Upfront EC2 Instance Savings Plan for a 1-year term.

Correct Answer: B

Community vote distribution

B (100%)

- Guru4Cloud

Highly Voted

8 months, 2 weeks ago

Selected Answer: B

The key considerations are:  
  
The company needs flexibility to change EC2 instance types and families every 2-3 months. This rules out Reserved Instances which lock you into an instance type and family for 1-3 years.  
A Compute Savings Plan allows switching instance types and families freely within the term as needed. No Upfront is more flexible than All Upfront.  
A 1-year term balances commitment and flexibility better than a 3-year term given the company's changing needs.  
With No Upfront, the company only pays for usage monthly without an upfront payment. This optimizes cost.  
upvoted 10 times
- Kiki\_Pass

Highly Voted

9 months ago

Selected Answer: B

"EC2 Instance Savings Plans give you the flexibility to change your usage between instances WITHIN a family in that region. "  
<https://aws.amazon.com/savingsplans/compute-pricing/>  
upvoted 5 times
- TariqKipkemei

Most Recent

5 months, 2 weeks ago

Selected Answer: B

Only Compute Savings Plan allows you to change instance family.  
upvoted 4 times
- avkya

8 months, 3 weeks ago

Selected Answer: B

" needs to change the type and family of its EC2 instances". that means B I think.  
upvoted 3 times
- mrsoa

9 months ago

Selected Answer: B

B is the right answer  
upvoted 3 times
- Bmaster

9 months ago

B is correct..  
'EC2 Instance Savings Plans' can't change 'family'.  
upvoted 2 times
- Josantru

9 months, 1 week ago

Correct B.  
To change 'Family' always Compute saving plan, right?  
upvoted 4 times

A solutions architect needs to review a company's Amazon S3 buckets to discover personally identifiable information (PII). The company stores the PII data in the us-east-1 Region and us-west-2 Region.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure Amazon Macie in each Region. Create a job to analyze the data that is in Amazon S3.
- B. Configure AWS Security Hub for all Regions. Create an AWS Config rule to analyze the data that is in Amazon S3.
- C. Configure Amazon Inspector to analyze the data that is in Amazon S3.
- D. Configure Amazon GuardDuty to analyze the data that is in Amazon S3.

**Correct Answer: A**

*Community vote distribution*

A (100%)

-  **Guru4Cloud** Highly Voted  1 year, 2 months ago

Selected Answer: A

The key reasons are:

Amazon Macie is designed specifically for discovering and classifying sensitive data like PII in S3. This makes it the optimal service to use. Macie can be enabled directly in the required Regions rather than enabling it across all Regions which is unnecessary. This minimizes overhead. Macie can be set up to automatically scan the specified S3 buckets on a schedule. No need to create separate jobs. Security Hub is for security monitoring across AWS accounts, not specific for PII discovery. More overhead than needed. Inspector and GuardDuty are not built for PII discovery in S3 buckets. They provide broader security capabilities.

upvoted 6 times
-  **awsgeek75** Most Recent  9 months, 3 weeks ago

Selected Answer: A

PII = Macie  
Security Hub: Organisation security and logging not for PII  
Inspector: Infra vulnerability management  
GuardDuty: Network protection

upvoted 4 times
-  **TariqKipkemei** 11 months, 2 weeks ago

Selected Answer: A

Amazon Macie = PII

upvoted 2 times
-  **mrsoa** 1 year, 3 months ago

Selected Answer: A

AWS Macie = PII detection

upvoted 4 times
-  **Deepakin96** 1 year, 3 months ago

Selected Answer: A

Amazon Macie will identify all PII

upvoted 3 times

A company's SAP application has a backend SQL Server database in an on-premises environment. The company wants to migrate its on-premises application and database server to AWS. The company needs an instance type that meets the high demands of its SAP database. On-premises performance data shows that both the SAP application and the database have high memory utilization.



Which solution will meet these requirements?


- A. Use the compute optimized instance family for the application. Use the memory optimized instance family for the database.
- B. Use the storage optimized instance family for both the application and the database.
- C. Use the memory optimized instance family for both the application and the database.
- D. Use the high performance computing (HPC) optimized instance family for the application. Use the memory optimized instance family for the database.

**Correct Answer: C**

*Community vote distribution*

C (100%)



-   **Guru4Cloud**


Highly Voted 

 8 months, 2 weeks ago

Selected Answer: C



Since both the app and database have high memory needs, the memory optimized family like R5 instances meet those requirements well. Using the same instance family simplifies management and operations, rather than mixing instance types. Compute optimized instances may not provide enough memory for the SAP app's needs. Storage optimized is overkill for the database's compute and memory needs. HPC is overprovisioned for the SAP app.


upvoted 16 times
-   **manOfThePeople**

Highly Voted 

 8 months ago

High memory utilization = memory optimized.  
C is the answer



upvoted 5 times
-   **LeonSauveterre**

Most Recent 

 4 months, 4 weeks ago



Selected Answer: C

Why D is incorrect: HPC family is geared toward compute-heavy workloads such as simulations or scientific calculations, not appropriate for a memory-intensive application like SAP.

upvoted 1 times
-   **TariqKipkemei** 5 months, 2 weeks ago

Selected Answer: C

Use the memory optimized instance family for both the application and the database

upvoted 3 times
-   **mrsoa** 9 months ago

Selected Answer: C

I thyink its C

upvoted 3 times

A company runs an application in a VPC with public and private subnets. The VPC extends across multiple Availability Zones. The application runs on Amazon EC2 instances in private subnets. The application uses an Amazon Simple Queue Service (Amazon SQS) queue.

A solutions architect needs to design a secure solution to establish a connection between the EC2 instances and the SQS queue.

Which solution will meet these requirements?

- A. Implement an interface VPC endpoint for Amazon SQS. Configure the endpoint to use the private subnets. Add to the endpoint a security group that has an inbound access rule that allows traffic from the EC2 instances that are in the private subnets.
- B. Implement an interface VPC endpoint for Amazon SQS. Configure the endpoint to use the public subnets. Attach to the interface endpoint a VPC endpoint policy that allows access from the EC2 instances that are in the private subnets.
- C. Implement an interface VPC endpoint for Amazon SQS. Configure the endpoint to use the public subnets. Attach an Amazon SQS access policy to the interface VPC endpoint that allows requests from only a specified VPC endpoint.
- D. Implement a gateway endpoint for Amazon SQS. Add a NAT gateway to the private subnets. Attach an IAM role to the EC2 instances that allows access to the SQS queue.

**Correct Answer: A**

*Community vote distribution*

A (100%)

 **Guru4Cloud** Highly Voted 1 year, 2 months ago

**Selected Answer: A**

An interface VPC endpoint is a private way to connect to AWS services without having to expose your VPC to the public internet. This is the most secure way to connect to Amazon SQS from the private subnets.

Configuring the endpoint to use the private subnets ensures that the traffic between the EC2 instances and the SQS queue is only within the VPC. This helps to protect the traffic from being intercepted by a malicious actor.

Adding a security group to the endpoint that has an inbound access rule that allows traffic from the EC2 instances that are in the private subnets further restricts the traffic to only the authorized sources. This helps to prevent unauthorized access to the SQS queue.

upvoted 11 times

 **Bmaster** Highly Voted 1 year, 3 months ago

A is correct.

B,C: 'Configuring endpoints to use public subnets' --> Invalid

D: No Gateway Endpoint for SQS.

upvoted 6 times

 **awsgeek75** Most Recent 9 months, 3 weeks ago

**Selected Answer: A**

BC are using public subnets so not useful for security

D uses gateway endpoint which is not useful to connect to SQS

A: <https://docs.aws.amazon.com/vpc/latest/privatelink/aws-services-privatelink-support.html>

upvoted 2 times

 **awsgeek75** 9 months, 2 weeks ago

Sorry, the link is wrong for A. Please ignore it!

upvoted 2 times

 **ShawnTang** 10 months, 3 weeks ago

A seems the most suitable,

but security group can't add to the endpoint directly, right?

upvoted 2 times

 **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: A**

Answer is A

upvoted 2 times

 **TariqKipkemei** 11 months, 2 weeks ago

Interface endpoints enable connectivity to services over AWS PrivateLink. It is a collection of one or more elastic network interfaces with a private IP address that serves as an entry point for traffic destined to a supported service.



Implement an interface VPC endpoint for Amazon SQS. Configure the endpoint to use the private subnets. Add to the endpoint a security group that has an inbound access rule that allows traffic from the EC2 instances that are in the private subnets.

upvoted 2 times

  **potomac** 1 year ago

A is correct

upvoted 2 times

  **mrsoa** 1 year, 3 months ago

**Selected Answer: A**

I think its A

upvoted 2 times

A solutions architect is using an AWS CloudFormation template to deploy a three-tier web application. The web application consists of a web tier and an application tier that stores and retrieves user data in Amazon DynamoDB tables. The web and application tiers are hosted on Amazon EC2 instances, and the database tier is not publicly accessible. The application EC2 instances need to access the DynamoDB tables without exposing API credentials in the template.

What should the solutions architect do to meet these requirements?

- A. Create an IAM role to read the DynamoDB tables. Associate the role with the application instances by referencing an instance profile.
- B. Create an IAM role that has the required permissions to read and write from the DynamoDB tables. Add the role to the EC2 instance profile, and associate the instance profile with the application instances.
- C. Use the parameter section in the AWS CloudFormation template to have the user input access and secret keys from an already-created IAM user that has the required permissions to read and write from the DynamoDB tables.
- D. Create an IAM user in the AWS CloudFormation template that has the required permissions to read and write from the DynamoDB tables. Use the GetAtt function to retrieve the access and secret keys, and pass them to the application instances through the user data.

**Correct Answer: B**

*Community vote distribution*

B (87%)

13%

  **upliftinghut** 9 months, 2 weeks ago



**Selected Answer: B**

best practice is using IAM role for database access. From app to DB => need both read & write, only B meets these 2  
upvoted 3 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: B**

Application "stores and retrieves" data in DynamoDB while A grants only access "to read".  
upvoted 3 times

  **Nisarg2121** 1 year ago

**Selected Answer: B**

B is correct, A total wrong because "read the DynamoDB tables", so what about write in database.  
upvoted 4 times

  **darekw** 1 year, 2 months ago

question says: ...application tier stores and retrieves user data in Amazon DynamoDB tables... so it needs read and write access  
A) is only read access  
B) seems to be the right answer  
upvoted 2 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**

Option B is the correct approach to meet the requirements:

Create an IAM role with permissions to access DynamoDB  
Add the IAM role to an EC2 Instance Profile  
Associate the Instance Profile with the application EC2 instances  
This allows the instances to assume the IAM role to obtain temporary credentials to access DynamoDB.  
upvoted 3 times

  **anibinaadi** 1 year, 2 months ago

Explanation. Both A and B seems suitable. But Option A is incorrect because it says "Associate the role with the application instances by referencing an instance profile". Which just only a Part of the solution.  
In API/AWS CLI following steps are required to complete the Role-> instance profile association-> to instance.  
1. Create an IAM Role  
2. add-role-to-instance-profile (aws iam add-role-to-instance-profile --role-name S3Access --instance-profile-name Webserver)  
3. associate-iam-instance-profile (aws ec2 associate-iam-instance-profile --instance-id i-123456789abcde123 --iam-instance-profile Name=admin-role)  
hence Option B is correct.  
upvoted 3 times

  **DannyKang5649** 1 year, 2 months ago

**Selected Answer: B**

Why "No read and write" ? The question clearly states that application tier STORE and RETRIEVE the data from DynamoDB. Which means write and read... I think answer should be B

upvoted 3 times

  **xyb** 1 year, 2 months ago

**Selected Answer: B**

<https://www.examttopics.com/discussions/amazon/view/80755-exam-aws-certified-solutions-architect-associate-saa-c02/>

upvoted 2 times

  **Ale1973** 1 year, 2 months ago

**Selected Answer: B**

My rationl: Option A is wrong because the scenario says "stores and retrieves user data in Amazon DynamoDB tables", STORES and RETRIVE, if you set a role to READ, you can write on DinamoDB database

upvoted 2 times

  **mrsoa** 1 year, 2 months ago

**Selected Answer: A**



AAAAAAAAA

upvoted 2 times

  **pentium75** 10 months, 1 week ago

No because it grants only read access

upvoted 3 times

  **kangho** 1 year, 2 months ago

**Selected Answer: A**

A is correct

upvoted 1 times

  **pentium75** 10 months, 1 week ago

No because it grants only read access

upvoted 3 times

A solutions architect manages an analytics application. The application stores large amounts of semistructured data in an Amazon S3 bucket. The solutions architect wants to use parallel data processing to process the data more quickly. The solutions architect also wants to use information that is stored in an Amazon Redshift database to enrich the data.

Which solution will meet these requirements?

- A. Use Amazon Athena to process the S3 data. Use AWS Glue with the Amazon Redshift data to enrich the S3 data.
- B. Use Amazon EMR to process the S3 data. Use Amazon EMR with the Amazon Redshift data to enrich the S3 data.
- C. Use Amazon EMR to process the S3 data. Use Amazon Kinesis Data Streams to move the S3 data into Amazon Redshift so that the data can be enriched.
- D. Use AWS Glue to process the S3 data. Use AWS Lake Formation with the Amazon Redshift data to enrich the S3 data.

**Correct Answer: B**

*Community vote distribution*

B (79%)

A (21%)

  **Guru4Cloud** Highly Voted 1 year, 2 months ago

**Selected Answer: B**

Option B is the correct solution that meets the requirements:

Use Amazon EMR to process the semi-structured data in Amazon S3. EMR provides a managed Hadoop framework optimized for processing large datasets in S3.

EMR supports parallel data processing across multiple nodes to speed up the processing.

EMR can integrate directly with Amazon Redshift using the EMR-Redshift integration. This allows querying the Redshift data from EMR and joining it with the S3 data.

This enables enriching the semi-structured S3 data with the information stored in Redshift

upvoted 18 times

  **zjcorpuz** Highly Voted 1 year, 3 months ago

By combining AWS Glue and Amazon Redshift, you can process the semistructured data in parallel using Glue ETL jobs and then store the processed and enriched data in a structured format in Amazon Redshift. This approach allows you to perform complex analytics efficiently and at scale.

upvoted 9 times

  **upliftinghut** Most Recent 9 months, 2 weeks ago

**Selected Answer: B**

D: not relevant, data is semistructured and Glue is more batch than stream data

A: not correct, Athena is for querying data

B & C look ok but C is out => redundant with Kinesis data stream; EMR already processed data as input into Redshift for parallel processing

Only B is most logical

upvoted 4 times

  **awsgeek75** 9 months, 2 weeks ago

**Selected Answer: B**

Key requirement: parallel data processing

parallel data processing is EMR (Kind of Apache Hadoop) so it only leave B and C

C is Kinesis to Redshift which is pointless logic here

B EMR for S3 and EMR for Redshift gives maximum parallel processing here

upvoted 3 times

  **pentium75** 10 months, 1 week ago

**Selected Answer: B**

A has a pitfall, "use Amazon Athena to PROCESS the data". With Athena you can query, not process, data.

C is wrong because Kinesis has no place here.

D is wrong because it does not process the Redshift data, and Glue does ETL, not analyze

Thus it's B. EMR can use semi-structured data from from S3 and structured data from Redshift and is ideal for "parallel data processing" of "large amounts" of data.

upvoted 7 times

  **aws94** 10 months, 3 weeks ago

**Selected Answer: B**

large amount of data + parallel data processing = EMR

upvoted 3 times

  **[Removed]** 10 months, 4 weeks ago

**Selected Answer: A**

Amazon Athena is an interactive query service that makes it easy to analyze data directly in Amazon Simple Storage Service (Amazon S3) using standard SQL.

upvoted 1 times

  **pentium75** 10 months, 1 week ago

Y, but A says "process", not "query" data with Athena.

upvoted 2 times

  **SHAAHIBHUSHANAWS** 11 months ago

Selected Answer: D

Glue use apache pyspark cluster for parallel processing. EMR or Glue are possible options. Glue is serverless so better use this plus pyspark is in memory parallel processing.

upvoted 1 times

  **aragornfsm** 11 months, 1 week ago

i think a is correct

semistructured data ==> Athena


upvoted 1 times

  **pentium75** 10 months, 1 week ago

"Hadoop [as used by EMR] helps you turn petabytes of un-structured or semi-structured data into useful insights"

<https://aws.amazon.com/emr/features/hadoop/>

upvoted 2 times

  **riyasara** 11 months, 2 weeks ago

Athena is not designed for parallel data processing. So it's B

upvoted 3 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: A**

Answer is A

upvoted 1 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: B**

From this documentation looks like EMR cannot interface with S3.

<https://aws.amazon.com/emr/>

I will settle with option A.

upvoted 2 times

  **pentium75** 10 months, 1 week ago

Of course EMR can access S3

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-file-systems.html>



upvoted 2 times

  **bogobob** 11 months, 3 weeks ago

**Selected Answer: B**

For those answering A, AWS Glue can directly query S3, it can't use Athena as a source of data. The questions say the Redshift data should be user to "enrich" which means thats the redshift data needs to be "added" to the s3 data. A doesn't allow that.

upvoted 2 times

  **hungta** 11 months, 3 weeks ago

**Selected Answer: B**

Choose option B.

Option A is not correct. Amazon Athena is suitable for querying data directly from S3 using SQL and allows parallel processing of S3 data.

AWS Glue can be used for data preparation and enrichment but might not directly integrate with Amazon Redshift for enrichment.


upvoted 2 times

  **potomac** 12 months ago

**Selected Answer: A**

Athena and Redshift both do SQL query

upvoted 1 times

  **Sab123** 1 year ago

**Selected Answer: A**

semi-structure supported by Athena not by EMR

upvoted 4 times

  **pentium75** 10 months, 1 week ago

"Hadoop helps you turn petabytes of un-structured or semi-structured data into useful insights about your applications or users."

[https://aws.amazon.com/emr/features/hadoop/?nc1=h\\_ls](https://aws.amazon.com/emr/features/hadoop/?nc1=h_ls)

upvoted 2 times

  **JKevin778** 1 year, 1 month ago

**Selected Answer: A**

athena for s3

upvoted 1 times

A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month.

What is the MOST cost-effective solution to connect these VPCs?

- A. Implement AWS Transit Gateway to connect the VPCs. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
- B. Implement an AWS Site-to-Site VPN tunnel between the VPCs. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
- C. Set up a VPC peering connection between the VPCs. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
- D. Set up a 1 GB AWS Direct Connect connection between the VPCs. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Correct Answer: C

Community vote distribution

C (100%)

  **luiscc** Highly Voted 9 months ago

**Selected Answer: C**

C is the correct answer.

VPC peering is the most cost-effective way to connect two VPCs within the same region and AWS account. There are no additional charges for VPC peering beyond standard data transfer rates.

Transit Gateway and VPN add additional hourly and data processing charges that are not necessary for simple VPC peering.

Direct Connect provides dedicated network connectivity, but is overkill for the relatively low inter-VPC data transfer needs described here. It has high fixed costs plus data transfer rates.

For occasional inter-VPC communication of moderate data volumes within the same region and account, VPC peering is the most cost-effective solution. It provides simple private connectivity without transfer charges or network appliances.

upvoted 5 times

  **Guru4Cloud** Highly Voted 8 months, 2 weeks ago

**Selected Answer: C**

The key reasons are:

VPC peering provides private connectivity between VPCs without using public IP space.

Data transferred between peered VPCs is free as long as they are in the same region.

500 GB/month inter-VPC data transfer fits within peering free tier.

Transit Gateway (Option A) incurs hourly charges plus data transfer fees. More costly than peering.

Site-to-Site VPN (Option B) incurs hourly charges and data transfer fees. More expensive than peering.

Direct Connect (Option D) has high hourly charges and would be overkill for this use case.

upvoted 5 times

  **TariqKipkemei** Most Recent 5 months, 2 weeks ago

**Selected Answer: C**

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account. The VPCs can be in different Regions (also known as an inter-Region VPC peering connection).

<https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html#:~:text=A-,VPC%20peering,-connection%20is%20a>

upvoted 3 times

  **BrijMohan08** 8 months, 1 week ago

**Selected Answer: C**

Transit Gateway network peering.

VPC Peering to peer 2 or more VPC in the same region.

upvoted 4 times

  **mrsoa** 9 months ago

**Selected Answer: C**

VPC peering is the most cost-effective solution  
upvoted 2 times

  **Deepakin96** 9 months ago

**Selected Answer: C**

Communicating with two VPC in same account = VPC Peering  
upvoted 2 times

A company hosts multiple applications on AWS for different product lines. The applications use different compute resources, including Amazon EC2 instances and Application Load Balancers. The applications run in different AWS accounts under the same organization in AWS Organizations across multiple AWS Regions. Teams for each product line have tagged each compute resource in the individual accounts.

The company wants more details about the cost for each product line from the consolidated billing feature in Organizations.



Which combination of steps will meet these requirements? (Choose two.)


- A. Select a specific AWS generated tag in the AWS Billing console.
- B. Select a specific user-defined tag in the AWS Billing console.
- C. Select a specific user-defined tag in the AWS Resource Groups console.
- D. Activate the selected tag from each AWS account.
- E. Activate the selected tag from the Organizations management account.

**Correct Answer:** *BE*

*Community vote distribution*

BE (100%)

-   **Guru4Cloud**



Highly Voted 

 1 year, 8 months ago



Selected Answer: BE


The reasons are:

User-defined tags were created by each product team to identify resources. Selecting the relevant tag in the Billing console will group costs. The tag must be activated from the Organizations management account to consolidate billing across all accounts. AWS generated tags are predefined by AWS and won't align to product lines. Resource Groups (Option C) helps manage resources but not billing. Activating the tag from each account (Option D) is not needed since Organizations centralizes billing.

upvoted 11 times
-   **adiorsdior999** 6 months ago

Resource Groups helps grouping, editing, searching tag for resources, and do not use for viewing cost details



upvoted 2 times
-   **potomac**

Most Recent 

 1 year, 5 months ago



Selected Answer: BE

Your user-defined cost allocation tags represent the tag key, which you activate in the Billing console.

upvoted 2 times
-   **mrsoa** 1 year, 9 months ago

Selected Answer: BE

BE BE BE BE

upvoted 3 times
-   **Kiki\_Pass** 1 year, 9 months ago

Selected Answer: BE

"Only a management account in an organization and single accounts that aren't members of an organization have access to the cost allocation tags manager in the Billing and Cost Management console."

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html>

upvoted 4 times

A company's solutions architect is designing an AWS multi-account solution that uses AWS Organizations. The solutions architect has organized the company's accounts into organizational units (OUs).

The solutions architect needs a solution that will identify any changes to the OU hierarchy. The solution also needs to notify the company's operations team of any changes.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Provision the AWS accounts by using AWS Control Tower. Use account drift notifications to identify the changes to the OU hierarchy.
- B. Provision the AWS accounts by using AWS Control Tower. Use AWS Config aggregated rules to identify the changes to the OU hierarchy.
- C. Use AWS Service Catalog to create accounts in Organizations. Use an AWS CloudTrail organization trail to identify the changes to the OU hierarchy.
- D. Use AWS CloudFormation templates to create accounts in Organizations. Use the drift detection operation on a stack to identify the changes to the OU hierarchy.

**Correct Answer: A**

*Community vote distribution*

A (70%)

B (20%)

10%

 **Guru4Cloud** Highly Voted 1 year, 8 months ago

**Selected Answer: A**

The key advantages you highlight of Control Tower are convincing:

Fully managed service simplifies multi-account setup.  
Built-in account drift notifications detect OU changes automatically.  
More scalable and less complex than Config rules or CloudTrail.  
Better security and compliance guardrails than custom options.  
Lower operational overhead compared to other solution  
upvoted 12 times

 **Bmaster** Highly Voted 1 year, 9 months ago

A is correct.

<https://docs.aws.amazon.com/controltower/latest/userguide/what-is-control-tower.html>  
<https://docs.aws.amazon.com/controltower/latest/userguide/prevention-and-notification.html>  
upvoted 8 times

 **1166ae3** Most Recent 10 months ago

**Selected Answer: C**

Create Accounts using AWS Service Catalog:

Utilize AWS Service Catalog to provision AWS accounts within AWS Organizations. This ensures standardized account creation and management.  
Enable AWS CloudTrail Organization Trail:

Set up an AWS CloudTrail organization trail that records all API calls across all accounts in the organization.  
This trail will capture changes to the OU hierarchy, including any modifications to organizational units.  
upvoted 2 times

 **chickenmf** 1 year, 1 month ago

**Selected Answer: B**

AWS Config helps you maintain a detailed inventory of your resources and their configurations, track changes over time, and ensure compliance with your organization's policies and industry regulations.  
upvoted 4 times

 **chickenmf** 1 year, 1 month ago


Furthermore, AWS Config Aggregated Rules are a feature within AWS Config that enables you to evaluate compliance with desired configurations or compliance standards across multiple AWS accounts and regions. They are particularly useful in scenarios where you want to enforce consistent rules and compliance checks across an entire organization with multiple AWS accounts.  
upvoted 1 times

 **chickenmf** 1 year, 1 month ago

NVM - This is such a stupid question lol changing my answer to A due to the following:  
Account drift notifications in AWS are a feature provided by AWS Control Tower. These notifications help organizations identify and respond

to changes made to an AWS account that deviate from the established baseline configuration created during the initial setup by AWS Control Tower. Drift refers to any configuration changes that have been made to an AWS account after it was provisioned by Control Tower.

upvoted 5 times

  **Avyay** 1 year, 1 month ago

This was in my exam today..I selected Answer A

upvoted 3 times

  **chickenmf** 1 year, 1 month ago

what percentage of all these questions would you say were in the exam?

upvoted 2 times

  **wizcloudifa** 1 year ago

I read in one of the earlier questions, its around 75%, someone who gave the exam said so

upvoted 2 times

  **SHAAHIBHUSHANAWS** 1 year, 5 months ago

A

<https://docs.aws.amazon.com/controltower/latest/userguide/drift.html>



upvoted 2 times

  **potomac** 1 year, 5 months ago

**Selected Answer: A**

AWS Control Tower provides passive and active methods of drift monitoring protection for preventive controls.

upvoted 2 times

  **darekw** 1 year, 8 months ago

<https://docs.aws.amazon.com/controltower/latest/userguide/prevention-and-notification.html>

upvoted 2 times

A company's website handles millions of requests each day, and the number of requests continues to increase. A solutions architect needs to improve the response time of the web application. The solutions architect determines that the application needs to decrease latency when retrieving product details from the Amazon DynamoDB table.

Which solution will meet these requirements with the LEAST amount of operational overhead?

- A. Set up a DynamoDB Accelerator (DAX) cluster. Route all read requests through DAX.
- B. Set up Amazon ElastiCache for Redis between the DynamoDB table and the web application. Route all read requests through Redis.
- C. Set up Amazon ElastiCache for Memcached between the DynamoDB table and the web application. Route all read requests through Memcached.
- D. Set up Amazon DynamoDB Streams on the table, and have AWS Lambda read from the table and populate Amazon ElastiCache. Route all read requests through ElastiCache.

**Correct Answer: A**

*Community vote distribution*

A (100%)

- mrsoa** Highly Voted

1 year, 3 months ago

Selected Answer: A

A , because B,C and D contains Elasticache which required a heavy code changes, so more operational overhead

upvoted 9 times
- TariqKipkemei** Highly Voted

11 months, 2 weeks ago

Selected Answer: A

decrease latency when retrieving product details from the Amazon DynamoDB = Amazon DynamoDB Accelerator (DAX)

upvoted 5 times
- sandordini** Most Recent

6 months, 2 weeks ago

Selected Answer: A

DAX to reduce latency

upvoted 3 times
- Guru4Cloud**

1 year, 2 months ago

Selected Answer: A

The key reasons:  
  
DAX provides a DynamoDB-compatible caching layer to reduce read latency. It is purpose-built for accelerating DynamoDB workloads. Using DAX requires minimal application changes - only read requests are routed through it. DAX handles caching logic automatically without needing complex integration code. ElastiCache Redis/Memcached (Options B/C) require more integration work to sync DynamoDB data. Using Lambda and Streams to populate ElastiCache (Option D) is a complex event-driven approach requiring ongoing maintenance. DAX plugs in seamlessly to accelerate DynamoDB with very little operational overhead

upvoted 3 times
- Deepakin96**

1 year, 3 months ago

Selected Answer: A

DynamoDB = DAX

upvoted 3 times
- Bmaster**

1 year, 3 months ago

only A

upvoted 3 times

A solutions architect needs to ensure that API calls to Amazon DynamoDB from Amazon EC2 instances in a VPC do not travel across the internet.

Which combination of steps should the solutions architect take to meet this requirement? (Choose two.)



- A. Create a route table entry for the endpoint.
- B. Create a gateway endpoint for DynamoDB.
- C. Create an interface endpoint for Amazon EC2.
- D. Create an elastic network interface for the endpoint in each of the subnets of the VPC.
- E. Create a security group entry in the endpoint's security group to provide access.


**Correct Answer:** AB

*Community vote distribution*

AB (74%)

BE (23%)



-   **ukivanlamlpi**


Highly Voted 

 1 year, 8 months ago

Selected Answer: AB

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-ddb.html>

upvoted 11 times
-   **Guru4Cloud**

Highly Voted 

 1 year, 8 months ago

Selected Answer: BE

The reasons are:



A gateway endpoint for DynamoDB enables private connectivity between DynamoDB and the VPC. This allows EC2 instances to access DynamoDB APIs without traversing the internet.

A security group entry is needed to allow the EC2 instances access to the DynamoDB endpoint over the VPC.

An interface endpoint is used for services like S3 and Systems Manager, not DynamoDB.



Route table entries route traffic within a VPC but do not affect external connectivity.


Elastic network interfaces are not needed for gateway endpoints.

upvoted 10 times
-   **unbendable** 1 year, 6 months ago

"The outbound rules for the security group for instances that access DynamoDB through the gateway endpoint must allow traffic to DynamoDB", <https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-ddb.html>

The option however is talking about the security group of the endpoint

upvoted 2 times
-   **JoeTromundo**

Most Recent 

 7 months, 1 week ago

Selected Answer: AB



A & B are correct

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

E is incorrect. There's no need for security group.



From the URL above:

"Once the VPC subnet's gateway endpoint has been granted access to DynamoDB, any AWS account with access to that subnet can use DynamoDB."



upvoted 2 times
-   **a7md0** 10 months ago

Selected Answer: AB

Creating the gateway endpoint and edit the route table is enough, there are no security group involved



upvoted 2 times
-   **osmk** 1 year, 3 months ago

AB <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

upvoted 3 times
-   **upliftinghut** 1 year, 3 months ago

Selected Answer: AB

C & D are both not relevant. D looks ok but DynamoDB doesn't go with security group, it only allows route table for VPC endpoint. Link here: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

upvoted 2 times
-   **upliftinghut** 1 year, 3 months ago

Sorry, E not D. E looks ok but DynamoDB doesn't go with security group, it only allows route table for VPC endpoint. Link here: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: AB**

DynamoDB can only be connected via Gateway endpoint (just like S3)  
route table for connecting the VPC tor the endpoint  
So do B then A

C: interface endpoint for EC2 to what?  
D: ENI not applicable here for VPC  
E: Incomplete option as to access to what?  
upvoted 4 times



  **theonlyhero** 1 year, 3 months ago

go through this video it will show the answer is AB  
<https://www.youtube.com/watch?v=8FTnyhklEvU>  
upvoted 5 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: AB**

Gateway Endpoint does not have an ENI, thus it has no security group. Instances have security groups and those must allow access to DynamoDB.  
upvoted 6 times

  **aws94** 1 year, 4 months ago

**Selected Answer: BE**

A. Create a route table entry for the endpoint: This is not necessary, as the gateway endpoint itself automatically creates the required route table entries.  
upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: AB**

Create a gateway endpoint for DynamoDB then create a route table entry for the endpoint  
upvoted 3 times

  **EdenWang** 1 year, 5 months ago

**Selected Answer: BE**

refer to question 555  
upvoted 2 times

  **cciesam** 1 year, 5 months ago

**Selected Answer: AB**

<https://docs.aws.amazon.com/vpc/latest/privatelink/gateway-endpoints.html#vpc-endpoints-routing>  
Traffic from your VPC to Amazon S3 or DynamoDB is routed to the gateway endpoint. Each subnet route table must have a route that sends traffic destined for the service to the gateway endpoint using the prefix list for the service.  
upvoted 2 times

  **potomac** 1 year, 5 months ago

**Selected Answer: AB**

You can access Amazon DynamoDB from your VPC using gateway VPC endpoints. After you create the gateway endpoint, you can add it as a target in your route table for traffic destined from your VPC to DynamoDB.  
upvoted 3 times

  **danielmakita** 1 year, 6 months ago

It is A and B. Not E because security group does not span VPCs.  
upvoted 3 times

  **iwannabeawsgod** 1 year, 6 months ago

**Selected Answer: AB**

A and B for sure  
upvoted 4 times

  **loveaws** 1 year, 7 months ago

B and D.  
upvoted 1 times

A company runs its applications on both Amazon Elastic Kubernetes Service (Amazon EKS) clusters and on-premises Kubernetes clusters. The company wants to view all clusters and workloads from a central location.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon CloudWatch Container Insights to collect and group the cluster information.
- B. Use Amazon EKS Connector to register and connect all Kubernetes clusters.
- C. Use AWS Systems Manager to collect and view the cluster information.
- D. Use Amazon EKS Anywhere as the primary cluster to view the other clusters with native Kubernetes commands.

Correct Answer: B

Community vote distribution

B (96%)

4%

  **ErnShm** Highly Voted 1 year, 8 months ago  
B

You can use Amazon EKS Connector to register and connect any conformant Kubernetes cluster to AWS and visualize it in the Amazon EKS console. After a cluster is connected, you can see the status, configuration, and workloads for that cluster in the Amazon EKS console. You can use this feature to view connected clusters in Amazon EKS console, but you can't manage them. The Amazon EKS Connector requires an agent that is an open source project on Github. For additional technical content, including frequently asked questions and troubleshooting, see Troubleshooting issues in Amazon EKS Connector

The Amazon EKS Connector can connect the following types of Kubernetes clusters to Amazon EKS.

On-premises Kubernetes clusters

Self-managed clusters that are running on Amazon EC2

Managed clusters from other cloud providers  
upvoted 7 times



  **mrsoa** Highly Voted 1 year, 9 months ago  
seems B

<https://docs.aws.amazon.com/eks/latest/userguide/eks-connector.html>  
upvoted 5 times



  **Danilus** Most Recent 5 months, 3 weeks ago  
**Selected Answer: B**



KEY-he company wants to view all clusters and workloads from a central location.

the answer is not D because amazon EKS anywhere allow you to run cluster of kubernetes in on-premises or in other clouds  
the answer is B EKS connector helps u to connect kubernetes clusters on premises or other cloud and manage from a single ubication  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago  
**Selected Answer: B**

<https://docs.aws.amazon.com/eks/latest/userguide/eks-connector.html>  
"You can use Amazon EKS Connector to register and connect any conformant Kubernetes cluster to AWS and visualize it in the Amazon EKS console. "  
B is the right product for this.  
upvoted 4 times

  **pentium75** 1 year, 4 months ago  
**Selected Answer: B**  
EKS Connector -> 'view clusters and workloads' as requested  
EKS Anywhere -> create and manage on-premises EKS clusters  
upvoted 4 times

  **SHAAHIBHUSHANAWS** 1 year, 5 months ago  
B  
EKS connector helps to integrate multiple cluster with EKS console. EKS anywhere is Kubernetes Ditro cluster to be deployed on-prem. It is not for integrating with other cluster.  
upvoted 1 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: B**

View all clusters and workloads (incl on-prem) from a central location = Amazon EKS Connector  
Create and operate Kubernetes clusters on your own infrastructure = Amazon EKS Anywhere

<https://aws.amazon.com/eks/eks-anywhere/#:~:text=Amazon-,EKS%20Anywhere,-lets%20you%20create>

<https://docs.aws.amazon.com/eks/latest/userguide/eks-connector.html#:~:text=You%20can%20use-,Amazon%20EKS%20Connector,-to%20register%20and>

upvoted 2 times

  **potomac** 1 year, 5 months ago

**Selected Answer: B**

It is B

upvoted 2 times

  **thainguyensunya** 1 year, 8 months ago

**Selected Answer: B**

Definitely B.

"You can use Amazon EKS Connector to register and connect any conformant Kubernetes cluster to AWS and visualize it in the Amazon EKS console. After a cluster is connected, you can see the status, configuration, and workloads for that cluster in the Amazon EKS console. "

<https://docs.aws.amazon.com/eks/latest/userguide/eks-connector.html>

upvoted 3 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

The key reasons:

EKS Connector allows registering external Kubernetes clusters (on-premises and otherwise) with Amazon EKS

This provides a unified view and management of all clusters within the EKS console.

EKS Connector handles keeping resources in sync across connected clusters.

This centralized approach minimizes operational overhead compared to using separate tools.

CloudWatch Container Insights (Option A) only provides metrics and logs, not cluster management.

Systems Manager (Option C) is more general purpose and does not natively integrate with EKS.



EKS Anywhere (Option D) would not provide a single pane of glass for external clusters.

upvoted 4 times

  **RealMarcus** 1 year, 8 months ago

Amazon EKS Connector enables you to create and manage a centralized view of all your Kubernetes clusters, regardless of whether they are Amazon EKS clusters or on-premises Kubernetes clusters. It allows you to register these clusters with your Amazon EKS control plane, providing a unified management interface for all clusters.

upvoted 2 times

  **avkya** 1 year, 8 months ago

**Selected Answer: B**

You can use Amazon EKS Connector to register and connect any conformant Kubernetes cluster to AWS and visualize it in the Amazon EKS console.

After a cluster is connected, you can see the status, configuration, and workloads for that cluster in the Amazon EKS console. You can use this feature to view connected clusters in Amazon EKS console, but you can't manage them

upvoted 1 times

  **ukivanlamlpi** 1 year, 8 months ago

**Selected Answer: D**

only D can connect to on-perm

upvoted 1 times

  **pentium75** 1 year, 4 months ago

No.

"The Amazon EKS Connector can connect the following types of Kubernetes clusters to Amazon EKS.

On-premises Kubernetes clusters"


<https://aws.amazon.com/de/eks/eks-anywhere/>

upvoted 2 times

  **pentium75** 1 year, 4 months ago

Wrong link, statement is from <https://docs.aws.amazon.com/eks/latest/userguide/eks-connector.html>

upvoted 2 times

  **Bmaster** 1 year, 9 months ago

Only B

<https://docs.aws.amazon.com/eks/latest/userguide/eks-connector.html>

upvoted 3 times



A company is building an ecommerce application and needs to store sensitive customer information. The company needs to give customers the ability to complete purchase transactions on the website. The company also needs to ensure that sensitive customer data is protected, even from database administrators.

Which solution meets these requirements?

- A. Store sensitive data in an Amazon Elastic Block Store (Amazon EBS) volume. Use EBS encryption to encrypt the data. Use an IAM instance role to restrict access.
- B. Store sensitive data in Amazon RDS for MySQL. Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.
- C. Store sensitive data in Amazon S3. Use AWS Key Management Service (AWS KMS) server-side encryption to encrypt the data. Use S3 bucket policies to restrict access.
- D. Store sensitive data in Amazon FSx for Windows Server. Mount the file share on application servers. Use Windows file permissions to restrict access.

Correct Answer: B

Community vote distribution

B (100%)

 **Guru4Cloud** Highly Voted 1 year, 2 months ago

Selected Answer: B


The key reasons are:

RDS MySQL provides a fully managed database service well suited for an ecommerce application. AWS KMS client-side encryption allows encrypting sensitive data before it hits the database. The data remains encrypted at rest. This protects sensitive customer data from database admins and privileged users. EBS encryption (Option A) protects data at rest but not in use. IAM roles don't prevent admin access. S3 (Option C) encrypts data at rest on the server side. Bucket policies don't restrict admin access. FSx file permissions (Option D) don't prevent admin access to unencrypted data.  
upvoted 12 times

 **pentium75** Highly Voted 10 months, 1 week ago

Selected Answer: B

A, C and D would allow the administrator of the storage to access the data. Besides, it is data about "purchase transactions" which is usually stored in a transactional database (such as RDS for MySQL), not in a file or object storage.  
upvoted 8 times

 **sidharthwader** 5 months, 4 weeks ago  
Good thought purchase transactions I had missed that part  
upvoted 2 times


 **SHAAHIBHUSHANAWS** Most Recent 11 months ago

B  
I want to go with B as question is for database administrator. Also client key encryption is possible in code and KMS can be used for encryption but not using KMS keys. Encrypted data available in DB is of no use to DB admin.  
upvoted 2 times

 **riyasara** 11 months, 2 weeks ago



Answer is option C. option B is not ideal because Amazon RDS for MySQL is a relational database service that is optimized for structured data, not for storing sensitive customer information. Moreover, by using client-side encryption with AWS KMS, you need to encrypt and decrypt the data in your application code, which increases the risk of exposing your data in transit or at rest. You also need to manage the encryption keys yourself, which adds complexity and overhead to your application.  
upvoted 2 times

 **awsgeek75** 9 months, 3 weeks ago  
eCommerce data and transaction data are ideal for RDS which, when encrypted, is secure even from the DBA.  
upvoted 2 times

 **pentium75** 10 months, 1 week ago  
"optimized for structured data, not for storing sensitive customer information" ... Data related to "purchase transactions" is usually structured; that it contains "sensitive customer information" doesn't change the structured nature.  
upvoted 4 times

 **wsdasdasdqwaw** 1 year ago

I would go for B, because RDS (database admins), but I would like to see as well encryption at rest as well, not only in transit.  
upvoted 2 times

  **mrsoa** 1 year, 2 months ago



**Selected Answer: B**

Using client-side encryption we can protect specific fields and guarantee only decryption if the client has access to an API key, we can protect specific fields even from database admins  
upvoted 3 times

  **D10SJoker** 1 year, 2 months ago

**Selected Answer: B**

For me it's B because of "client-side encryption to encrypt the data"  
upvoted 2 times



  **h8er** 1 year, 2 months ago

keyword - database administrators  
upvoted 5 times

  **Kiki\_Pass** 1 year, 2 months ago

**Selected Answer: B**

"even from database administrators" -> "Client Side encryption"  
upvoted 4 times

  **Bmaster** 1 year, 3 months ago

My choice is B  
upvoted 4 times

A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.

Which migration solution will meet these requirements?

- A. Use native MySQL tools to migrate the database to Amazon RDS for MySQL. Configure elastic storage scaling.
- B. Migrate the database to Amazon Redshift by using the mysqldump utility. Turn on Auto Scaling for the Amazon Redshift cluster.
- C. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Aurora. Turn on Aurora Auto Scaling.
- D. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoDB. Configure an Auto Scaling policy.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Guru4Cloud** Highly Voted  1 year, 2 months ago

**Selected Answer: C**

The key reasons are:

DMS provides an easy migration path from MySQL to Aurora while minimizing downtime.

Aurora is a MySQL-compatible relational database service that will maintain compatibility with the company's applications.

Aurora Auto Scaling allows the database to automatically scale up and down based on demand to handle increased workloads.

RDS MySQL (Option A) does not scale as well as the Aurora architecture.

Redshift (Option B) is for analytics, not transactional data, and may not be compatible.

DynamoDB (Option D) is a NoSQL datastore and lacks MySQL compatibility.

upvoted 9 times

  **awsgeek75** Most Recent  9 months, 3 weeks ago

**Selected Answer: C**

A is wrong as you cannot use native MySQL tools for migration. Happy to be corrected though!

B Redshift is not compatible with MySQL

D is DynamoDB

C Aurora MySQL is compatible and supports auto scaling

upvoted 3 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: C**

on-premises MySQL database, transactional data, maintain compatibility, scale automatically = Amazon Aurora

migrating the database to the AWS Cloud = AWS Database Migration Service

upvoted 2 times

  **potomac** 12 months ago

**Selected Answer: C**

Aurora is a MySQL-compatible relational database service

upvoted 2 times

  **mrsoa** 1 year, 3 months ago

**Selected Answer: C**

Aurora is better in autoscaling then RDS

upvoted 2 times

  **Bmaster** 1 year, 3 months ago

C is correct

A is incorrect. RDS for MySQL does not scale automatically during periods of increased demand.

B is incorrect. Redshift is used for data sharing purposes.

D is incorrect. you muse change application codes.

upvoted 2 times

  **Eminenza22** 1 year, 3 months ago

Amazon RDS now supports Storage Auto Scaling

upvoted 2 times

A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 bucket. Allow access from all the EC2 instances in the VPC.
- B. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system from each EC2 instance.
- C. Create a file system on a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume. Attach the EBS volume to all the EC2 instances.
- D. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance. Synchronize the EBS volumes across the different EC2 instances.

**Correct Answer: B**

*Community vote distribution*

B (96%)

4%

- Josantru**

Highly Voted

1 year, 9 months ago

Correct B.

How is Amazon EFS different than Amazon S3?  
Amazon EFS provides shared access to data using a traditional file sharing permissions model and hierarchical directory structure via the NFSv4 protocol. Applications that access data using a standard file system interface provided through the operating system can use Amazon EFS to take advantage of the scalability and reliability of file storage in the cloud without writing any new code or adjusting applications.

Amazon S3 is an object storage platform that uses a simple API for storing and accessing data. Applications that do not require a file system structure and are designed to work with object storage can use Amazon S3 as a massively scalable, durable, low-cost object storage solution.

upvoted 13 times
- Gape4**

Most Recent

10 months, 1 week ago

**Selected Answer: B**

<https://aws.amazon.com/efs/when-to-choose-efs/>

upvoted 2 times
- zinabu**

1 year ago

hierarchical structure =EFS

upvoted 3 times
- pentium75**

1 year, 4 months ago

**Selected Answer: B**

Not A because S3 does not allow a "hierarchical directory structure"  
Not C because Multi-attach does not work "across two Availability Zones"  
Not D because we need "shared", not synchronized, storage.

upvoted 4 times
- TariqKipkemei**

1 year, 5 months ago

**Selected Answer: B**

hierarchical directory structure, read and write rapidly and concurrently to shared storage = Amazon Elastic File System

upvoted 2 times
- zits88**

9 months ago

Best and most concise explanation here. All true.

upvoted 1 times
- potomac**

1 year, 5 months ago

**Selected Answer: B**

Amazon EFS simultaneously supports on-premises servers using a traditional file permissions model, file locking, and hierarchical directory structure through the NFS v4 protocol.

upvoted 2 times
- Guru4Cloud**

1 year, 8 months ago

**Selected Answer: B**

The key reasons:

EFS provides a scalable, high performance NFS file system that can be concurrently accessed from multiple EC2 instances. It supports the hierarchical directory structure needed by the applications.  
EFS is elastic, growing and shrinking automatically as needed.  
It can be accessed from instances across AZs, meeting the shared storage requirement.  
S3 object storage (option A) lacks the file system semantics needed by the apps.  
EBS volumes (options C and D) are attached to a single instance and would require replication and syncing to share across instances.  
EFS is purpose-built for this use case of a shared file system across Linux instances and aligns best with the performance, concurrency, and availability needs.

upvoted 4 times

  **barracouto** 1 year, 8 months ago

**Selected Answer: B**

Going with b  
upvoted 2 times

  **Bennyboy789** 1 year, 8 months ago

**Selected Answer: B**

C and D involve using Amazon EBS volumes, which are block storage. While they can be attached to EC2 instances, they might not provide the same level of shared concurrent access as Amazon EFS. Additionally, synchronizing EBS volumes across different EC2 instances (as in option D) can be complex and error-prone.

Therefore, for a scenario where multiple EC2 instances need to rapidly and concurrently access shared storage with a hierarchical directory structure, Amazon EFS is the best solution.

upvoted 3 times

  **ukivanlamlpi** 1 year, 8 months ago

**Selected Answer: B**


s3 is flat structure. EBS multi mount only for same available zone  
upvoted 2 times

  **Dana12345** 1 year, 9 months ago

**Selected Answer: B**

Because Amazon EBS Multi-Attach enables you to attach a single Provisioned IOPS SSD (io1 or io2) volume to multiple instances that are in the same Availability Zone. The infra contains 2 AZ's.



upvoted 2 times

  **mrsoa** 1 year, 9 months ago

**Selected Answer: B**

B is the correct answer

<https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html>  
upvoted 2 times

  **mrsoa** 1 year, 9 months ago

B is the correct answer

<https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html>  
upvoted 2 times

  **RazSteel** 1 year, 9 months ago

**Selected Answer: C**

I think that C is the best option coz io2 can share storage and multi attach.  
upvoted 1 times

  **PLN6302** 1 year, 8 months ago

hierarchial directory structure is present in EFS  
upvoted 3 times

  **pentium75** 1 year, 4 months ago

Multi-attach does not work "across two Availability Zones".  
upvoted 3 times

A solutions architect is designing a workload that will store hourly energy consumption by business tenants in a building. The sensors will feed a database through HTTP requests that will add up usage for each tenant. The solutions architect must use managed services when possible. The workload will receive more features in the future as the solutions architect adds independent components.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon API Gateway with AWS Lambda functions to receive the data from the sensors, process the data, and store the data in an Amazon DynamoDB table.
- B. Use an Elastic Load Balancer that is supported by an Auto Scaling group of Amazon EC2 instances to receive and process the data from the sensors. Use an Amazon S3 bucket to store the processed data.
- C. Use Amazon API Gateway with AWS Lambda functions to receive the data from the sensors, process the data, and store the data in a Microsoft SQL Server Express database on an Amazon EC2 instance.
- D. Use an Elastic Load Balancer that is supported by an Auto Scaling group of Amazon EC2 instances to receive and process the data from the sensors. Use an Amazon Elastic File System (Amazon EFS) shared file system to store the processed data.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **Guru4Cloud** Highly Voted  1 year, 8 months ago

**Selected Answer: A**

The key reasons are:

- ° API Gateway removes the need to manage servers to receive the HTTP requests from sensors
  - ° Lambda functions provide a serverless compute layer to process data as needed
  - ° DynamoDB is a fully managed NoSQL database that scales automatically
  - ° This serverless architecture has minimal operational overhead to manage
  - ° Options B, C, and D all require managing EC2 instances which increases ops workload
  - ° Option C also adds SQL Server admin tasks and licensing costs
  - ° Option D uses EFS file storage which requires capacity planning and management
- upvoted 6 times

  **mrsoa** Highly Voted  1 year, 9 months ago

**Selected Answer: A**

A seems to be the right answer

upvoted 5 times

  **AWSSURI** Most Recent  8 months, 1 week ago

Options B,C,D involves unwanted operational overheads due to EC2

So A is the right answer

upvoted 2 times

  **Mikado211** 1 year, 4 months ago

Thinking of that, there is not many questions about IoT Core, but this product could be an excellent answer for the need.

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

Workload runs every hour, must use managed services, more features in the future, LEAST operational overhead = AWS Lambda functions.

HTTP requests, must use managed services, more features in the future, LEAST operational overhead = API Gateway.

Must use managed services, more features in the future, LEAST operational overhead =Amazon DynamoDB.

upvoted 4 times

  **ersin13** 1 year, 9 months ago

key word is "must use managed services when possible" api ,lambda dynamodb are serverless. so answer is A

upvoted 2 times

  **Kiki\_Pass** 1 year, 9 months ago

**Selected Answer: A**

"The workload will receive more features in the future ..." -> DynamoDB

upvoted 4 times

  **Bmaster** 1 year, 9 months ago

A is correct.

upvoted 3 times

A solutions architect is designing the storage architecture for a new web application used for storing and viewing engineering drawings. All application components will be deployed on the AWS infrastructure.

The application design must support caching to minimize the amount of time that users wait for the engineering drawings to load. The application must be able to store petabytes of data.

Which combination of storage and caching should the solutions architect use?

- A. Amazon S3 with Amazon CloudFront
- B. Amazon S3 Glacier with Amazon ElastiCache
- C. Amazon Elastic Block Store (Amazon EBS) volumes with Amazon CloudFront
- D. AWS Storage Gateway with Amazon ElastiCache

**Correct Answer: A**

*Community vote distribution*

A (100%)

- Guru4Cloud**

Highly Voted

1 year, 8 months ago

Selected Answer: A

The key reasons are:

S3 provides highly durable and scalable object storage capable of handling petabytes of data cost-effectively. CloudFront can be used to cache S3 content at the edge, minimizing latency for users and speeding up access to the engineering drawings. The global CloudFront edge network is ideal for caching large amounts of static media like drawings. EBS provides block storage but lacks the scale and durability of S3 for large media files. Glacier is cheaper archival storage but has higher latency unsuited for frequent access. Storage Gateway and ElastiCache may play a role but do not align as well to the main requirements.

upvoted 6 times
- Bmaster**

Highly Voted

1 year, 9 months ago

A is correct

upvoted 5 times
- Gape4**

Most Recent

10 months, 1 week ago

Selected Answer: A

I think the answer is A.

upvoted 2 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: A

Petabyte data on AWS infra with high performance

B is Glacier so slow

C EBS for petabyte data doesn't work

D Storage gateway is for on premise connectivity which is not required

upvoted 4 times
- TariqKipkemei**

1 year, 5 months ago

Selected Answer: A

Storing and viewing engineering drawings = Amazon S3

Support caching to minimize the amount of time that users wait for the engineering drawings to load = Amazon CloudFront

upvoted 2 times
- wsdasdasdqwda**

1 year, 6 months ago

CF caching and S3 supports petabytes data

upvoted 3 times
- lemur88**

1 year, 8 months ago

Selected Answer: A

CF allows caching

upvoted 2 times
- mrsoa**

1 year, 9 months ago

**Selected Answer: A**

The answer seems A:

B : Glacier for archiving

C : i dont think EBS scale to petabytes (I am not sure about that)

D : it incorrect becasseAll application components will be deployed on the AWS infrastrucur  
upvoted 3 times

An Amazon EventBridge rule targets a third-party API. The third-party API has not received any incoming traffic. A solutions architect needs to determine whether the rule conditions are being met and if the rule's target is being invoked.

Which solution will meet these requirements?

- A. Check for metrics in Amazon CloudWatch in the namespace for AWS/Events.
- B. Review events in the Amazon Simple Queue Service (Amazon SQS) dead-letter queue.
- C. Check for the events in Amazon CloudWatch Logs.
- D. Check the trails in AWS CloudTrail for the EventBridge events.

Correct Answer: A

Community vote distribution

A (76%)

D (15%)

10%

- lemur88

Highly Voted

1 year, 2 months ago

Selected Answer: A

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-monitoring.html>

upvoted 8 times
- awsgeek75

Highly Voted

9 months, 3 weeks ago

Selected Answer: A

"EventBridge sends metrics to Amazon CloudWatch every minute for everything from the number of matched events to the number of times a target is invoked by a rule."  
from <https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-monitoring.html>  
B: SQS, irrelevant  
C: 'Check for events', this wording is confusing but could mean something in wrong context. I would have chosen C if A wasn't an option  
D: CloudTrail is for AWS resource monitoring so irrelevant

upvoted 7 times
- pentium75

Most Recent

10 months, 1 week ago

Selected Answer: A

A per <https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-monitoring.html>  
Not B because SQS is not even involved here  
Not C because EventBridge sends only metrics, not detailed logs, to CloudWatch  
Not D, many fall for CloudTrail supposedly recording "API calls", but this is about calls for the EventBridge API to AWS, not calls to 3rd party APIs by EventBridge.

upvoted 6 times
- Min\_93

10 months, 1 week ago

Selected Answer: C

Option A, "Check for metrics in Amazon CloudWatch in the namespace for AWS/Events," primarily provides aggregated metrics related to EventBridge, but it may not give detailed information about individual events or their specific content. Metrics in CloudWatch can give you an overview of how many events are being processed, but for detailed inspection of events and their conditions, checking CloudWatch Logs (option C) is more appropriate.

CloudWatch Logs allow you to see the actual event data and details, providing a more granular view that is useful for troubleshooting and understanding the specifics of why a third-party API is not receiving incoming traffic.

upvoted 1 times
- SHAAHIBHUSHANAWS

11 months ago

A

Events not generating logs in cloudwatch and cloudtrail. only metric data is available.

upvoted 2 times
- TariqKipkemei

11 months, 2 weeks ago

Selected Answer: D

CloudWatch is a monitoring service for AWS resources and applications. CloudTrail is a web service that records API activity in your AWS account. CloudWatch monitors applications and infrastructure performance in the AWS environment. CloudTrail monitors actions in the AWS environment.

upvoted 1 times
- pentium75

10 months, 1 week ago

"API activity", referring to AWS APIs. This would record if someone modifies the EventBridge configuration.

upvoted 2 times

🗨️ 👤 **EdenWang** 11 months, 3 weeks ago

**Selected Answer: C**

C should be correct, I check in AWS management concole.  
upvoted 1 times

🗨️ 👤 **potomac** 12 months ago

**Selected Answer: A**

should be A  
upvoted 2 times

🗨️ 👤 **ibu007** 1 year, 2 months ago

**Selected Answer: D**

Check the trails in AWS CloudTrail for the EventBridge events.  
upvoted 1 times

🗨️ 👤 **pentium75** 10 months, 1 week ago

I think CloudTrail captures management events (such as modifying the EventBridge configuration)  
upvoted 2 times

🗨️ 👤 **Eminenza22** 1 year, 2 months ago

**Selected Answer: C**

Amazon CloudWatch Logs is a service that collects and stores logs from Amazon Web Services (AWS) resources. These logs can be used to troubleshoot problems, monitor performance, and audit activity.  
The other options are incorrect:

Option A: CloudWatch metrics are used to track the performance of AWS resources. They are not used to store events.  
Option B: Amazon SQS dead-letter queues are used to store messages that cannot be delivered to their intended recipients. They are not used to store events.  
Option D: AWS CloudTrail is a service that records AWS API calls. It can be used to track the activity of EventBridge rules, but it does not store the events themselves.  
upvoted 2 times

🗨️ 👤 **Eminenza22** 1 year, 2 months ago

\*Errata Corrige\*  
A

EventBridge sends metrics to Amazon CloudWatch every minute for everything from the number of matched events to the number of times a target is invoked by a rule.  
<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-monitoring.html>  
upvoted 2 times

🗨️ 👤 **Eminenza22** 1 year, 2 months ago

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/CloudWatch-Events-Monitoring-CloudWatch-Metrics.html>  
upvoted 2 times

🗨️ 👤 **jayce5** 1 year, 2 months ago

**Selected Answer: D**

The answer is D:  
"CloudTrail captures API calls made by or on behalf of your AWS account from the EventBridge console and to EventBridge API operations."  
(<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-logging-monitoring.html>)  
upvoted 2 times

🗨️ 👤 **pentium75** 10 months, 1 week ago

"API calls" to AWS for managing EventBridge. Not "API calls" BY EventBridge to 3rd party APIs.  
upvoted 1 times

🗨️ 👤 **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: D**

The key reasons:  
  
AWS CloudTrail provides visibility into EventBridge operations by logging API calls made by EventBridge.  
Checking the CloudTrail trails will show the PutEvents API calls made when EventBridge rules match an event pattern.  
CloudTrail will also log the Invoke API call when the rule target is triggered.  
CloudWatch metrics and logs contain runtime performance data but not info on rule evaluation and targeting.  
SQS dead letter queues collect failed event deliveries but won't provide insights on successful invocations.  
CloudTrail is purpose-built to log operational events and API activity so it can confirm if the EventBridge rule is being evaluated and triggering the target as expected.  
upvoted 2 times

🗨️ 👤 **Eminenza22** 1 year, 2 months ago

Amazon CloudWatch Logs is a service that collects and stores logs from Amazon Web Services (AWS) resources. These logs can be used to troubleshoot problems, monitor performance, and audit activity.  
The other options are incorrect:  
Option A: CloudWatch metrics are used to track the performance of AWS resources. They are not used to store events.

Option B: Amazon SQS dead-letter queues are used to store messages that cannot be delivered to their intended recipients. They are not used to store events.

Option D: AWS CloudTrail is a service that records AWS API calls. It can be used to track the activity of EventBridge rules, but it does not store the events themselves.



upvoted 1 times

  **Bennyboy789** 1 year, 2 months ago

**Selected Answer: A**

Option A is the most appropriate solution because Amazon EventBridge publishes metrics to Amazon CloudWatch. You can find relevant metrics in the "AWS/Events" namespace, which allows you to monitor the number of events matched by the rule and the number of invocations to the rule's target.

upvoted 6 times

  **h8er** 1 year, 2 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/CloudWatch-Events-Monitoring-CloudWatch-Metrics.html>

upvoted 2 times

A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a reminder in Amazon EventBridge to scale the instances.
- B. Create an Auto Scaling group that has a scheduled action.
- C. Create an Auto Scaling group that uses manual scaling.
- D. Create an Auto Scaling group that uses automatic scaling.

**Correct Answer: B**

*Community vote distribution*

B (100%)

  **Bmaster** Highly Voted  1 year, 3 months ago

B is correct.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-scheduled-scaling.html>  
upvoted 9 times

  **wizcloudifa** Most Recent  6 months ago

**Selected Answer: B**

A - too much operation overhead, manually provisioning the instances after you receive the reminder from eventbridge  
B - right answer, as you can scale up the EC2 instances and keep them ready before large overload time  
C - too much operation overhead in manually scaling  
D - automatic scaling will scale up the instances after some duration after it has encountered the heavy workload traffic, not ideal  
upvoted 3 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: B**

runs every Friday evening = an Auto Scaling group that has a scheduled action  
upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: B**



The key reasons:

Auto Scaling scheduled actions allow defining specific dates/times to scale out or in. This can be used to scale to 6 instances every Friday evening automatically.  
Scheduled scaling removes the need for manual intervention to scale up/down for the workload.  
EventBridge reminders and manual scaling require human involvement each week adding overhead.  
Automatic scaling responds to demand and may not align perfectly to scale out every Friday without additional tuning.  
Scheduled Auto Scaling actions provide the automation needed to scale for the weekly workload without ongoing operational overhead.  
upvoted 4 times

  **Sat897** 1 year, 3 months ago

**Selected Answer: B**

Predicted period.. So schedule the instance  
upvoted 4 times

  **mrsoa** 1 year, 3 months ago



**Selected Answer: B**

B seems to be correct  
upvoted 2 times

  **Deepakin96** 1 year, 3 months ago

**Selected Answer: B**

When we know the run time is Friday, we can schedule the instance to 6  
upvoted 3 times

  **Josantru** 1 year, 3 months ago

Correct B.

upvoted 4 times

A company is creating a REST API. The company has strict requirements for the use of TLS. The company requires TLSv1.3 on the API endpoints. The company also requires a specific public third-party certificate authority (CA) to sign the TLS certificate.

Which solution will meet these requirements?

- A. Use a local machine to create a certificate that is signed by the third-party CA. Import the certificate into AWS Certificate Manager (ACM). Create an HTTP API in Amazon API Gateway with a custom domain. Configure the custom domain to use the certificate.
- B. Create a certificate in AWS Certificate Manager (ACM) that is signed by the third-party CA. Create an HTTP API in Amazon API Gateway with a custom domain. Configure the custom domain to use the certificate.
- C. Use AWS Certificate Manager (ACM) to create a certificate that is signed by the third-party CA. Import the certificate into AWS Certificate Manager (ACM). Create an AWS Lambda function with a Lambda function URL. Configure the Lambda function URL to use the certificate.
- D. Create a certificate in AWS Certificate Manager (ACM) that is signed by the third-party CA. Create an AWS Lambda function with a Lambda function URL. Configure the Lambda function URL to use the certificate.

Correct Answer: A

Community vote distribution

A (78%)

B (22%)

  **bjexamprep** Highly Voted 1 year, 8 months ago

Selected Answer: A

I don't understand why some many people vote B. In ACM, you can either request certificate from Amazon CA or import an existing certificate. There is no option in ACM that allow you to request a certificate that can be signed by third party CA.  
upvoted 21 times

  **markoniz** 1 year, 7 months ago

I fully agree  
upvoted 6 times

  **wsdasdasdqwda** 1 year, 6 months ago

Hmm AWS is saying:  
ACM certificates can be used to establish secure communications across the internet or within an internal network. You can request a publicly trusted certificate directly from ACM (an "ACM certificate") or import a publicly trusted certificate issued by a third party. Self-signed certificates are also supported. To provision your organization's internal PKI, you can issue ACM certificates signed by a private certificate authority (CA) created and managed by AWS Private CA. The CA may either reside in your account or be shared with you by a different account.

<https://docs.aws.amazon.com/acm/latest/userguide/gs.html>  
upvoted 5 times

  **pentium75** 1 year, 4 months ago

Exactly. You can "import [not create] a publicly trusted certificate issued by a third party".  
upvoted 5 times

  **luiscc** Highly Voted 1 year, 9 months ago

Selected Answer: B

AWS Certificate Manager (ACM) is a service that lets you easily provision, manage, and deploy SSL/TLS certificates for use with AWS services and your internal resources. By creating a certificate in ACM that is signed by the third-party CA, the company can meet its requirement for a specific public third-party CA to sign the TLS certificate.  
upvoted 9 times

  **pentium75** 1 year, 4 months ago

Sounds like ChatGPT answer, "creating a certificate in ACM that is signed by the third-party CA" is not possible.  
upvoted 5 times

  **emakid** Most Recent 10 months, 1 week ago

Selected Answer: A

A. Use a local machine to create a certificate that is signed by the third-party CA. Import the certificate into AWS Certificate Manager (ACM). Create an HTTP API in Amazon API Gateway with a custom domain. Configure the custom domain to use the certificate.

Reason:

Custom Certificate: Allows you to use a certificate signed by the third-party CA.  
TLSv1.3 Support: API Gateway supports TLSv1.3 for custom domains.

Configuration: You can import the third-party CA certificate into ACM and configure API Gateway to use this certificate with a custom domain.

This approach meets all the specified requirements by allowing the use of a third-party CA-signed certificate and ensuring the API endpoints use TLSv1.3.

upvoted 2 times

🗲️ 👤 **awsgeek75** 1 year, 3 months ago

**Selected Answer: A**

A is logical answer.

BCD are either misworded here or intentionally confusing. Regardless, you cannot create a cert in ACM that is signed by 3rd party CA. You can only import these certs to ACM.

upvoted 4 times

🗲️ 👤 **Shubhi\_08** 1 year, 4 months ago

**Selected Answer: A**

We can't create third party certificates in ACM.

upvoted 2 times

🗲️ 👤 **foha2012** 1 year, 4 months ago

Is this a question from the associate or professional exam ??

upvoted 2 times

🗲️ 👤 **pentium75** 1 year, 4 months ago

**Selected Answer: A**

ACM can import, but not create, 3rd party certificates. Leaves only A.

upvoted 2 times

🗲️ 👤 **maged123** 1 year, 4 months ago

**Selected Answer: A**

You have already a publicly trusted certificate issued by a third party and you just need to import it in ACM not to creat a new one. So, the correct answer is A which is the only one that importing the certificate in ACM while B, C and D are creating a new one.

upvoted 2 times

🗲️ 👤 **sparun1607** 1 year, 4 months ago

The answer must be A,

You can't create a certificate in ACM, read the below link

<https://docs.aws.amazon.com/acm/latest/userguide/setup.html>

upvoted 2 times

🗲️ 👤 **numark** 1 year, 5 months ago

Answer is A: Can I import a third-party certificate and use it with AWS services?

Yes. If you want to use a third-party certificate with Amazon CloudFront, Elastic Load Balancing, or Amazon API Gateway, you may import it into ACM using the AWS Management Console, AWS CLI, or ACM APIs. ACM does not manage the renewal process for imported certificates. You can use the AWS Management Console to monitor the expiration dates of an imported certificates and import a new third-party certificate to replace an expiring one.

upvoted 2 times

🗲️ 👤 **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

It's 22/Nov/2023 and from the console you cant create a certificate in AWS Certificate Manager (ACM) that is signed by the third-party CA. But you could obtain it externally then import it into ACM.

upvoted 2 times

🗲️ 👤 **Tshring** 1 year, 5 months ago

**Selected Answer: B**

Option B meets these requirements:

- API Gateway HTTP APIs support TLS 1.3
- ACM can import certificates signed by third-party CAs
- API Gateway provides REST APIs

upvoted 1 times

🗲️ 👤 **pentium75** 1 year, 4 months ago

"ACM can import (!) certificates signed by third-party CA", but not create (!) them as B suggests.

upvoted 2 times

🗲️ 👤 **NickGordon** 1 year, 5 months ago

**Selected Answer: A**

In ACM you can't create a cert signed by another CA. Dude, try it by yourself. There is no such option!

upvoted 2 times

🗲️ 👤 **chen0305\_099** 1 year, 8 months ago

WHY NOT A?

upvoted 1 times



**Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

Use ACM to create a certificate signed by the third-party CA. ACM integrates with external CAs.

Create an API Gateway HTTP API with a custom domain name.

Configure the custom domain to use the ACM certificate. API Gateway supports configuring custom domains with ACM certificates.

This allows serving the API over TLS using the required third-party certificate and TLS 1.3 support.

upvoted 2 times



**pentium75** 1 year, 4 months ago

"ACM integrates with external CAs." no

upvoted 2 times



**taustin2** 1 year, 8 months ago

**Selected Answer: A**

You can provide certificates for your integrated AWS services either by issuing them directly with ACM or by importing third-party certificates into the ACM management system.

upvoted 2 times



**vini15** 1 year, 8 months ago

Should be A.

We need to import third-party certificate to ACM.

upvoted 5 times

A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL-compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory.

The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Provision an Amazon DynamoDB database with default read and write capacity settings.
- B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacity unit (ACU).
- C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
- D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

**Correct Answer: C**

*Community vote distribution*

C (100%)

- Guru4Cloud**

Highly Voted

1 year, 8 months ago

Selected Answer: C

The key reasons:

Aurora Serverless v2 provides auto-scaling so the database can handle inconsistent workloads and spikes automatically without admin intervention.

It can scale down to zero when not in use to minimize costs.

The minimum 1 ACU capacity is sufficient to replace the on-prem 2 GiB database based on the info given.

Serverless capabilities reduce admin overhead for capacity management.

DynamoDB lacks MySQL compatibility and requires more hands-on management.

RDS and provisioned Aurora require manually resizing instances to scale, increasing admin overhead.

upvoted 11 times
- dkw2342**

1 year, 1 month ago

> It can scale down to zero when not in use to minimize costs.

This part is not correct. Aurora Serverless v1 was able to scale to zero.

upvoted 2 times
- kambarami**

Highly Voted

1 year, 7 months ago

the questions are hard from 500 +

upvoted 7 times
- foha2012**

1 year, 4 months ago

I dont think these are associate exam questions rather are from AWS professional exam

upvoted 3 times
- awsgeek75**

1 year, 3 months ago

Yes, I agree. I have been reading the pro questions and these are copy paste. On the bright side, it prepares you for the next step!

upvoted 7 times
- emakid**

Most Recent

10 months, 1 week ago

Selected Answer: C

C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).

Suitability: Amazon Aurora Serverless v2 is a good option for applications with variable workloads because it automatically adjusts capacity based on demand. It can handle MySQL-compatible databases and supports auto-scaling. You can set the minimum and maximum capacity based on your needs, making it highly suitable for handling unexpected workload increases with minimal administrative overhead.

upvoted 2 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: C



LEAST administrative overhead = Aurora Serverless

upvoted 3 times
- TariqKipkemei**

1 year, 5 months ago

Selected Answer: C

LEAST administrative overhead = Serverless  
upvoted 2 times



  **ibu007** 1 year, 8 months ago

**Selected Answer: C**

serverless = LEAST overhead  
upvoted 3 times



  **D10SJoker** 1 year, 9 months ago

Why not D?  
upvoted 2 times

  **wizcloudifa** 1 year ago  
no autoscaling with RDS  
upvoted 2 times

  **awsgeek75** 1 year, 3 months ago

Because "LEAST administrative overhead" is a requirement. RDS configured with mem requirements is an admin overhead  
upvoted 2 times

  **mrsoa** 1 year, 9 months ago

**Selected Answer: C**


C seems to be the right answer

Instead of provisioning and managing database servers, you specify Aurora capacity units (ACUs). Each ACU is a combination of approximately 2 gigabytes (GB) of memory, corresponding CPU, and networking. Database storage automatically scales from 10 gibibytes (GiB) to 128 tebibytes (TiB), the same as storage in a standard Aurora DB cluster

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v1.how-it-works.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.html>

upvoted 2 times

  **Bmaster** 1 year, 9 months ago

C is correct.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.how-it-works.html#aurora-serverless-v2.how-it-works.capacity>

upvoted 3 times

A company wants to use an event-driven programming model with AWS Lambda. The company wants to reduce startup latency for Lambda functions that run on Java 11. The company does not have strict latency requirements for the applications. The company wants to reduce cold starts and outlier latencies when a function scales up.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure Lambda provisioned concurrency.
- B. Increase the timeout of the Lambda functions.
- C. Increase the memory of the Lambda functions.
- D. Configure Lambda SnapStart.

**Correct Answer: D**

*Community vote distribution*

D (100%)

🗳️  **Guru4Cloud** Highly Voted 👍 1 year, 2 months ago

**Selected Answer: D**

The key reasons:

SnapStart keeps functions initialized and ready to respond quickly, eliminating cold starts.  
SnapStart is optimized for applications without aggressive latency needs, reducing costs.  
It scales automatically to match traffic spikes, eliminating outliers when scaling up.  
SnapStart is a native Lambda feature with no additional charges, keeping costs low.  
Provisioned concurrency incurs charges for always-on capacity reserved. More costly than SnapStart.  
Increasing timeout and memory do not directly improve startup performance like SnapStart.  
upvoted 15 times

🗳️  **anikety123** Highly Voted 👍 1 year, 2 months ago

**Selected Answer: D**

Both Lambda SnapStart and provisioned concurrency can reduce cold starts and outlier latencies when a function scales up. SnapStart helps you improve startup performance by up to 10x at no extra cost. Provisioned concurrency keeps functions initialized and ready to respond in double-digit milliseconds. Configuring provisioned concurrency incurs charges to your AWS account. Use provisioned concurrency if your application has strict cold start latency requirements. You can't use both SnapStart and provisioned concurrency on the same function version.  
upvoted 5 times

🗳️  **awsgeek75** Most Recent 🕒 9 months, 3 weeks ago

**Selected Answer: D**

<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html>

"Lambda SnapStart for Java can improve startup performance for latency-sensitive applications by up to 10x at no extra cost, typically with no changes to your function code."  
upvoted 5 times

🗳️  **awsgeek75** 9 months, 3 weeks ago

Also

A: Solves concurrency issues not startup

B is for execution timeout (don't think that possible if I understand the option correctly)

C Memory is not the issue here

upvoted 2 times

🗳️  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: D**

Lambda SnapStart it is.

<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html#:~:text=RSS-,Lambda%20SnapStart,-for%20Java%20can>

upvoted 2 times

🗳️  **TariqKipkemei** 11 months, 2 weeks ago

only because its a Java 11 app...if it were any other besides Java I believe Provisioned concurrency could help.

upvoted 2 times

🗳️  **potomac** 12 months ago

**Selected Answer: D**

Lambda SnapStart for Java can improve startup performance for latency-sensitive applications by up to 10x at no extra cost, typically with no changes to your function code.

<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html>  
upvoted 3 times

  **BrijMohan08** 1 year, 2 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html>  
upvoted 2 times



  **skyphilip** 1 year, 2 months ago

**Selected Answer: D**

D is correct  
Lambda SnapStart for Java can improve startup performance for latency-sensitive applications by up to 10x at no extra cost, typically with no changes to your function code. The largest contributor to startup latency (often referred to as cold start time) is the time that Lambda spends initializing the function, which includes loading the function's code, starting the runtime, and initializing the function code.

With SnapStart, Lambda initializes your function when you publish a function version. Lambda takes a Firecracker microVM snapshot of the memory and disk state of the initialized execution environment, encrypts the snapshot, and caches it for low-latency access. When you invoke the function version for the first time, and as the invocations scale up, Lambda resumes new execution environments from the cached snapshot instead of initializing them from scratch, improving startup latency.

upvoted 2 times



  **avkya** 1 year, 2 months ago

"SnapStart does not support provisioned concurrency, the arm64 architecture, Amazon Elastic File System (Amazon EFS), or ephemeral storage greater than 512 MB." The question says "The company wants to reduce cold starts" This means provisioned concurrency. I'm a little bit confused with D.

upvoted 3 times

  **Woodlawn5700** 1 year, 2 months ago

D  
<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html>  
upvoted 2 times


  **mrsoa** 1 year, 3 months ago

**Selected Answer: D**

D is the answer


Lambda SnapStart for Java can improve startup performance for latency-sensitive applications by up to 10x at no extra cost, typically with no changes to your function code. The largest contributor to startup latency (often referred to as cold start time) is the time that Lambda spends initializing the function, which includes loading the function's code, starting the runtime, and initializing the function code.

<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html>  
upvoted 3 times

  **Bmaster** 1 year, 3 months ago

D is best!!  
A is not MOST cost effectly.  
lambda snapshot is new feature for lambda.

<https://docs.aws.amazon.com/lambda/latest/dg/snapstart.html>  
upvoted 4 times

  **Bmaster** 1 year, 3 months ago

misspell.... lambda snapstart  
upvoted 2 times

  **RaksAWS** 1 year, 3 months ago

why not D  
It should work  
upvoted 3 times

A financial services company launched a new application that uses an Amazon RDS for MySQL database. The company uses the application to track stock market trends. The company needs to operate the application for only 2 hours at the end of each week. The company needs to optimize the cost of running the database.

Which solution will meet these requirements MOST cost-effectively?

- A. Migrate the existing RDS for MySQL database to an Aurora Serverless v2 MySQL database cluster.
- B. Migrate the existing RDS for MySQL database to an Aurora MySQL database cluster.
- C. Migrate the existing RDS for MySQL database to an Amazon EC2 instance that runs MySQL. Purchase an instance reservation for the EC2 instance.
- D. Migrate the existing RDS for MySQL database to an Amazon Elastic Container Service (Amazon ECS) cluster that uses MySQL container images to run tasks.

**Correct Answer: A**

*Community vote distribution*

A (92%)

8%

 **Guru4Cloud**

Highly Voted 


1 year, 8 months ago

Selected Answer: A

The key reasons are:

Aurora Serverless v2 scales compute capacity automatically based on actual usage, down to zero when not in use. This minimizes costs for intermittent usage.  
Since it only runs for 2 hours per week, the application is ideal for a serverless architecture like Aurora Serverless.  
Aurora Serverless v2 charges per second when the database is active, unlike RDS which charges hourly.  
Aurora Serverless provides higher availability than self-managed MySQL on EC2 or ECS.  
Using reserved EC2 instances or ECS still incurs charges when not in use versus the fine-grained scaling of serverless.  
Standard Aurora clusters have a minimum capacity unlike the auto-scaling serverless architecture.


upvoted 13 times


 **dkw2342**

1 year, 1 month ago

A is correct, but Aurora Serverless v2 only scales down to 0.5 ACU, not to zero.

upvoted 2 times

 **pentium75**


Highly Voted 


1 year, 4 months ago

Selected Answer: A

2 hours per week = Serverless = A. Recommended for "infrequent, intermittent, or unpredictable workloads"

upvoted 5 times

 **LeonSauveterre**

Most Recent 

4 months, 3 weeks ago

Selected Answer: A


A: Aurora Serverless v2 offers the most cost-effective and operationally efficient solution for intermittent workloads like this.

B: over-provisioned and too expensive for this workload

C: Also expensive. You gotta pay for the instance 24/7, regardless of usage.

D: Introduces complexity and operational overhead for database management tasks. Overall, ECS is never easy to handle.

upvoted 1 times

 **awsgeek75**


1 year, 3 months ago

Selected Answer: A

B is wrong because Aurora MySQL cluster will just keep on running for the rest of the week and will be costly.

C and D have too much infra bloating so costly

upvoted 2 times

 **TariqKipkemei**

1 year, 5 months ago

Selected Answer: A

Answer is A.

Here are the key distinctions:

Amazon Aurora: provides built-in security, continuous backups, serverless compute, up to 15 read replicas, automated multi-Region replication, and integrations with other AWS services.

Amazon Aurora Serverless: is an on-demand, auto-scaling configuration for Aurora where the database automatically starts up, shuts down, and

scales capacity up or down based on your application's needs.

With serverless the db will shut down when not in use.

upvoted 5 times

  **anikety123** 1 year, 8 months ago

**Selected Answer: A**

Option is A

upvoted 3 times



  **hachiri** 1 year, 8 months ago

**Selected Answer: A**

### Aurora Serverless

- Automated database instantiation and auto-scaling based on actual usage
- Good for infrequent, intermittent or unpredictable workloads
- No capacity planning needed
- Pay per second, can be more cost-effective

upvoted 3 times

  **vini15** 1 year, 8 months ago

will go with A

Amazon Aurora Serverless v2 is suitable for the most demanding, highly variable workloads. For example, your database usage might be heavy for a short period of time, followed by long periods of light activity or no activity at all.

upvoted 3 times


  **msdnpro** 1 year, 9 months ago

**Selected Answer: A**

"Amazon Aurora Serverless v2 is suitable for the most demanding, highly variable workloads. For example, your database usage might be heavy for a short period of time, followed by long periods of light activity or no activity at all. "

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.how-it-works.html>

upvoted 3 times

  **ersin13** 1 year, 9 months ago

A. Migrate the existing RDS for MySQL database to an Aurora Serverless v2 MySQL database cluster.

upvoted 2 times

  **mrsoa** 1 year, 9 months ago

**Selected Answer: B**

B seems to be the correct answer, because if we have a predictable workload Aurora database seems to be most cost effective however if we have unpredictable workload aurora serverless seems to be more cost effective because our database will scale up and down

for more informations please read this article

<https://medium.com/trackit/aurora-or-aurora-serverless-v2-which-is-more-cost-effective-bcd12e172dcf>

upvoted 3 times

  **Chef\_couincouin** 1 year, 5 months ago


according to the link, i understand that Aurora Serverless is ideal for sudden peaks in database usage with moderate or minimal usage during other periods of the day. So Answer is A

upvoted 3 times

  **Smart** 1 year, 8 months ago

True but due to autoscaling - it will be cheaper...check example#1 in the your link.

upvoted 1 times

  **Smart** 1 year, 8 months ago

Correct Answer is A

upvoted 2 times

  **pentium75** 1 year, 4 months ago

Provisioned RDS (as in B) is good for steady (not "predictable") workloads. In this case, the workload is predictable, but the prediction is that it will be used only 2 hours per week.

upvoted 3 times

  **pentium75** 1 year, 4 months ago

Aurora Serverless is for "infrequent, intermittent, OR unpredictable workloads"

upvoted 3 times

A company deploys its applications on Amazon Elastic Kubernetes Service (Amazon EKS) behind an Application Load Balancer in an AWS Region. The application needs to store data in a PostgreSQL database engine. The company wants the data in the database to be highly available. The company also needs increased capacity for read workloads.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Create an Amazon DynamoDB database table configured with global tables.
- B. Create an Amazon RDS database with Multi-AZ deployments.
- C. Create an Amazon RDS database with Multi-AZ DB cluster deployment.
- D. Create an Amazon RDS database configured with cross-Region read replicas.

**Correct Answer: C**

*Community vote distribution*

C (100%)

- Guru4Cloud**

Highly Voted

1 year, 8 months ago

Selected Answer: C

RDS Multi-AZ DB cluster deployments provide high availability, automatic failover, and increased read capacity. A multi-AZ cluster automatically handles replicating data across AZs in a single region. This maintains operational efficiency as it is natively managed by RDS without needing external replication. DynamoDB global tables involve complex provisioning and requires app changes. RDS read replicas require manual setup and management of replication. RDS Multi-AZ clustering is purpose-built by AWS for HA PostgreSQL deployments and balancing read workloads.

upvoted 9 times
- luiscc**

Highly Voted

1 year, 9 months ago

Selected Answer: C

DB cluster deployment can scale read workloads by adding read replicas. This provides increased capacity for read workloads without impacting the write workload.

upvoted 5 times
- MatAlves**

Most Recent

7 months, 3 weeks ago

Selected Answer: C

"A Multi-AZ DB cluster deployment is a semisynchronous, high availability deployment mode of Amazon RDS with two readable replica DB instances."

upvoted 2 times
- upliftinghut**

1 year, 3 months ago

Selected Answer: C

multi-AZ addresses both HA & increased read capacity with synchronous data replication between main DB & standby. Read replica is not enough because only increased read capacity not enabling HA, besides the data replication is async

upvoted 2 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/multi-az-db-clusters-concepts.html>  
"A Multi-AZ DB cluster deployment is a semisynchronous, high availability deployment mode of Amazon RDS with two readable standby DB instances"  
A: DynamoDB is not Postgres  
B: Although HA is achieve but it does not increase the read capacity as much as C without additional operational complexity  
D: Cross region is not a requirement and won't solve the same region HA or read issues

upvoted 2 times
- aws94**

1 year, 4 months ago

Selected Answer: C

Multi-AZ DB Cluster Deployment = Aurora

upvoted 2 times
- TariqKipkemei**

1 year, 5 months ago

Selected Answer: C



Multi-AZ DB cluster deployments provides two readable DB instances if you need additional read capacity

upvoted 2 times

  **potomac** 1 year, 5 months ago



**Selected Answer: C**

C is correct  
upvoted 2 times

  **avkya** 1 year, 8 months ago

**Selected Answer: C**

Multi-AZ DB clusters provide high availability, increased capacity for read workloads, and lower write latency when compared to Multi-AZ DB instance deployments.  
upvoted 2 times

  **mrsoa** 1 year, 9 months ago

**Selected Answer: C**

CCCCCCCCCCCCcCCcCCCCccccCc  
upvoted 2 times

A company is building a RESTful serverless web application on AWS by using Amazon API Gateway and AWS Lambda. The users of this web application will be geographically distributed, and the company wants to reduce the latency of API requests to these users.

Which type of endpoint should a solutions architect use to meet these requirements?

- A. Private endpoint
- B. Regional endpoint
- C. Interface VPC endpoint
- D. Edge-optimized endpoint

**Correct Answer: D**

*Community vote distribution*

D (100%)

  **mrsoa** Highly Voted 1 year, 3 months ago

**Selected Answer: D**

The correct answer is D

API Gateway - Endpoint Types

- Edge-Optimized (default): For global clients
- Requests are routed through the CloudFront Edge locations (improves latency)
- The API Gateway still lives in only one region
- Regional:
- For clients within the same region
- Could manually combine with CloudFront (more control over the caching strategies and the distribution)
- Private:
- Can only be accessed from your VPC using an interface VPC endpoint (ENI)
- Use a resource policy to define access

upvoted 8 times

  **potomac** Highly Voted 12 months ago

**Selected Answer: D**

An edge-optimized API endpoint typically routes requests to the nearest CloudFront Point of Presence (POP), which could help in cases where your clients are geographically distributed. This is the default endpoint type for API Gateway REST APIs.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-endpoint-types.html>

upvoted 5 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

**Selected Answer: D**

geographically distributed users + low latency = Edge optimized ednpoint

upvoted 2 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: D**

An edge-optimized API endpoint typically routes requests to the nearest CloudFront Point of Presence (POP), which could help in cases where your clients are geographically distributed. This is the default endpoint type for API Gateway REST APIs.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-endpoint-types.html#:~:text=API%20endpoint%20typically-,routes,-requests%20to%20the>

upvoted 3 times

  **dilaaziz** 12 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-endpoint-types.html>

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: D**

Edge-optimized endpoint

upvoted 3 times

  **Josantru** 1 year, 3 months ago

Correct D.

Edge-optimized API endpoints

An edge-optimized API endpoint is best for geographically distributed clients. API requests are routed to the nearest CloudFront Point of Presence (POP). This is the default endpoint type for API Gateway REST APIs.

upvoted 3 times

A company uses an Amazon CloudFront distribution to serve content pages for its website. The company needs to ensure that clients use a TLS certificate when accessing the company's website. The company wants to automate the creation and renewal of the TLS certificates.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Use a CloudFront security policy to create a certificate.
- B. Use a CloudFront origin access control (OAC) to create a certificate.
- C. Use AWS Certificate Manager (ACM) to create a certificate. Use DNS validation for the domain.
- D. Use AWS Certificate Manager (ACM) to create a certificate. Use email validation for the domain.

Correct Answer: C

Community vote distribution

C (100%)

Bmaster

Highly Voted

1 year, 3 months ago

C is correct.

"ACM provides managed renewal for your Amazon-issued SSL/TLS certificates. This means that ACM will either renew your certificates automatically (if you are using DNS validation), or it will send you email notices when expiration is approaching. These services are provided for both public and private ACM certificates."

<https://docs.aws.amazon.com/acm/latest/userguide/managed-renewal.html>

upvoted 10 times

Guru4Cloud

Highly Voted

1 year, 2 months ago

Selected Answer: C

The key reasons are:

AWS Certificate Manager (ACM) provides free public TLS/SSL certificates and handles certificate renewals automatically. Using DNS validation with ACM is operationally efficient since it automatically makes changes to Route 53 rather than requiring manual validation steps. ACM integrates natively with CloudFront distributions for delivering HTTPS content. CloudFront security policies and origin access controls do not issue TLS certificates. Email validation requires manual steps to approve the domain validation emails for each renewal.

upvoted 6 times

awsgeek75

Most Recent

9 months, 3 weeks ago

Selected Answer: C

For me, C is the only realistic option as I don't think you can do AB without a lot of complexity. D just makes no sense.

upvoted 2 times

ibu007

1 year, 2 months ago

Selected Answer: C

Use AWS Certificate Manager (ACM) to create a certificate. Use DNS validation for the domain

upvoted 4 times

chen0305\_099

1 year, 2 months ago

Selected Answer: C

C 似乎是正確的

upvoted 4 times

Kiki\_Pass

1 year, 2 months ago

Selected Answer: C

"DNS Validation is preferred for automation purposes" -- Stephane's course on Udemy

upvoted 3 times

mrsoa

1 year, 3 months ago

Selected Answer: C

C seems to be correct

upvoted 2 times

nananashi

1 year, 3 months ago

I think the general product uses DNS rather than email to automate, is the given answer correct?

upvoted 1 times

A company deployed a serverless application that uses Amazon DynamoDB as a database layer. The application has experienced a large increase in users. The company wants to improve database response time from milliseconds to microseconds and to cache requests to the database.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use DynamoDB Accelerator (DAX).
- B. Migrate the database to Amazon Redshift.
- C. Migrate the database to Amazon RDS.
- D. Use Amazon ElastiCache for Redis.

**Correct Answer: A**

*Community vote distribution*

A (91%)

9%

- h8er**

Highly Voted

1 year, 9 months ago

Selected Answer: A

Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for Amazon DynamoDB that delivers up to a 10 times performance improvement—from milliseconds to microseconds—even at millions of requests per second.

[https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20\(DAX\)%20is,millions%20of%20requests%20per%20second.](https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20(DAX)%20is,millions%20of%20requests%20per%20second.)

upvoted 11 times
- MatAlves**

Most Recent

7 months, 3 weeks ago

Amazon ElastiCache for Redis would help with "caching requests", but not "improve database response" itself.

upvoted 2 times
- awsgeek75**

1 year, 3 months ago

Selected Answer: A

DAX is least operations overhead.  
B: Redshift, although powerful, but is for analytics  
C: Downgrading to RDS won't help  
D: EC for Redis is more for persistent caching so would be good but lot of operational overhead

upvoted 3 times
- TariqKipkemei**

1 year, 5 months ago

Selected Answer: A

improve DynamoDB response time from milliseconds to microseconds and to cache requests to the database = DynamoDB Accelerator (DAX)

upvoted 3 times
- Guru4Cloud**

1 year, 8 months ago

Selected Answer: C

Use DynamoDB Accelerator (DAX).

upvoted 2 times
- pentium75**

1 year, 4 months ago

Which is A, not C.

upvoted 3 times
- awsgeek75**

1 year, 3 months ago

Quote A but mark C. You need more coffee mate :)

upvoted 5 times
- mrsoa**

1 year, 9 months ago

Selected Answer: A

A is the right answer

upvoted 3 times
- Bmaster**

1 year, 9 months ago

Correct A.

upvoted 2 times

A company runs an application that uses Amazon RDS for PostgreSQL. The application receives traffic only on weekdays during business hours. The company wants to optimize costs and reduce operational overhead based on this usage.

Which solution will meet these requirements?

- A. Use the Instance Scheduler on AWS to configure start and stop schedules.
- B. Turn off automatic backups. Create weekly manual snapshots of the database.
- C. Create a custom AWS Lambda function to start and stop the database based on minimum CPU utilization.
- D. Purchase All Upfront reserved DB instances.

**Correct Answer: A**

*Community vote distribution*

A (96%)

4%

  **potomac** Highly Voted 12 months ago

**Selected Answer: A**

The Instance Scheduler on AWS solution automates the starting and stopping of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Relational Database Service (Amazon RDS) instances.

This solution helps reduce operational costs by stopping resources that are not in use and starting them when they are needed. The cost savings can be significant if you leave all of your instances running at full utilization continuously.

<https://aws.amazon.com/solutions/implementations/instance-scheduler-on-aws/>

upvoted 6 times

  **pentium75** Most Recent 10 months, 1 week ago

**Selected Answer: A**

B increases operational overhead

C Lambda functions could work but NOT "based on minimum CPU utilization"

D might save cost but not as much as A

upvoted 4 times

  **ibu007** 1 year, 2 months ago

**Selected Answer: A**

A. Use the Instance Scheduler on AWS to configure start and stop schedules

upvoted 4 times

  **baba365** 1 year, 1 month ago

Why not D?

upvoted 3 times

  **AndreiWebNet** 11 months ago

How do you actually reduce costs enough to buy upfront instances that you pay for them if you use them 1 h or 24 and it is payed to run 24h. It says that you use this instances 8 hours a day 5 days a week, totaling 40h a week.... so is it the difference from 40h to 168 h?



upvoted 3 times

  **master9** 10 months, 2 weeks ago

When you buy Reserved Instances, the larger the upfront payment, the greater the discount. To maximize your savings, you can pay all up-front and receive the largest discount. Partial up-front RI's offer lower discounts but give you the option to spend less up front. Lastly, you can choose to spend nothing up front and receive a smaller discount, but allowing you to free up capital to spend in other projects.

But you need some mechanism to stop on weekend and in night to save cost.

upvoted 2 times

  **ErnShm** 1 year, 2 months ago

A

<https://docs.aws.amazon.com/solutions/latest/instance-scheduler-on-aws/solution-overview.html>

upvoted 3 times

  **Guru4Cloud** 1 year, 2 months ago

**Selected Answer: A**

Purpose-built scheduling minimizes operational overhead.

Aligns instance running time precisely with business hour demands.

Maintains backups unlike disabling auto backups.

More cost effective and flexible than reserved instances.  
Simpler to implement than a custom Lambda function.

upvoted 4 times

  **anikety123** 1 year, 2 months ago

**Selected Answer: B**

Its B. Check the AWS link



[https://aws.amazon.com/solutions/implementations/instance-scheduler-on-aws/?nc1=h\\_ls](https://aws.amazon.com/solutions/implementations/instance-scheduler-on-aws/?nc1=h_ls)

upvoted 1 times

  **anikety123** 1 year, 2 months ago

Sorry I wanted to select A.

upvoted 5 times



  **mrsoa** 1 year, 3 months ago

**Selected Answer: A**

A

<https://aws.amazon.com/solutions/implementations/instance-scheduler-on-aws/>

upvoted 2 times

  **luiscc** 1 year, 3 months ago

**Selected Answer: A**

Scheduler do the job

upvoted 4 times

A company uses locally attached storage to run a latency-sensitive application on premises. The company is using a lift and shift method to move the application to the AWS Cloud. The company does not want to change the application architecture.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Auto Scaling group with an Amazon EC2 instance. Use an Amazon FSx for Lustre file system to run the application.
- B. Host the application on an Amazon EC2 instance. Use an Amazon Elastic Block Store (Amazon EBS) GP2 volume to run the application.
- C. Configure an Auto Scaling group with an Amazon EC2 instance. Use an Amazon FSx for OpenZFS file system to run the application.
- D. Host the application on an Amazon EC2 instance. Use an Amazon Elastic Block Store (Amazon EBS) GP3 volume to run the application.

**Correct Answer: D**

*Community vote distribution*

D (100%)

🗲️ 👤 **TariqKipkemei** Highly Voted 👍 5 months, 1 week ago

**Selected Answer: D**

MOST cost-effectively =GP3  
upvoted 5 times

🗲️ 👤 **potomac** Most Recent ⌚ 5 months, 4 weeks ago

**Selected Answer: D**

gp3 offers SSD-performance at a 20% lower cost per GB than gp2 volumes.  
upvoted 3 times

🗲️ 👤 **bojila** 8 months ago

GP3 is the latest version  
upvoted 3 times

🗲️ 👤 **Hades2231** 8 months, 1 week ago

**Selected Answer: D**

GP3 is the latest version, and it is cost effective  
upvoted 3 times

🗲️ 👤 **Guru4Cloud** 8 months, 2 weeks ago

**Selected Answer: D**

GP3 is preferable over GP2, FSx for Lustre, and FSx for OpenZFS is clear and convincing:

GP3 offers identical latency performance to GP2 at a lower price point.  
FSx options are higher performance but more expensive and require application changes.  
GP3 aligns better with lift and shift needs as a directly attached block storage volume.  
upvoted 4 times

🗲️ 👤 **taustin2** 8 months, 3 weeks ago

**Selected Answer: D**

Migrate your Amazon EBS volumes from gp2 to gp3 and save up to 20% on costs.  
upvoted 3 times

🗲️ 👤 **Vadbro7** 8 months, 3 weeks ago

Y not gp2  
upvoted 1 times

🗲️ 👤 **Ale1973** 8 months, 4 weeks ago

**Selected Answer: D**

My rational:  
Options A y C are based on autoscaling-group and no make sense for me on this scenery.  
Then, use Amazon EBS is the solution and GP2 or GP3 is the question.  
Requirement requires the most COST effective solution, then, I choose GP3  
upvoted 4 times

A company runs a stateful production application on Amazon EC2 instances. The application requires at least two EC2 instances to always be running.

A solutions architect needs to design a highly available and fault-tolerant architecture for the application. The solutions architect creates an Auto Scaling group of EC2 instances.

Which set of additional steps should the solutions architect take to meet these requirements?

- A. Set the Auto Scaling group's minimum capacity to two. Deploy one On-Demand Instance in one Availability Zone and one On-Demand Instance in a second Availability Zone.
- B. Set the Auto Scaling group's minimum capacity to four. Deploy two On-Demand Instances in one Availability Zone and two On-Demand Instances in a second Availability Zone.
- C. Set the Auto Scaling group's minimum capacity to two. Deploy four Spot Instances in one Availability Zone.
- D. Set the Auto Scaling group's minimum capacity to four. Deploy two On-Demand Instances in one Availability Zone and two Spot Instances in a second Availability Zone.

**Correct Answer: B**

*Community vote distribution*

B (75%)

A (25%)

  **luiscc** Highly Voted 1 year, 9 months ago

**Selected Answer: B**

By setting the Auto Scaling group's minimum capacity to four, the architect ensures that there are always at least two running instances. Deploying two On-Demand Instances in each of two Availability Zones ensures that the application is highly available and fault-tolerant. If one Availability Zone becomes unavailable, the application can still run in the other Availability Zone.

upvoted 19 times

  **Ale1973** Highly Voted 1 year, 8 months ago

**Selected Answer: A**

My rational is: Highly available = 2 AZ, and then 2 EC2 instances always running is 1 EC2 in each AZ. If an entire AZ fails, SacalinGroup deploy the minimun instances (2) on the running AZ

upvoted 14 times

  **baba365** 1 year, 7 months ago

Ans: A.

The application requires at least two EC2 instances to always be running = 2 minimum capacity... minimum cap of 4 ec2 will work but a waste of resources that doesn't follow well archi. framework.

upvoted 2 times

  **Ramdi1** 1 year, 7 months ago

it says always have to have two running, hence you need 4. two in each AV. it might be a waste of resource but if that what is required by the company then so be it. Also you out the 4 you cannot use spot instances because if the two instances on the on demand go down and you need to use the spot instance they could be called back at any point.

upvoted 5 times

  **Ramdi1** 1 year, 7 months ago

AZ \* not AV

upvoted 3 times

  **bora4motion** Most Recent 2 days, 18 hours ago

**Selected Answer: B**

A and C are not compatible with the "fault tolerant" condition.

upvoted 1 times


  **emakid** 10 months, 1 week ago

Option A:

Set the Auto Scaling group's minimum capacity to two. Deploy one On-Demand Instance in one Availability Zone and one On-Demand Instance in a second Availability Zone.

This configuration ensures that you have two instances running across two different AZs, which provides high availability. However, it does not take advantage of additional capacity to handle failures or spikes in demand. If either AZ becomes unavailable, you will have one running instance, but this does not meet the requirement of having at least two running instances at all times.

upvoted 1 times

  **emakid** 10 months, 1 week ago

**Selected Answer: B**

Set the Auto Scaling group's minimum capacity to four. Deploy two On-Demand Instances in one Availability Zone and two On-Demand Instances in a second Availability Zone.

This configuration provides high availability with four instances distributed across two AZs. The minimum capacity of four ensures that even if one instance fails, there are still two instances in each AZ to handle the load. This option is highly available and fault-tolerant but may be more than required if only two instances are needed to be always running.

upvoted 2 times

  **cjace** 10 months, 3 weeks ago

Answer is A: If one Availability Zone fails, the Auto Scaling group will automatically launch a new instance in a different, healthy Availability Zone to maintain the desired capacity of two instances. This is one of the key benefits of using Auto Scaling groups—they automatically maintain the desired number of instances across multiple Availability Zones, ensuring that your application is highly available and fault-tolerant. So even in the event of a failure in one Availability Zone, your application will continue to run on the required number of instances. This is why it's recommended to distribute instances across multiple Availability Zones when designing architectures for high availability and fault tolerance.

upvoted 1 times

  **LuongTo** 6 months ago

that is reasonable. but what do you think about "always"? With A, there's a moment there would be only one EC2 healthy.

upvoted 2 times

  **Marco\_St** 1 year, 3 months ago

**Selected Answer: B**

so indeed ASG can set up a new EC2 instance in another AZ if there is one AZ failed with fault but it failed to meet the need of always having 2 instance running before the new instance replacement is done in the working AZ. so this is why we deploy 2 instances per AZ

upvoted 2 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: B**

If it would not mention the "stateful" application, and if it would only have to be "highly available" but NOT "fault-tolerant", A would be fine.

upvoted 5 times

  **1rob** 1 year, 5 months ago

**Selected Answer: B**

From <<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-best-practices.html> : Spot Instances are not suitable for workloads that are inflexible, stateful, fault-intolerant, or tightly coupled between instance nodes. So C and D don't fit.

From <<https://docs.aws.amazon.com/whitepapers/latest/real-time-communication-on-aws/use-multiple-availability-zones.html> : Within the constructs of AWS, customers are encouraged to run their workloads in more than one Availability Zone. This ensures that customer applications can withstand even a complete Availability Zone failure - a very rare event in itself.

So a HA solution in this case implies a total of 4 instances, 2 per AZ.

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: B**

The main requirement here is a 'highly available and fault-tolerant architecture for the application', this covered by option B.

The application requires at least two EC2 instances to always be running, main word here being 'atleast' which means more than two is ok.

upvoted 2 times

  **Ramdi1** 1 year, 7 months ago

**Selected Answer: B**

B - Need 2 in each AZ and you cant use spot instances as it could be recalled.


upvoted 2 times

  **Mandar15** 1 year, 7 months ago

**Selected Answer: B**

Stateful is keyword here. 2 is minimum required all time.

upvoted 2 times

  **JA2018** 5 months, 1 week ago

minimum of 2x EC2 instances per AZ at all times.

upvoted 1 times

  **MII1975** 1 year, 7 months ago

**Selected Answer: A**

If a complete AZ fails, autoscale will lunch a second EC2 in the running AZ. If that short period of time is not always, which is not, then the answer is B, but I would take my chances and select A in the exam xD because the application is highly available and fault-tolerant.

upvoted 1 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: B**

- ° Minimum of 4 ensures at least 2 instances are always running in each AZ, meeting the HA requirement.
- ° On-Demand instances provide consistent performance and availability, unlike Spot.
- ° Spreading across 2 AZs adds fault tolerance, protecting from AZ failure.

upvoted 3 times

  **darkknight23** 1 year, 8 months ago

**Selected Answer: B**

While Spot Instances can be used to reduce costs, they might not provide the same level of availability and guaranteed uptime that On-Demand Instances offer. So I will go with B and not D.

upvoted 2 times

  **Sat897** 1 year, 9 months ago

**Selected Answer: B**

Highly available - 2 AZ and then 2 EC2 instances always running. 2 in each AZ.

upvoted 3 times

  **Sat897** 1 year, 9 months ago

Highly available - 2 AZ and then 2 EC2 instances always running. 2 in each AZ..

upvoted 2 times

An ecommerce company uses Amazon Route 53 as its DNS provider. The company hosts its website on premises and in the AWS Cloud. The company's on-premises data center is near the us-west-1 Region. The company uses the eu-central-1 Region to host the website. The company wants to minimize load time for the website as much as possible.

Which solution will meet these requirements?

- A. Set up a geolocation routing policy. Send the traffic that is near us-west-1 to the on-premises data center. Send the traffic that is near eu-central-1 to eu-central-1.
- B. Set up a simple routing policy that routes all traffic that is near eu-central-1 to eu-central-1 and routes all traffic that is near the on-premises datacenter to the on-premises data center.
- C. Set up a latency routing policy. Associate the policy with us-west-1.
- D. Set up a weighted routing policy. Split the traffic evenly between eu-central-1 and the on-premises data center.

**Correct Answer: A**

*Community vote distribution*

A (89%)

11%

  **Guru4Cloud** Highly Voted 1 year, 2 months ago

**Selected Answer: A**

The key reasons are:


Geolocation routing allows you to route users to the closest endpoint based on their geographic location. This will provide the lowest latency. Routing us-west-1 traffic to the on-premises data center minimizes latency for those users since it is also located near there. Routing eu-central-1 traffic to the eu-central-1 AWS region minimizes latency for users nearby. This achieves routing users to the closest endpoint on a geographic basis to optimize for low latency.

upvoted 5 times

  **PLN6302** 1 year, 2 months ago

why can't be the option C

upvoted 2 times

  **lemur88** 1 year, 2 months ago

You cannot associate the policy to us-west-1 as the AWS account is in eu-central-1

upvoted 4 times

  **awsgeek75** Most Recent 9 months, 3 weeks ago

**Selected Answer: A**

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-geo.html>

B can be done but definition of "near" is ambiguous

C wrong region

D wrong solution as splitting evenly does not reduce latency for on-prem server users

upvoted 2 times

  **Cyberkayu** 10 months, 2 weeks ago

**Selected Answer: A**

not C. Client do not have AWS us-west-1 region. Client have a on prem DC near west-1

not D. 2 people visit the site together near eu-central-1, one of the user may be thrown to west-1 due to load balancing on split even weighted policy.

A and B are both valid, latency = how soon user reach the datacenter and received a responses from the DC, round trip. So in short, geolocation or send user to the nearest DC will improve latency.

upvoted 2 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: A**

Geolocation routing policy allows you to route traffic based on the location of your users.

upvoted 4 times

  **t0nx** 11 months, 2 weeks ago

**Selected Answer: C**

C. Set up a latency routing policy. Associate the policy with us-west-1.

Explanation:



A latency routing policy directs traffic based on the lowest network latency to the specified AWS endpoint. Since the on-premises data center is near the us-west-1 Region, associating the policy with us-west-1 ensures that users near that region will be directed to the on-premises data center.

This allows for optimal routing, minimizing the load time for users based on their geographical proximity to the respective hosting locations (us-west-1 and eu-central-1).

Options A, B, and D do not explicitly consider latency or are not optimal for minimizing load time:

Option A (geolocation routing policy) would direct traffic based on the geographic location of the user but may not necessarily optimize for the lowest latency.

upvoted 2 times

  **awsgeek75** 9 months, 2 weeks ago

There is nothing in us-west-1 as the company's data centre is near us-west-1.

upvoted 1 times

  **Chiquitabandita** 11 months, 2 weeks ago

except I don't think that it should be applied to the west region. If Geolocation is applied and the west is closer to the client, but the west is having intermittent issues at the time, they will have a longer latency even though closer to that region. this is why I would apply latency in a real world solution.

upvoted 1 times

  **Chiquitabandita** 11 months, 2 weeks ago

in real world I think it should use latency routing if the main concern is to lower the latency but AWS likes to promote geolocation and if that is in the question I think that will be the answer so I choose A.

upvoted 1 times

  **baba365** 1 year, 1 month ago

The company wants to minimize load time for the website as much as possible... between data Centre and website or between users and website?

upvoted 1 times

  **Hades2231** 1 year, 2 months ago

**Selected Answer: A**

Geolocation is the key word

upvoted 2 times

  **Iemur88** 1 year, 2 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-geo.html>

upvoted 2 times

A company has 5 PB of archived data on physical tapes. The company needs to preserve the data on the tapes for another 10 years for compliance purposes. The company wants to migrate to AWS in the next 6 months. The data center that stores the tapes has a 1 Gbps uplink internet connectivity.

Which solution will meet these requirements MOST cost-effectively?

- A. Read the data from the tapes on premises. Stage the data in a local NFS storage. Use AWS DataSync to migrate the data to Amazon S3 Glacier Flexible Retrieval.
- B. Use an on-premises backup application to read the data from the tapes and to write directly to Amazon S3 Glacier Deep Archive.
- C. Order multiple AWS Snowball devices that have Tape Gateway. Copy the physical tapes to virtual tapes in Snowball. Ship the Snowball devices to AWS. Create a lifecycle policy to move the tapes to Amazon S3 Glacier Deep Archive.
- D. Configure an on-premises Tape Gateway. Create virtual tapes in the AWS Cloud. Use backup software to copy the physical tape to the virtual tape.

**Correct Answer: C**

*Community vote distribution*

C (98%)

  **adeyinkaamole** Highly Voted 1 year, 8 months ago

If you have made it to the end of the exam dump, you will definitely pass your exams in Jesus name. After over a year of Procrastination, I am finally ready to write my AWS Solutions Architect Exam. Thank you Exam Topics

upvoted 30 times

  **Hades2231** Highly Voted 1 year, 8 months ago

**Selected Answer: C**



Ready for the exam tomorrow. Wish you guys all the best. BTW Snowball Device comes in handy when you need to move a huge amount of data but cant afford any bandwidth loss

upvoted 12 times

  **MatAlves** Most Recent 7 months, 3 weeks ago

Oh, to think now we have to study 904 questions instead of just 583 lol

upvoted 3 times

  **JA2018** 5 months, 1 week ago

currently (22 Nov 2024), question count now stands at 1.019

upvoted 3 times

  **XXXXXINN** 7 months, 1 week ago

well, 2 weeks later today, it says 981 questions...

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: C**

5PB over 1GB connection will take approximately 15 months so anything with "transfer" is invalid. ABD are not practical.

C: Just order snowball

upvoted 5 times

  **pentium75** 1 year, 4 months ago

**Selected Answer: C**

Though we'll need more than 60 Snowball devices, C is the only option that works. The internet uplink could transport less than 2 PB in 6 months (otherwise, say with a 10 Gb uplink, D would work).

upvoted 6 times

  **Cyberkayu** 1 year, 4 months ago

transfer 5 PB data in 1Gbps link, assume 0 overhead and drop packet, need 485 days, 10 hours, 50 minutes, 40 seconds to complete.

Snowball it is. C

upvoted 1 times

  **SHAAHIBHUSHANAWS** 1 year, 5 months ago

C

<https://docs.aws.amazon.com/storagegateway/latest/tgw/using-tape-gateway-snowball.html>



upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: C**

Migrate petabyte-scale data stored on physical tapes to AWS using AWS Snowball  
<https://aws.amazon.com/snowball/#:~:text=Migrate-,petabyte%2Dscale,-data%20stored%20on>

upvoted 2 times

  **hungta** 1 year, 5 months ago

**Selected Answer: C**

5 PB data is too huge for using 1Gbps uplink. With this uplink, it takes more than 1 year to migrate this data.

upvoted 2 times

  **baba365** 1 year, 7 months ago

Answer: D for most cost effective.

If you are looking for a cost-effective, durable, long-term, offsite alternative for data archiving, deploy a Tape Gateway. With its virtual tape library (VTL) interface, you can use your existing tape-based backup software infrastructure to store data on virtual tape cartridges that you create -

<https://docs.aws.amazon.com/storagegateway/latest/tgw/WhatIsStorageGateway.html>

upvoted 1 times

  **Devsin2000** 1 year, 7 months ago

D  
<https://aws.amazon.com/storagegateway/vtl/>  
the bandwidth and available time is ample

upvoted 1 times

  **nnencode** 1 year, 7 months ago

**Selected Answer: A**

The most cost-effective solution to meet the requirements is to read the data from the tapes on premises. Stage the data in a local NFS storage. Use AWS DataSync to migrate the data to Amazon S3 Glacier Flexible Retrieval.

This solution is the most cost-effective because it uses the least amount of bandwidth. AWS DataSync is a service that transfers data between on-premises storage and Amazon S3. It uses a variety of techniques to optimize the transfer speed and reduce c

upvoted 1 times

  **lemur88** 1 year, 8 months ago

**Selected Answer: C**

Only thing that makes sense given the 1Gbps limitation

upvoted 1 times

  **Guru4Cloud** 1 year, 8 months ago

**Selected Answer: C**

Option C is likely the most cost-effective solution given the large data size and limited internet bandwidth. The physical data transfer and integration with the existing tape infrastructure provides efficiency benefits that can optimize the cost.

upvoted 3 times

  **barracouto** 1 year, 8 months ago

**Selected Answer: C**

Went through this dump twice now. Exam is in about an hour. Will update with results.

upvoted 2 times

  **Vaishali12** 1 year, 8 months ago

how was ur exam?  
was these dump que helpful?

upvoted 1 times

  **riccardoto** 1 year, 8 months ago

Finished the dump today - taking my exam tomorrow :-) Wish me luck!

upvoted 4 times

  **Ale1973** 1 year, 8 months ago

My rational: question is about which solution will meet these requirements MOST cost-effectively, not MOST time or effectively, then, my response is D (using Tape Gateways)

upvoted 4 times

A company is deploying an application that processes large quantities of data in parallel. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to prevent groups of nodes from sharing the same underlying hardware.

Which networking solution meets these requirements?

- A. Run the EC2 instances in a spread placement group.
- B. Group the EC2 instances in separate accounts.
- C. Configure the EC2 instances with dedicated tenancy.
- D. Configure the EC2 instances with shared tenancy.

Correct Answer: A

Community vote distribution

A (87%)

13%

  **czyboi** Highly Voted 1 year, 8 months ago

Selected Answer: A

A spread placement group is a group of instances that are each placed on distinct hardware.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

upvoted 12 times

  **pentium75** Highly Voted 1 year, 4 months ago

Selected Answer: A

A, spread placement group does exactly what is required here.  
Not C and D, tenancy determines whether the hardware is shared with other customers or not, it has nothing to with your own instances sharing hardware. (On the contrary, dedicated tenancy would spread your EC2 instances across as little nodes as possible.)

Not B, accounts have nothing to do with the issue.

upvoted 5 times

  **emakid** Most Recent 10 months, 1 week ago

Selected Answer: A

Option A: Run the EC2 instances in a spread placement group.

Spread Placement Group: This placement group strategy ensures that EC2 instances are distributed across distinct hardware to reduce the risk of correlated failures. Instances in a spread placement group are placed on different underlying hardware, which aligns with the requirement to prevent groups of nodes from sharing the same underlying hardware. This is a good fit for the scenario where you need to ensure high availability and fault tolerance.

upvoted 2 times

  **Gape4** 10 months, 1 week ago

Selected Answer: A

Spread – Strictly places a small group of instances across distinct underlying hardware to reduce correlated failures.

upvoted 2 times

  **Marco\_St** 1 year, 3 months ago

Selected Answer: A

dedicated tenancy cannot ensure the instances share the same hardware. So A

upvoted 3 times

  **maged123** 1 year, 4 months ago

Selected Answer: A

Let's assume that you have two groups of instances, group A and group B and you have two physical hardware X and Y. With spread placement group, you can have group A of instances on hardware X and group B on hardware Y but this will not prevent hardware X to host other instances of other customers because your only requirement is to separate group A from group B. On the other hand, the dedicated tenancy means that AWS will dedicate the physical hardware only for you. So, the correct answer is A.

upvoted 3 times

  **Murtadhaceit** 1 year, 4 months ago

Question is ambiguous and confusing. Is it asking about the EC2 instance of the same application not sharing hardware? or EC2 instance not sharing hardware with other EC2 from other applications?



upvoted 2 times


  **Mikado211** 1 year, 5 months ago

**Selected Answer: A**

Spread placement group allows you to isolate your instances on hardware level.  
Dedicated tenancy allows you to be sure that you are the only customer on the hardware.

The correct answer is A.  
upvoted 2 times

  **Mikado211** 1 year, 5 months ago  
A : Spread placement group  
upvoted 2 times

  **lucasbg** 1 year, 5 months ago

**Selected Answer: A**

Def is A: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>  
upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: A**

Keywords 'prevent groups of nodes from sharing the same underlying hardware'.  
Spread Placement Group strictly places a small group of instances across distinct underlying hardware to reduce correlated failures.  
upvoted 2 times

  **cciesam** 1 year, 5 months ago

**Selected Answer: A**

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>  
Each instances is placed on seven different racks, each rack has its own network and power source.  
upvoted 2 times

  **wsdasdasdqwda** 1 year, 6 months ago

Another tricky question, but I would go for A because:

Dedicated instances:  
Dedicated Instances are EC2 instances that run on hardware that's dedicated to a single customer. Dedicated Instances that belong to different AWS accounts are physically isolated at a hardware level, even if those accounts are linked to a single payer account. However, Dedicated Instances might share hardware with other instances from the same AWS account that are not Dedicated Instances.  
Which is not the desired option.

Spread – strictly places a small group of instances across distinct underlying hardware to reduce correlated failures.

That's why A.  
upvoted 3 times

  **garuta** 1 year, 7 months ago

**Selected Answer: C**

C is clear.  
upvoted 1 times

  **pentium75** 1 year, 4 months ago

"Dedicated tenancy" means that all your nodes run on hardware that is not shared with other customers. This is counter-productive to the objective here.  
upvoted 2 times

  **Devsin2000** 1 year, 7 months ago

A  
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.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>  
upvoted 3 times

  **taustin2** 1 year, 7 months ago

**Selected Answer: A**

Spread Placement Group strictly places a small group of instances across distinct underlying hardware to reduce correlated failures.  
upvoted 1 times

  **Guru4Cloud** 1 year, 7 months ago

**Selected Answer: C**

C is the correct answer.

Configuring the EC2 instances with dedicated tenancy ensures that each instance will run on isolated, single-tenant hardware. This meets the requirement to prevent groups of nodes from sharing underlying hardware.

A spread placement group only provides isolation at the Availability Zone level. Instances could still share hardware within an AZ.  
upvoted 5 times

  **pentium75** 1 year, 4 months ago

No. C ensures that your EC2 instances run on hardware that is not shared with other customers (!). It is still shared among YOUR instances.  
upvoted 6 times

  **Eminenza22** 1 year, 8 months ago

**Selected Answer: A**

Option A is the correct answer. It suggests running the EC2 instances in a spread placement group. This solution is cost-effective and requires minimal development effort .

upvoted 3 times

  **Eminenza22** 1 year, 8 months ago

The placement group reduces the risk of simultaneous failures by spreading the instances across distinct underlying hardware

upvoted 2 times

A solutions architect is designing a disaster recovery (DR) strategy to provide Amazon EC2 capacity in a failover AWS Region. Business requirements state that the DR strategy must meet capacity in the failover Region.

Which solution will meet these requirements?

- A. Purchase On-Demand Instances in the failover Region.
- B. Purchase an EC2 Savings Plan in the failover Region.
- C. Purchase regional Reserved Instances in the failover Region.
- D. Purchase a Capacity Reservation in the failover Region.

**Correct Answer: D**

*Community vote distribution*

D (93%)

7%

  **EllenLiu** 4 months, 2 weeks ago

**Selected Answer: D**

it is called on-demand capacity reservation actually . no discount. must specify the following attributes:  
Availability Zone & Instance type & os type & tenancy


<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>

upvoted 2 times

  **EllenLiu** 4 months, 2 weeks ago

DR is a classic use case for capacity reservation

upvoted 2 times

  **emakid** 10 months, 1 week ago

**Selected Answer: D**

Option D: Purchase a Capacity Reservation in the failover Region.

Capacity Reservation: Capacity Reservations ensure that you have reserved capacity in a specific region for your instances, regardless of whether you are using On-Demand or Reserved Instances. This is ideal for DR scenarios because it guarantees that the required EC2 capacity will be available when needed.

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

**Selected Answer: D**

"Business requirements state that the DR strategy must meet capacity in the failover Region"  
so only D meets these requirements

- A. No reservation of capacity
- B. Saving plans don't guarantee capacity
- C. Can be possible but it's like an active instance so doesn't really make sense

upvoted 3 times

  **awsgeek75** 1 year, 3 months ago

Correction on C (I mixed it up with tenancy!). Reserved instance are not really for capacity, its for type of instance which gives good discount but that is not required here.

upvoted 2 times

  **Derek\_G** 1 year, 4 months ago

**Selected Answer: D**

Purchase a Capacity Reservation in the failover Region:

A Capacity Reservation allows you to reserve a specific amount of EC2 instance capacity in a given region without purchasing specific instances. This reserved capacity is dedicated to your account and can be utilized for launching instances when needed. Capacity Reservations offer flexibility, allowing you to launch different instance types and sizes within the reserved capacity.

Purchase regional Reserved Instances in the failover Region:

Regional Reserved Instances involve paying an upfront fee to reserve a certain number of specific EC2 instances in a particular region. These reserved instances are of a predefined type and size, providing a more traditional reservation model. Regional Reserved Instances are specific to a designated region and ensure that the reserved instances of a particular specification are available when needed.

upvoted 4 times

  **TheLaPlanta** 1 year, 1 month ago

What I don't get is... can't you accomplish that by using on-demand? I understood that you can scale infinitely

upvoted 2 times

  **Tanidanindo** 1 year, 1 month ago

It won't guarantee that you have the capacity when you need it. If available, like in most cases, it'll work. But the scenario requires a guarantee that the capacity will be available.

upvoted 2 times

  **SHAAHIBHUSHANAWS** 1 year, 5 months ago

D

Ask is to reserve capacity with RI capacity is not reserved also you can reserve capacity along with RI but only in AZ .

<https://repost.aws/knowledge-center/ri-reserved-capacity>

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: D**

Capacity Reservations mitigate against the risk of being unable to get On-Demand capacity in case there are capacity constraints. If you have strict capacity requirements, and are running business-critical workloads that require a certain level of long or short-term capacity assurance, create a Capacity Reservation to ensure that you always have access to Amazon EC2 capacity when you need it, for as long as you need it.

upvoted 2 times

  **potomac** 1 year, 5 months ago

**Selected Answer: D**

Capacity Reservations enable you to reserve capacity for your Amazon EC2 instances in a specific Availability Zone for any duration. This gives you the flexibility to selectively add capacity reservations and still get the Regional RI discounts for that usage. By creating Capacity Reservations, you ensure that you always have access to Amazon EC2 capacity when you need it, for as long as you need it.

upvoted 3 times

  **potomac** 1 year, 5 months ago

Savings Plans does not provide a capacity reservation.

upvoted 2 times

  **Guru4Cloud** 1 year, 7 months ago

**Selected Answer: D**

Capacity Reservations allocate EC2 capacity in a specific AWS Region for you to launch instances.

The capacity is reserved and available to be utilized when needed, meeting the requirement to provide EC2 capacity in the failover region.

Other options do not reserve capacity. On-Demand provides flexible capacity but does not reserve capacity upfront. Savings Plans and Reserved Instances provide discounts but do not reserve capacity.

Capacity Reservations allow defining instance attributes like instance type, platform, Availability Zone so the reserved capacity matches the production environment.

upvoted 4 times

  **Eminenza22** 1 year, 8 months ago

**Selected Answer: D**

A regional Reserved Instance does not reserve capacity

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/reserved-instances-scope.html>

upvoted 3 times

  **judyda** 1 year, 8 months ago

**Selected Answer: D**

reserved instances for price discount. need capacity reservation.

upvoted 3 times

  **gispankaj** 1 year, 8 months ago

**Selected Answer: C**

The Reserved Instance discount applies to instance usage within the instance family, regardless of size.



upvoted 2 times

  **pentium75** 1 year, 4 months ago

"Reserved Instances are not physical instances, but rather a billing discount "

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-reserved-instances.html>

upvoted 2 times

  **ErnShm** 1 year, 8 months ago

D

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>

upvoted 2 times

A company has five organizational units (OUs) as part of its organization in AWS Organizations. Each OU correlates to the five businesses that the company owns. The company's research and development (R&D) business is separating from the company and will need its own organization. A solutions architect creates a separate new management account for this purpose.

What should the solutions architect do next in the new management account?

- A. Have the R&D AWS account be part of both organizations during the transition.
- B. Invite the R&D AWS account to be part of the new organization after the R&D AWS account has left the prior organization.
- C. Create a new R&D AWS account in the new organization. Migrate resources from the prior R&D AWS account to the new R&D AWS account.
- D. Have the R&D AWS account join the new organization. Make the new management account a member of the prior organization.

**Correct Answer: B**

*Community vote distribution*

B (90%)

10%

- awsgeek75**

Highly Voted

1 year, 3 months ago

Selected Answer: B

An account can only join another org when it leaves the first org.  
A is wrong as it's not possible  
C that's a new account so not really a migration  
D The R&D department is separating from the company so you don't want the OU to join via nesting

upvoted 9 times
- pentium75**

Highly Voted

1 year, 4 months ago

Selected Answer: B

B as exactly described here: <https://repost.aws/knowledge-center/organizations-move-accounts>

upvoted 7 times
- emakid**

Most Recent

10 months, 1 week ago

Selected Answer: B

Option B: Invite the R&D AWS account to be part of the new organization after the R&D AWS account has left the prior organization is the appropriate approach. This option ensures that the R&D AWS account transitions smoothly from the old organization to the new one. The steps involved are:  
  
Remove the R&D AWS account from the existing organization: This is done from the existing organization's management account.  
  
Invite the R&D AWS account to join the new organization: Once the R&D account is no longer part of the previous organization, it can be invited to and accepted into the new organization.

upvoted 2 times
- Marco\_St**

1 year, 3 months ago

Selected Answer: B

<https://aws.amazon.com/blogs/mt/migrating-accounts-between-aws-organizations-with-consolidated-billing-to-all-features/>

upvoted 4 times
- ale\_brd\_111**

1 year, 4 months ago

Selected Answer: B

<https://repost.aws/knowledge-center/organizations-move-accounts>  
Remove the member account from the old organization.  
Send an invite to the member account from the new organization.  
Accept the invite to the new organization from the member account.

upvoted 3 times
- Derek\_G**

1 year, 4 months ago

Selected Answer: C

C is better. first migrate , then delete. avoid the data lose.

upvoted 1 times
- pentium75**

1 year, 4 months ago

What kind of "data lose" would happen when you change the account to a new organization? And why should you migrate ALL RESOURCES of the account to a new account?

upvoted 2 times

- Derek\_G

1 year, 4 months ago

C is better. first migrate , then delete. avoid the data lose.

upvoted 1 times
- TariqKipkemei

1 year, 5 months ago

Selected Answer: B

As per this document, B is clearly the answer.  
<https://repost.aws/knowledge-center/organizations-move-accounts#:~:text=In%20either%20case%2C-,perform%20these%20actions,-for%20each%20member>

upvoted 2 times
- Joben

1 year, 7 months ago

Selected Answer: B

In either case, perform these actions for each member account:
  - Remove the member account from the old organization.
  - Send an invite to the member account from the new organization.
  - Accept the invite to the new organization from the member account.  
<https://repost.aws/knowledge-center/organizations-move-accounts>

upvoted 6 times
- Guru4Cloud

1 year, 7 months ago

Selected Answer: C

Creating a brand new AWS account in the new organization (Option C) allows for a clean separation and migration of only the necessary resources from the old account to the new.

upvoted 2 times
- pentium75

1 year, 4 months ago

"A clean separation" is already existing, they have their own account. "Migration of only the necessary resources from the old account to the new" is not asked for. They have an account in an existing organization, they need their own organization, thus move the existing account to a new organisation (B), done.

upvoted 2 times
- Guru4Cloud

1 year, 7 months ago

Selected Answer: C

When separating a business unit from an AWS Organizations structure, best practice is to:  
  
Create a new AWS account dedicated for the business unit in the new organization  
Migrate resources from the old account to the new account  
Remove the old account from the original organization  
This allows a clean break between the organizations and avoids any linking between them after separation.

upvoted 1 times
- pentium75

1 year, 4 months ago

Says who?

upvoted 2 times
- ErnShm

1 year, 8 months ago

B  
<https://aws.amazon.com/blogs/mt/migrating-accounts-between-aws-organizations-with-consolidated-billing-to-all-features/>

upvoted 3 times
- gispankaj

1 year, 8 months ago

Selected Answer: B

account can leave current organization and then join new organization.

upvoted 4 times

A company is designing a solution to capture customer activity in different web applications to process analytics and make predictions. Customer activity in the web applications is unpredictable and can increase suddenly. The company requires a solution that integrates with other web applications. The solution must include an authorization step for security purposes.

Which solution will meet these requirements?



- A. Configure a Gateway Load Balancer (GWLB) in front of an Amazon Elastic Container Service (Amazon ECS) container instance that stores the information that the company receives in an Amazon Elastic File System (Amazon EFS) file system. Authorization is resolved at the GWLB.
- B. Configure an Amazon API Gateway endpoint in front of an Amazon Kinesis data stream that stores the information that the company receives in an Amazon S3 bucket. Use an AWS Lambda function to resolve authorization.
- C. Configure an Amazon API Gateway endpoint in front of an Amazon Kinesis Data Firehose that stores the information that the company receives in an Amazon S3 bucket. Use an API Gateway Lambda authorizer to resolve authorization.
- D. Configure a Gateway Load Balancer (GWLB) in front of an Amazon Elastic Container Service (Amazon ECS) container instance that stores the information that the company receives on an Amazon Elastic File System (Amazon EFS) file system. Use an AWS Lambda function to resolve authorization.

**Correct Answer: C**

*Community vote distribution*



C (86%)

14%

-   **ralfj** Highly Voted 1 year, 8 months ago



**Selected Answer: C**

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-use-lambda-authorizer.html>

upvoted 9 times
-   **FlyingHawk** Most Recent 3 months, 1 week ago

**Selected Answer: C**

<https://docs.aws.amazon.com/firehose/latest/dev/basic-deliver.html>

upvoted 1 times
-   **emakid** 10 months, 1 week ago

**Selected Answer: C**

Option C: Configure an Amazon API Gateway endpoint in front of an Amazon Kinesis Data Firehose that stores the information that the company receives in an Amazon S3 bucket. Use an API Gateway Lambda authorizer to resolve authorization.



This solution meets the requirements in the following ways:

Handles Unpredictable Traffic: Amazon Kinesis Data Firehose can handle variable amounts of streaming data and automatically scales to accommodate sudden increases in traffic.

Integration with Web Applications: Amazon API Gateway provides a RESTful API endpoint for integrating with web applications.



Authorization: An API Gateway Lambda authorizer provides the necessary authorization step to secure API access.

Data Storage: Amazon Kinesis Data Firehose can deliver data directly to an Amazon S3 bucket for storage, making it suitable for long-term analytics and predictions.



upvoted 4 times
-   **Matte\_** 11 months, 2 weeks ago

**Selected Answer: B**



<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-use-lambda-authorizer.html>

upvoted 2 times
-   **MatAlves** 7 months, 3 weeks ago

You cannot use Kinesis Data stream to store data in S3. You need Firehose for that.

upvoted 3 times
-   **Mr\_Marcus** 11 months, 1 week ago

Ummm. This link (Use API Gateway Lambda authorizers) helps to validate "C" as the correct answer, not "B".

upvoted 2 times
-   **4fad2f8** 1 year, 3 months ago

**Selected Answer: B**

B. Amazon Kinesis Data Firehose does not save anything  
upvoted 2 times

  **FlyingHawk** 3 months, 1 week ago

Amazon Data Firehose can send data to S3: <https://docs.aws.amazon.com/firehose/latest/dev/s3-object-name.html>  
upvoted 1 times

  **jaswantn** 1 year, 2 months ago

option C...Amazon Kinesis Data Firehose that stores the information (that the company receives) in an Amazon S3 bucket.  
This answer statement is worded in a complex way. It means to say that Firehose stores the data in S3 ...which company receives from API Gateway.  
upvoted 3 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: C**

Configure an Amazon API Gateway endpoint in front of an Amazon Kinesis Data Firehose that stores the information that the company receives in an Amazon S3 bucket. Use an API Gateway Lambda authorizer to resolve authorization.  
upvoted 4 times

  **wsdasdasdqwda** 1 year, 6 months ago

Using ECS just to store the information is an overkill. So B or C then, lambda authoriser is the key word => C  
upvoted 4 times

  **Eminenza22** 1 year, 8 months ago

**Selected Answer: C**

<https://docs.aws.amazon.com/lambda/latest/dg/services-kinesisfirehose.html>  
upvoted 3 times

  **ErnShm** 1 year, 8 months ago

C  
  
authorizer is configured for the method. If it is, API Gateway calls the Lambda function. The Lambda function authenticates the caller by means such as the following: Calling out to an OAuth provider to get an OAuth access token  
upvoted 3 times

  **gispankaj** 1 year, 8 months ago

**Selected Answer: C**

lambda authoriser seems to be logical solution.  
upvoted 3 times

An ecommerce company wants a disaster recovery solution for its Amazon RDS DB instances that run Microsoft SQL Server Enterprise Edition. The company's current recovery point objective (RPO) and recovery time objective (RTO) are 24 hours.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a cross-Region read replica and promote the read replica to the primary instance.
- B. Use AWS Database Migration Service (AWS DMS) to create RDS cross-Region replication.
- C. Use cross-Region replication every 24 hours to copy native backups to an Amazon S3 bucket.
- D. Copy automatic snapshots to another Region every 24 hours.

Correct Answer: D

Community vote distribution

D (100%)

- TiagueteVital

Highly Voted

1 year, 8 months ago

Selected Answer: D

Snapshots are always a cost-efficiency way to have a DR plan.  
upvoted 5 times
- awsgeek75

Highly Voted

1 year, 3 months ago

Selected Answer: D

Cross region data transfer is billable so think of smallest amount of data to transfer every 24 hours  
upvoted 5 times
- MatAlves

Most Recent

7 months, 3 weeks ago

A, B, C => cross-region \$\$  
D => copy snapshots -> most cost-effectively.  
upvoted 2 times
- emakid

10 months, 1 week ago

Selected Answer: D

Option D: Copy automatic snapshots to another Region every 24 hours.  
  
Explanation: This option involves copying RDS automatic snapshots to another Region. It is a straightforward way to ensure that snapshots are available in the event of a disaster. Since RDS snapshots are typically incremental and copied periodically, this solution matches the 24-hour RPO requirement effectively and is cost-effective compared to maintaining constant cross-Region replication.  
upvoted 2 times
- potomac

1 year, 5 months ago

Selected Answer: D

Amazon RDS creates and saves automated backups of your DB instance or Multi-AZ DB cluster during the backup window of your DB instance. RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. RDS saves the automated backups of your DB instance according to the backup retention period that you specify. If necessary, you can recover your DB instance to any point in time during the backup retention period.  
upvoted 3 times
- wsdasdasdqwda

1 year, 6 months ago

most cost-effective way is just copying the snapshot (24h delta in the storage). => D  
upvoted 3 times
- Guru4Cloud

1 year, 7 months ago

Selected Answer: D

Dddddddddd  
upvoted 4 times

 **Eminenza22** 1 year, 8 months ago

**Selected Answer: D**

This is the most cost-effective solution because it does not require any additional AWS services. Amazon RDS automatically creates snapshots of your DB instances every hour. You can copy these snapshots to another Region every 24 hours to meet your RPO and RTO requirements.

The other solutions are more expensive because they require additional AWS services. For example, AWS DMS is a more expensive service than AWS RDS.

upvoted 4 times

A company runs a web application on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer that has sticky sessions enabled. The web server currently hosts the user session state. The company wants to ensure high availability and avoid user session state loss in the event of a web server outage.

Which solution will meet these requirements?

- A. Use an Amazon ElastiCache for Memcached instance to store the session data. Update the application to use ElastiCache for Memcached to store the session state.
- B. Use Amazon ElastiCache for Redis to store the session state. Update the application to use ElastiCache for Redis to store the session state.
- C. Use an AWS Storage Gateway cached volume to store session data. Update the application to use AWS Storage Gateway cached volume to store the session state.
- D. Use Amazon RDS to store the session state. Update the application to use Amazon RDS to store the session state.

**Correct Answer: B**

*Community vote distribution*

B (91%)

9%

- Guru4Cloud**

Highly Voted

1 year, 7 months ago

Selected Answer: B

The key points are:

ElastiCache Redis provides in-memory caching that can deliver microsecond latency for session data. Redis supports replication and multi-AZ which can provide high availability for the cache. The application can be updated to store session data in ElastiCache Redis rather than locally on the web servers. If a web server fails, the user can be routed via the load balancer to another web server which can retrieve their session data from the highly available ElastiCache Redis cluster.

upvoted 9 times
- czyboi**

Highly Voted

1 year, 8 months ago

Selected Answer: B

high availability => use redis instead of Elastich memcache

upvoted 5 times
- emakid**

Most Recent

10 months, 1 week ago

Selected Answer: B

Option B: Use Amazon ElastiCache for Redis to store the session state. Update the application to use ElastiCache for Redis to store the session state.

Explanation: Amazon ElastiCache for Redis is suitable for session state storage because Redis provides both in-memory data storage and persistence options. Redis supports features like replication, persistence, and high availability (through Redis Sentinel or clusters). This ensures that session state is preserved and available even if individual web servers fail.

upvoted 2 times
- pentium75**

1 year, 4 months ago

Selected Answer: B

As Memcached is not HA

upvoted 4 times
- SHAAHIBHUSHANAWS**

1 year, 5 months ago

A

As cache needs to be distributed as ALB is used.

upvoted 1 times
- potomac**

1 year, 5 months ago

Selected Answer: B

B is correct

upvoted 3 times
- franbarberan**

1 year, 7 months ago

Selected Answer: D

Elastic cache is Only for RDS

upvoted 3 times



**pentium75** 1 year, 4 months ago

Since when?

upvoted 4 times



**gispankaj** 1 year, 8 months ago

**Selected Answer: B**

redis is correct since it provides high availability and data persistance

upvoted 5 times



**Eminenza22** 1 year, 8 months ago

**Selected Answer: B**

B is the correct answer. It suggests using Amazon ElastiCache for Redis to store the session state. Update the application to use ElastiCache for Redis to store the session state. This solution is cost-effective and requires minimal development effort.

upvoted 4 times

Question #590

Topic 1

A company migrated a MySQL database from the company's on-premises data center to an Amazon RDS for MySQL DB instance. The company sized the RDS DB instance to meet the company's average daily workload. Once a month, the database performs slowly when the company runs queries for a report. The company wants to have the ability to run reports and maintain the performance of the daily workloads.

Which solution will meet these requirements?

- A. Create a read replica of the database. Direct the queries to the read replica.
- B. Create a backup of the database. Restore the backup to another DB instance. Direct the queries to the new database.
- C. Export the data to Amazon S3. Use Amazon Athena to query the S3 bucket.
- D. Resize the DB instance to accommodate the additional workload.

**Correct Answer: A**

*Community vote distribution*

A (100%)



**TariqKipkemei** 5 months, 1 week ago

**Selected Answer: A**

queries for reports = read replica

upvoted 3 times



**Guru4Cloud** 7 months, 3 weeks ago

**Selected Answer: A**

Create a read replica of the database. Direct the queries to the read replica.

upvoted 4 times



**Eminenza22** 8 months ago

**Selected Answer: A**

This is the most cost-effective solution because it does not require any additional AWS services. A read replica is a copy of a database that is synchronized with the primary database. You can direct the queries for the report to the read replica, which will not affect the performance of the daily workloads

upvoted 4 times



**TiagueteVital** 8 months ago

**Selected Answer: A**

Clearly the right choice, with a read replica all the queries needed for a report are done in the replica, leaving the primary on best perfomance for write

upvoted 3 times

A company runs a container application by using Amazon Elastic Kubernetes Service (Amazon EKS). The application includes microservices that manage customers and place orders. The company needs to route incoming requests to the appropriate microservices.

Which solution will meet this requirement MOST cost-effectively?

- A. Use the AWS Load Balancer Controller to provision a Network Load Balancer.
- B. Use the AWS Load Balancer Controller to provision an Application Load Balancer.
- C. Use an AWS Lambda function to connect the requests to Amazon EKS.
- D. Use Amazon API Gateway to connect the requests to Amazon EKS.

Correct Answer: B

Community vote distribution

B (77%)

D (23%)

  **awsgeek75** Highly Voted  1 year, 3 months ago

Selected Answer: B

"The company needs to route incoming requests to the appropriate microservices"  
In Kubernetes world, this would be called an Ingress Service so it will need B  
<https://kubernetes-sigs.github.io/aws-load-balancer-controller/v2.6/>  
<https://kubernetes.io/docs/concepts/services-networking/ingress/>  
upvoted 7 times

  **pentium75** Highly Voted  1 year, 4 months ago

Selected Answer: B

Not D because  
- even with an API gateway you'd need an ALB or ELB (so B+D would work, but D alone does not)  
- you would use AWS API Gateway Controller (not "Amazon API Gateway") to create the API Gateway  
upvoted 5 times

  **pentium75** 1 year, 4 months ago

<https://aws.amazon.com/blogs/containers/microservices-development-using-aws-controllers-for-kubernetes-ack-and-amazon-eks-blueprints/>

<https://aws.amazon.com/blogs/containers/integrate-amazon-api-gateway-with-amazon-eks/>

upvoted 4 times

  **emakid** Most Recent  10 months, 1 week ago

Selected Answer: B

Option B: Use the AWS Load Balancer Controller to provision an Application Load Balancer (ALB).

Explanation: The AWS Load Balancer Controller can provision ALBs, which operate at the application layer (Layer 7). ALBs support advanced routing capabilities such as routing based on HTTP paths or hostnames. This makes ALBs well-suited for routing requests to different microservices based on URL paths or domains. This approach integrates well with Kubernetes and is a common pattern for microservices architectures.  
upvoted 2 times

  **[Removed]** 1 year, 4 months ago

Selected Answer: B

ALB is cost-effectively  
upvoted 4 times

  **Mikado211** 1 year, 4 months ago



Selected Answer: B

ALB is considered as expensive than API Gateway, particularly on higher load.

If you do not need any specific functionalities of API Gateway so you must choose ALB because it will be cheaper.  
upvoted 4 times

  **Mikado211** 1 year, 4 months ago



ALB is considered as LESS expensive  
upvoted 2 times

  **riyasara** 1 year, 5 months ago

Selected Answer: B

API Gateway has a pricing model that includes a cost per API call, and depending on the volume of requests, this could potentially be more expensive than using an Application Load Balancer.

upvoted 4 times

  **1rob** 1 year, 5 months ago

**Selected Answer: B**

Routing requests to the appr. microserv. can easily be done with ALB and ingress. The ingress handles routing rules to the micro.serv. With answer D you will still need ALB or NLB as can be seen in the pics of <https://aws.amazon.com/blogs/containers/integrate-amazon-api-gateway-with-amazon-eks/> or <https://aws.amazon.com/blogs/containers/microservices-development-using-aws-controllers-for-kubernetes-ack-and-amazon-eks-blueprints/> so that is not the most cost-effectively.

upvoted 3 times

  **ale\_brd\_111** 1 year, 4 months ago

yeah, I was going with D than checked and seems that you are right to deploy API gateway you still a LB

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: D**

Both ALB and API gateway can be used to route traffic to the microservices, but the question seeks the most 'cost effective' option.

You are charged for each hour or partial hour that an Application Load Balancer is running, and the number of Load Balancer Capacity Units (LCU) used per hour.

With Amazon API Gateway, you only pay when your APIs are in use.

I say API gateway is the best option for this case.

upvoted 2 times



  **pentium75** 1 year, 4 months ago

But you still need an ALB or ELB

<https://aws.amazon.com/blogs/containers/microservices-development-using-aws-controllers-for-kubernetes-ack-and-amazon-eks-blueprints/>

<https://aws.amazon.com/blogs/containers/integrate-amazon-api-gateway-with-amazon-eks/>

upvoted 2 times

  **t0nx** 1 year, 5 months ago

**Selected Answer: B**

AWS Load Balancer Controller: The AWS Load Balancer Controller is a Kubernetes controller that makes it easy to set up an Application Load Balancer (ALB) or Network Load Balancer (NLB) for your Amazon EKS clusters. It simplifies the process of managing load balancers for applications running on EKS.

Application Load Balancer (ALB): ALB is a Layer 7 load balancer that is capable of routing requests based on content, such as URL paths or hostnames. This makes it suitable for routing requests to different microservices based on specific criteria.

Cost-Effectiveness: ALB is typically more cost-effective than an NLB, and it provides additional features at the application layer, which may be useful for routing requests to microservices based on specific conditions.

Option D: Amazon API Gateway is designed for creating, publishing, and managing APIs. While it can integrate with Amazon EKS, it may be more feature-rich and complex than needed for simple routing to microservices within an EKS cluster.

upvoted 4 times



  **potomac** 1 year, 5 months ago

**Selected Answer: D**

API Gateway provides an entry point to your microservices.



<https://aws.amazon.com/blogs/containers/integrate-amazon-api-gateway-with-amazon-eks/>

upvoted 3 times

  **ccmc** 1 year, 6 months ago

B is correct, it is a required before exposing through api gateway


upvoted 2 times

  **thanhnv142** 1 year, 6 months ago

B: is correct.

For EKS, use application load balancer to expose microservices



upvoted 4 times

  **KhasDenis** 1 year, 7 months ago

**Selected Answer: B**

Routing to ms in k8s -> Ingresses -> Ingress Controller -> AWS Load Balancer Controller <https://kubernetes-sigs.github.io/aws-load-balancer-controller/v2.6/>

upvoted 4 times

  **RDM10** 1 year, 7 months ago

Microservices--> API--> API GW

upvoted 3 times

 **Guru4Cloud** 1 year, 7 months ago

**Selected Answer: D**

D. Use Amazon API Gateway to connect the requests to Amazon EKS.

upvoted 3 times

 **MII1975** 1 year, 7 months ago

**Selected Answer: D**

API Gateway is a fully managed service that makes it easy for you to create, publish, maintain, monitor, and secure APIs at any scale. API Gateway provides an entry point to your microservices.

upvoted 1 times

 **Eminenza22** 1 year, 8 months ago

**Selected Answer: D**

<https://aws.amazon.com/blogs/containers/microservices-development-using-aws-controllers-for-kubernetes-ack-and-amazon-eks-blueprints/>

upvoted 1 times

A company uses AWS and sells access to copyrighted images. The company's global customer base needs to be able to access these images quickly. The company must deny access to users from specific countries. The company wants to minimize costs as much as possible.

Which solution will meet these requirements?

- A. Use Amazon S3 to store the images. Turn on multi-factor authentication (MFA) and public bucket access. Provide customers with a link to the S3 bucket.
- B. Use Amazon S3 to store the images. Create an IAM user for each customer. Add the users to a group that has permission to access the S3 bucket.
- C. Use Amazon EC2 instances that are behind Application Load Balancers (ALBs) to store the images. Deploy the instances only in the countries the company services. Provide customers with links to the ALBs for their specific country's instances.
- D. Use Amazon S3 to store the images. Use Amazon CloudFront to distribute the images with geographic restrictions. Provide a signed URL for each customer to access the data in CloudFront.

**Correct Answer: D**

*Community vote distribution*

D (100%)

🗲️ 👤 **ralfj** Highly Voted 👍 8 months ago

**Selected Answer: D**

Use Cloudfront and geographic restriction  
upvoted 5 times

🗲️ 👤 **TariqKipkemei** Most Recent ⌚ 5 months, 1 week ago

**Selected Answer: D**

Store images = Amazon S3  
global customer base needs to be able to access these images quickly = Amazon CloudFront  
deny access to users from specific countries = Amazon CloudFront geographic restrictions, signed URLs  
upvoted 4 times

🗲️ 👤 **Guru4Cloud** 7 months, 3 weeks ago

**Selected Answer: D**

D. Use Amazon S3 to store the images. Use Amazon CloudFront to distribute the images with geographic restrictions. Provide a signed URL for each customer to access the data in CloudFront.  
upvoted 4 times

🗲️ 👤 **Colz** 7 months, 3 weeks ago

Correct answer is D  
upvoted 1 times

🗲️ 👤 **hubbabubba** 8 months ago

**Selected Answer: D**

answer is D  
upvoted 2 times

🗲️ 👤 **Eminenza22** 8 months ago

**Selected Answer: D**

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/georestrictions.html>  
upvoted 3 times

A solutions architect is designing a highly available Amazon ElastiCache for Redis based solution. The solutions architect needs to ensure that failures do not result in performance degradation or loss of data locally and within an AWS Region. The solution needs to provide high availability at the node level and at the Region level.

Which solution will meet these requirements?

- A. Use Multi-AZ Redis replication groups with shards that contain multiple nodes.
- B. Use Redis shards that contain multiple nodes with Redis append only files (AOF) turned on.
- C. Use a Multi-AZ Redis cluster with more than one read replica in the replication group.
- D. Use Redis shards that contain multiple nodes with Auto Scaling turned on.

**Correct Answer: A**

*Community vote distribution*

A (81%)

Other

  **TariqKipkemei** Highly Voted  11 months, 2 weeks ago

**Selected Answer: A**

Multi-AZ is only supported on Redis clusters that have more than one node in each shard (node groups).

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html#:~:text=node%20in%20each-,shard.,-Topics>  
upvoted 5 times

  **pentium75** Highly Voted  10 months, 1 week ago

**Selected Answer: A**

It seems like "Multi-AZ Redis replication group" (A) and "Multi-AZ Redis cluster" (C) are different wordings for the same configuration. However, "to minimize the impact of a node failure, we recommend that your implementation use multiple nodes in each shard" - and that is mentioned only in A.  
upvoted 5 times

  **LocNV** Most Recent  10 months, 1 week ago

**Selected Answer: A**

high availability at the node level = shard and Multi A-Z = region level  
upvoted 5 times

  **Cyberkayu** 10 months, 2 weeks ago

did client ask for improved performance, unfortunately they didn't, so C is good to have but not part of the business requirement.

My answer A.  
upvoted 3 times

  **SHAAHIBHUSHANAWS** 11 months ago

A  
Multi-AZ is only option. It is regional service so can use backup to replicate but can not use for failover.  
upvoted 2 times

  **t0nx** 11 months, 2 weeks ago

**Selected Answer: C**

C. Use a Multi-AZ Redis cluster with more than one read replica in the replication group.  
In summary, option C, using a Multi-AZ Redis cluster with more than one read replica, is designed to provide both node-level and AWS Region-level high availability, making it the most suitable choice for the given requirements.  
upvoted 2 times

  **potomac** 12 months ago

**Selected Answer: A**

the replication structure is contained within a shard (called node group in the API/CLI) which is contained within a Redis cluster

A shard (in the API and CLI, a node group) is a hierarchical arrangement of nodes, each wrapped in a cluster. Shards support replication. Within a shard, one node functions as the read/write primary node. All the other nodes in a shard function as read-only replicas of the primary node.  
upvoted 1 times

  **thanhnv142** 1 year ago


C is correct.  
Not A because in replication mode, shard have multiple nodes by default.

B and D not correct because that not an option  
upvoted 1 times

  **iwannabeawsgod** 1 year ago

**Selected Answer: C**

its c for me  
upvoted 1 times

  **bsbs1234** 1 year ago

C:  
Cluster mode will create multiple shards, when node level failure, request of shard that not impacted will not has any performance impact. If the issue at AZ level, spread traffic between multiple shards shall also reduce the performance degrade.  
upvoted 1 times

  **loveaws** 1 year ago

c.  
Option A is not ideal because it doesn't mention read replicas, and it's generally better to have read replicas for both performance and high availability.  
Option B mentions Redis append-only files (AOF), but AOF alone doesn't provide high availability or fault tolerance.  
Option D mentions Auto Scaling, but this doesn't directly address high availability at the Region level or data replication  
upvoted 1 times

  **taustin2** 1 year, 1 month ago

Multi-AZ is only supported on Redis clusters that have more than one node in each shard.  
upvoted 1 times

  **taustin2** 1 year, 1 month ago



**Selected Answer: A**

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.html>  
upvoted 5 times

  **Guru4Cloud** 1 year, 1 month ago

**Selected Answer: A**

Multi-AZ replication groups provide automatic failover between AZs if there is an issue with the primary AZ. This provides high availability at the region level  
upvoted 3 times

  **pentium75** 10 months, 1 week ago  
What about "the node level"?  
upvoted 1 times

  **xyb** 1 year, 1 month ago

**Selected Answer: C**

Enabling ElastiCache Multi-AZ with automatic failover on your Redis cluster (in the API and CLI, replication group) improves your fault tolerance. This is true particularly in cases where your cluster's read/write primary cluster becomes unreachable or fails for any reason. Multi-AZ with automatic failover is only supported on Redis clusters that support replication  
upvoted 1 times

  **MII1975** 1 year, 1 month ago

**Selected Answer: A**

I would go with A too

I would go with A, Using AOF can't protect you from all failure scenarios.  
For example, if a node fails due to a hardware fault in an underlying physical server, ElastiCache will provision a new node on a different server. In this case, the AOF is not available and can't be used to recover the data.  
upvoted 2 times

  **hubbabubba** 1 year, 2 months ago

**Selected Answer: A**

Hate to say this, but I read the two docs linked below, and I still think the answer is A. Turning on AOF helps in data persistence after failure, but it does nothing for availability unless you use Multi-AZ replica groups.  
upvoted 3 times

A company plans to migrate to AWS and use Amazon EC2 On-Demand Instances for its application. During the migration testing phase, a technical team observes that the application takes a long time to launch and load memory to become fully productive.

Which solution will reduce the launch time of the application during the next testing phase?

- A. Launch two or more EC2 On-Demand Instances. Turn on auto scaling features and make the EC2 On-Demand Instances available during the next testing phase.
- B. Launch EC2 Spot Instances to support the application and to scale the application so it is available during the next testing phase.
- C. Launch the EC2 On-Demand Instances with hibernation turned on. Configure EC2 Auto Scaling warm pools during the next testing phase.
- D. Launch EC2 On-Demand Instances with Capacity Reservations. Start additional EC2 instances during the next testing phase.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **Guru4Cloud** Highly Voted  1 year, 1 month ago

**Selected Answer: C**

Using EC2 hibernation and Auto Scaling warm pools will help address this:

Hibernation saves the in-memory state of the EC2 instance to persistent storage and shuts the instance down. When the instance is started again, the in-memory state is restored, which launches much faster than launching a new instance.

Warm pools pre-initialize EC2 instances and keep them ready to fulfill requests, reducing launch time. The hibernated instances can be added to a warm pool.

When auto scaling scales out during the next testing phase, it will be able to launch instances from the warm pool rapidly since they are already initialized

upvoted 8 times

  **awsgeek75** Most Recent  9 months, 3 weeks ago

**Selected Answer: C**

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-warm-pools.html>

upvoted 3 times

  **riyasara** 11 months ago

**Selected Answer: C**

Amazon EC2 hibernation and warm pool

upvoted 2 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: C**

If an instance or application takes a long time to bootstrap and build a memory footprint in order to become fully productive, you can use hibernation to pre-warm the instance. To pre-warm the instance, you:

Launch it with hibernation enabled.

Bring it to a desired state.

Hibernate it so that it's ready to be resumed to the desired state whenever needed.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Hibernate.html#:~:text=you%20can%20use-,hibernation,-to%20pre%2Dwarm>

upvoted 4 times

  **potomac** 12 months ago

**Selected Answer: C**

With Amazon EC2 hibernation enabled, you can maintain your EC2 instances in a "pre-warmed" state so these can get to a productive state faster.

upvoted 3 times

  **tabbyDolly** 1 year, 1 month ago

C: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Hibernate.html>

upvoted 3 times

  **ralfj** 1 year, 2 months ago

**Selected Answer: C**

just use hibernation option so you won't load the full EC2 Instance

upvoted 2 times

A company's applications run on Amazon EC2 instances in Auto Scaling groups. The company notices that its applications experience sudden traffic increases on random days of the week. The company wants to maintain application performance during sudden traffic increases.

Which solution will meet these requirements MOST cost-effectively?

- A. Use manual scaling to change the size of the Auto Scaling group.
- B. Use predictive scaling to change the size of the Auto Scaling group.
- C. Use dynamic scaling to change the size of the Auto Scaling group.
- D. Use schedule scaling to change the size of the Auto Scaling group.

**Correct Answer: C**

*Community vote distribution*

C (100%)

  **ralfj** Highly Voted 1 year, 2 months ago

**Selected Answer: C**

Dynamic Scaling – This is yet another type of Auto Scaling in which the number of EC2 instances is changed automatically depending on the signals received. Dynamic Scaling is a good choice when there is a high volume of unpredictable traffic.



<https://www.developer.com/web-services/aws-auto-scaling-types-best-practices/#:~:text=Dynamic%20Scaling%20%E2%80%93%20This%20is%20yet,high%20volume%20of%20unpredictable%20traffic.>

upvoted 6 times

  **tabbyDolly** Highly Voted 1 year, 1 month ago

C - " sudden traffic increases on random days of the week" --> dynamic scaling

upvoted 5 times

  **awsgeek75** Most Recent 9 months, 2 weeks ago

**Selected Answer: C**

random = dynamic

A: Manual is never a solution

B: Predictive is not possible as it's random

D: Cannot schedule random

upvoted 4 times

  **TariqKipkemei** 11 months, 2 weeks ago

**Selected Answer: C**

Dynamic scaling

upvoted 2 times

  **dilaaziz** 1 year ago

**Selected Answer: C**

<https://aws.amazon.com/ec2/autoscaling/faqs/>

upvoted 3 times

  **Guru4Cloud** 1 year, 1 month ago

**Selected Answer: C**

C is the best answer here. Dynamic scaling is the most cost-effective way to automatically scale the Auto Scaling group to maintain performance during random traffic spikes.

upvoted 3 times

An ecommerce application uses a PostgreSQL database that runs on an Amazon EC2 instance. During a monthly sales event, database usage increases and causes database connection issues for the application. The traffic is unpredictable for subsequent monthly sales events, which impacts the sales forecast. The company needs to maintain performance when there is an unpredictable increase in traffic.

Which solution resolves this issue in the MOST cost-effective way?

- A. Migrate the PostgreSQL database to Amazon Aurora Serverless v2.
- B. Enable auto scaling for the PostgreSQL database on the EC2 instance to accommodate increased usage.
- C. Migrate the PostgreSQL database to Amazon RDS for PostgreSQL with a larger instance type.
- D. Migrate the PostgreSQL database to Amazon Redshift to accommodate increased usage.

**Correct Answer: A**

*Community vote distribution*

A (95%)5%

- Guru4Cloud**

Highly Voted

1 year, 1 month ago

Selected Answer: A

Answer is A.  
Aurora Serverless v2 got autoscaling, highly available and cheaper when compared to the other options.

upvoted 8 times
- pentium75**

Most Recent

10 months, 1 week ago

Selected Answer: A

Not B - we can auto-scale the EC2 instance, but not "the [self-managed] PostgreSQL database ON the EC2 instance"  
Not C - This does not mention scaling, so it would incur high cost and still it might not be able to keep up with the "unpredictable" spikes  
Not D - Redshift is OLAP Data Warehouse

upvoted 3 times
- TariqKipkemei**

11 months, 1 week ago

Selected Answer: A

Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Aurora where the database automatically starts up, shuts down, and scales capacity up or down based on your application's needs. This is the least costly option for unpredictable traffic.

upvoted 3 times
- tabbyDolly**

1 year, 1 month ago

A: "he traffic is unpredictable for subsequent monthly sales events" --> serverless

upvoted 3 times
- Wayne23Fang**

1 year, 1 month ago

Selected Answer: C

A is probably more expensive than C. Aurora is serverless and fast. But nevertheless it needs DB migration service. Not sure DMS may not be free.

upvoted 1 times
- danielmakita**

1 year ago

C is more expensive if you think the scenario where the traffic is low. You are paying for a larger hardware but not using it. That's why I think A is correct.

upvoted 4 times
- TiagueteVital**

1 year, 2 months ago

Selected Answer: A

A to autoscaling

upvoted 3 times
- manOfThePeople**

1 year, 2 months ago

Answer is A.  
Aurora Serverless v2 got autoscaling, highly available and cheaper when compared to the other options.

upvoted 2 times

  **anikety123** 1 year, 2 months ago

**Selected Answer: A**

The correct answer is A  
upvoted 2 times

A company hosts an internal serverless application on AWS by using Amazon API Gateway and AWS Lambda. The company’s employees report issues with high latency when they begin using the application each day. The company wants to reduce latency.

Which solution will meet these requirements?

- A. Increase the API Gateway throttling limit.
- B. Set up a scheduled scaling to increase Lambda provisioned concurrency before employees begin to use the application each day.
- C. Create an Amazon CloudWatch alarm to initiate a Lambda function as a target for the alarm at the beginning of each day.
- D. Increase the Lambda function memory.

Correct Answer: B

Community vote distribution

B (100%)

- oayoade

Highly Voted

1 year, 8 months ago

Selected Answer: B

<https://aws.amazon.com/blogs/compute/scheduling-aws-lambda-provisioned-concurrency-for-recurring-peak-usage/>  
upvoted 5 times
- Guru4Cloud

Highly Voted

1 year, 7 months ago

Selected Answer: B

Set up a scheduled scaling to increase Lambda provisioned concurrency before employees begin to use the application each day.  
upvoted 5 times
- emakid

Most Recent

10 months, 1 week ago

Selected Answer: B

Option B: Set up a scheduled scaling to increase Lambda provisioned concurrency before employees begin to use the application each day.  
  
Explanation: Provisioned concurrency ensures that a specified number of Lambda instances are initialized and ready to handle requests. By scheduling this scaling, you can pre-warm Lambda functions before peak usage times, reducing cold start latency. This solution directly addresses the latency issue caused by cold starts.  
upvoted 4 times
- TariqKipkemei

1 year, 5 months ago

Selected Answer: B

Provisioned concurrency pre-initializes execution environments for your functions. These execution environments are prepared to respond immediately to incoming function requests at start of day.  
upvoted 3 times
- potomac

1 year, 5 months ago

Selected Answer: B

A is wrong  
API Gateway throttling limit is for better throughput, not for latency  
upvoted 2 times
- MII1975

1 year, 7 months ago

Selected Answer: B

Provisioned Concurrency incurs additional costs, so it is cost-efficient to use it only when necessary. For example, early in the morning when activity starts, or to handle recurring peak usage.  
upvoted 4 times
- Eminenza22

1 year, 8 months ago

Selected Answer: B

B option setting up a scheduled scaling to increase Lambda provisioned concurrency before employees begin to use the application each day. This solution is cost-effective and requires minimal development effort.  
upvoted 2 times

A research company uses on-premises devices to generate data for analysis. The company wants to use the AWS Cloud to analyze the data. The devices generate .csv files and support writing the data to an SMB file share. Company analysts must be able to use SQL commands to query the data. The analysts will run queries periodically throughout the day.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose three.)

- A. Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode.
- B. Deploy an AWS Storage Gateway on premises in Amazon FSx File Gateway made.
- C. Set up an AWS Glue crawler to create a table based on the data that is in Amazon S3.
- D. Set up an Amazon EMR cluster with EMR File System (EMRFS) to query the data that is in Amazon S3. Provide access to analysts.
- E. Set up an Amazon Redshift cluster to query the data that is in Amazon S3. Provide access to analysts.
- F. Setup Amazon Athena to query the data that is in Amazon S3. Provide access to analysts.

**Correct Answer: ACF**

*Community vote distribution*

ACF (96%)4%

- awsgeek75**

Highly Voted

 1 year, 3 months ago

Selected Answer: ACF

SQL Queries is Athena so DE are wrong and we are now dependant on S3  
A to get files into S3  
C Glue to convert CSV to S3 table data  
B irrelevant as we don't have anything to consume data from FSx in other options

upvoted 6 times
- JA2018** 5 months, 1 week ago

err,,, pardon my ignorance but I did not recall seeing any mention of S3 in the stem

upvoted 1 times
- awsgeek75** 1 year, 3 months ago

My only reservation with this answer is C.  
CSV is technically a table and Athena can query multiple csv from S3. Glue just seems overengineering over here

upvoted 3 times
- pentium75**

Highly Voted

 1 year, 4 months ago

Selected Answer: ACF

A to upload the files to S3 via SMB  
C to convert the data from CSV format  
F to query with SQL

Not B (we need the data in S3, not in FSx)  
Not D or E (we should provide the ability to run SQL queries)

upvoted 5 times
- TariqKipkemei**

Most Recent

 1 year, 5 months ago

Selected Answer: ACF

SMB + use SQL commands to query the data = Amazon S3 File Gateway mode + Amazon Athena

upvoted 1 times
- wsdasdasdqwda** 1 year, 6 months ago

<https://aws.amazon.com/storagegateway/file/s3/#:~:text=Amazon%20S3%20File%20Gateway%20provides,Amazon%20S3%20with%20local%20caching.>

"Amazon S3 File Gateway provides a seamless way to connect to the cloud in order to store application data files and backup images as durable objects in Amazon S3 cloud storage. Amazon S3 File Gateway offers SMB or NFS-based access to data in Amazon S3 with local caching"

=> SMB and NFS is supported in Amazon S3 File Gateway => ACF

upvoted 2 times
- iwannabeawsgod** 1 year, 6 months ago

Selected Answer: ACF

ACF 100% sure



upvoted 3 times

  **Ramdi1** 1 year, 7 months ago

**Selected Answer: ACF**

I thought the correct answer was BCF however I have changed my mind to BCF  
FSx does support SMB protocol. However so does s3 file gateway which is version 2 and 3 of the SMB protocol. Hence using it with athena ACF should be correct

upvoted 4 times

  **RDM10** 1 year, 7 months ago

SMB file share- is B incorrect?

upvoted 1 times

  **pentium75** 1 year, 4 months ago

Yes because "FSx File Gateway" uploads the files to FSx, but we need it in S3.

upvoted 2 times

  **Guru4Cloud** 1 year, 7 months ago

**Selected Answer: BCE**

BCF is the correct

upvoted 1 times

  **pentium75** 1 year, 4 months ago

No because FSx File Gateway (B) uploads it to FSx while we we need it in S3. S3 File Gateway provides access to S3 via SMB.

upvoted 3 times

  **Eminenza22** 1 year, 8 months ago

**Selected Answer: ACF**

<https://docs.aws.amazon.com/glue/latest/dg/aws-glue-programming-etl-format-csv-home.html>  
<https://aws.amazon.com/blogs/aws/amazon-athena-interactive-sql-queries-for-data-in-amazon-s3/>  
<https://aws.amazon.com/storagegateway/faqs/>

upvoted 3 times

  **anikety123** 1 year, 8 months ago

**Selected Answer: ACF**

It should be ACF

upvoted 2 times

  **ralfj** 1 year, 8 months ago

**Selected Answer: ACF**

ACF use S3 File Gateway, Use Glue and Use Athena

upvoted 3 times

A company wants to use Amazon Elastic Container Service (Amazon ECS) clusters and Amazon RDS DB instances to build and run a payment processing application. The company will run the application in its on-premises data center for compliance purposes.

A solutions architect wants to use AWS Outposts as part of the solution. The solutions architect is working with the company's operational team to build the application.

Which activities are the responsibility of the company's operational team? (Choose three.)

- A. Providing resilient power and network connectivity to the Outposts racks
- B. Managing the virtualization hypervisor, storage systems, and the AWS services that run on Outposts
- C. Physical security and access controls of the data center environment
- D. Availability of the Outposts infrastructure including the power supplies, servers, and networking equipment within the Outposts racks
- E. Physical maintenance of Outposts components
- F. Providing extra capacity for Amazon ECS clusters to mitigate server failures and maintenance events

**Correct Answer:** ACF

*Community vote distribution*

ACF (55%)ACE (27%)Other

**taustin2**

Highly Voted

1 year, 6 months ago

Selected Answer: ACF

From <https://docs.aws.amazon.com/whitepapers/latest/aws-outposts-high-availability-design/aws-outposts-high-availability-design.html>

With Outposts, you are responsible for providing resilient power and network connectivity to the Outpost racks to meet your availability requirements for workloads running on Outposts. You are responsible for the physical security and access controls of the data center environment. You must provide sufficient power, space, and cooling to keep the Outpost operational and network connections to connect the Outpost back to the Region. Since Outpost capacity is finite and determined by the size and number of racks AWS installs at your site, you must decide how much EC2, EBS, and S3 on Outposts capacity you need to run your initial workloads, accommodate future growth, and to provide extra capacity to mitigate server failures and maintenance events.

upvoted 27 times

**ibu007**

Highly Voted

1 year, 8 months ago

Selected Answer: ACE

My exam is tomorrow. thank you all for the answers and links.

upvoted 16 times

**pentium75**

1 year, 4 months ago

"Physical maintenance" such as replacing faulty disks is NOT your responsibility.

upvoted 7 times

**FlyingHawk**

Most Recent

3 months, 1 week ago

Selected Answer: ACF

this is a tricky question. I selected ACE initially, but after reading the AWS doc:  
<https://docs.aws.amazon.com/whitepapers/latest/aws-outposts-high-availability-design/aws-outposts-high-availability-design.html>  
Based on it , capacity is customer's responsibility. E is AWS responsibility

upvoted 1 times

**FlyingHawk**

3 months, 1 week ago

Since Outpost capacity is finite and determined by the size and number of racks AWS installs at your site, you must decide how much EC2, EBS, and S3 on Outposts capacity you need to run your initial workloads, accommodate future growth, and to provide extra capacity to mitigate server failures and maintenance events.(F)  
AWS is responsible for the availability of the Outposts infrastructure including the power supplies, servers, and networking equipment within the AWS Outposts racks. (D, E) AWS also manages the virtualization hypervisor, storage systems, and the AWS services that run on Outposts.(B)

upvoted 1 times

**[Removed]**

8 months, 3 weeks ago

Selected Answer: ACE

The activities that are the responsibility of the company's operational team when using Amazon Elastic Container Service (Amazon ECS) clusters and Amazon RDS DB instances on AWS Outposts are:

Providing resilient power and network connectivity to the Outposts racks.  
Physical security and access controls of the data center environment.

Physical maintenance of Outposts components.  
The solutions architect is responsible for the following activities:

Managing the virtualization hypervisor, storage systems, and the AWS services that run on Outposts.  
Ensuring the availability of the Outposts infrastructure, including the power supplies, servers, and networking equipment within the Outposts racks.  
Providing extra capacity for Amazon ECS clusters to mitigate server failures and maintenance events.  
upvoted 2 times

  **pentium75** 1 year, 4 months ago


**Selected Answer: ACF**

F: "If there is no additional capacity on the Outpost, the instance remains in the stopped state. The Outpost owner can try to free up used capacity or request additional capacity for the Outpost so that the migration can complete."

Not D: "Equipment within the Outposts rack" is AWS' responsibility, you're not supposed to touch that

Not E: "When the AWS installation team arrives on site, they will replace the unhealthy hosts, switches, or rack elements"


upvoted 6 times

  **1rob** 1 year, 5 months ago

**Selected Answer: ACF**

From <<https://aws.amazon.com/outposts/rack/faqs/> : Your site must support the basic power, networking and space requirements to host an Outpost ==> A  
From <<https://docs.aws.amazon.com/whitepapers/latest/applying-security-practices-to-network-workload-for-csps/the-shared-responsibility-model.html> : In AWS Outposts, the customer takes the responsibility of securing the physical infrastructure to host the AWS Outposts equipment in their own data centers. ==> C

upvoted 2 times

  **1rob** 1 year, 5 months ago

and From <<https://docs.aws.amazon.com/whitepapers/latest/aws-outposts-high-availability-design/aws-outposts-high-availability-design.html> : Since Outpost capacity is finite and determined by the size and number of racks AWS installs at your site, you must decide how much EC2, EBS, and S3 on Outposts capacity you need to run your initial workloads, accommodate future growth, and to provide extra capacity to mitigate server failures and maintenance events. ==> F

upvoted 3 times

  **1rob** 1 year, 5 months ago

From <<https://docs.aws.amazon.com/whitepapers/latest/aws-outposts-high-availability-design/aws-outposts-high-availability-design.html>> AWS is responsible for the availability of the Outposts infrastructure including the power supplies, servers, and networking equipment within the AWS Outposts racks. AWS also manages the virtualization hypervisor, storage systems, and the AWS services that run on Outposts. So The customer isn't so not D.

upvoted 2 times

  **TariqKipkemei** 1 year, 5 months ago

**Selected Answer: AC**

Only A and C are correct.  
AWS is responsible for the hardware and software that run on AWS Outposts. This is a fully managed infrastructure service. AWS manages security patches, updates firmware, and maintains the Outpost equipment. AWS also monitors the performance, health, and metrics for your Outpost and determines whether any maintenance is required.

<https://docs.aws.amazon.com/outposts/latest/userguide/outpost-maintenance.html>

upvoted 1 times

  **pentium75** 1 year, 4 months ago

"Choose three".

You're missing F, you must order the Outposts rack with excess capacity

upvoted 2 times

  **[Removed]** 1 year, 5 months ago

**Selected Answer: ACE**

The role that physical companies will play is ACE.

upvoted 1 times

  **potomac** 1 year, 5 months ago

**Selected Answer: ACD**

E is wrong  
If there is a need to perform physical maintenance, AWS will reach out to schedule a time to visit your site.  
<https://aws.amazon.com/outposts/rack/faqs/#:~:text=As%20AWS%20Outposts%20rack%20runs,the%20Outpost%20for%20compliance%20certification.>

upvoted 1 times

  **pentium75** 1 year, 4 months ago

So is D, "equipment WITHIN the Outposts rack" is something that your Infra team should stay away from.

upvoted 2 times

  **beast2091** 1 year, 6 months ago

ACE  
AWS is responsible for the availability of the Outposts infrastructure including the power supplies, servers, and networking equipment within the AWS Outposts racks. AWS also manages the virtualization hypervisor, storage systems, and the AWS services that run on Outposts.

<https://d1.awsstatic.com/whitepapers/aws-outposts-high-availability-design-and-architecture-considerations.pdf>  
upvoted 1 times

  **dilaaziz** 1 year, 6 months ago

**Selected Answer: ACF**

<https://docs.aws.amazon.com/whitepapers/latest/aws-outposts-high-availability-design/aws-outposts-high-availability-design.html>  
upvoted 2 times

  **canonlycontainletters1** 1 year, 6 months ago

**Selected Answer: ACD**

I choose ACD  
upvoted 1 times

  **danielmakita** 1 year, 6 months ago

**Selected Answer: ACD**

I think ACD is correct  
upvoted 1 times

  **chris0975** 1 year, 6 months ago

**Selected Answer: ACF**

You get to choose the capacity. F  
upvoted 2 times

  **thanhnv142** 1 year, 6 months ago

A, C and D  
upvoted 1 times

  **aleksand41** 1 year, 6 months ago

ACD <https://docs.aws.amazon.com/outposts/latest/userguide/outpost-maintenance.html>  
upvoted 1 times

  **Ramdi1** 1 year, 7 months ago

**Selected Answer: ACD**

I think because of the shared responsibility model it is ACD  
upvoted 3 times

A company is planning to migrate a TCP-based application into the company's VPC. The application is publicly accessible on a nonstandard TCP port through a hardware appliance in the company's data center. This public endpoint can process up to 3 million requests per second with low latency. The company requires the same level of performance for the new public endpoint in AWS.



What should a solutions architect recommend to meet this requirement?


- A. Deploy a Network Load Balancer (NLB). Configure the NLB to be publicly accessible over the TCP port that the application requires.
- B. Deploy an Application Load Balancer (ALB). Configure the ALB to be publicly accessible over the TCP port that the application requires.
- C. Deploy an Amazon CloudFront distribution that listens on the TCP port that the application requires. Use an Application Load Balancer as the origin.
- D. Deploy an Amazon API Gateway API that is configured with the TCP port that the application requires. Configure AWS Lambda functions with provisioned concurrency to process the requests.

**Correct Answer: A**

*Community vote distribution*

A (100%)

  **Sugarbear\_01**

Highly Voted 

 1 year, 1 month ago

**Selected Answer: A**

Since the company requires the same level of performance for the new public endpoint in AWS.

A Network Load Balancer functions at the fourth layer of the Open Systems Interconnection (OSI) model. It can handle millions of requests per second. After the load balancer receives a connection request, it selects a target from the target group for the default rule. It attempts to open a TCP connection to the selected target on the port specified in the listener configuration.

Link;  
<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html>  
upvoted 11 times



  **TariqKipkemei**


Highly Voted 

 11 months, 1 week ago

**Selected Answer: A**

TCP = NLB  
upvoted 6 times



  **awsgeek75**

Most Recent 

 9 months, 3 weeks ago

**Selected Answer: A**

B: Is wrong as ALB is not going to help with TCP traffic  
C: CloudFront is CDN. There is no content here  
D: API Gateway is for HTTP web/API stuff, not custom TCP port applicationns  
upvoted 3 times

  **taustin2** 1 year, 1 month ago

**Selected Answer: A**

NLBs handle millions of requests per second. NLBs can handle general TCP traffic.  
upvoted 4 times